Hate Crimes and Student Outcomes*

Xuan Li † Ying Shi ‡ Hantao Wu §

Job Market Paper

This version: December 2025

Please click here for the most recent version

Abstract: This paper examines the causal impact of exposure to racially targeted hate crimes on students' academic and behavioral outcomes. We construct a dataset spanning nearly one and a half decades that links geocoded hate crime records from Los Angeles County to administrative data from the Los Angeles Unified School District. We estimate the effects of school- and neighborhood-level exposure using two complementary difference-in-differences designs: one comparing targeted students in treated schools with peers in similar control schools, and another exploiting hyperlocal variation in students' proximity to hate crimes within the same neighborhood. Exposure to anti-Hispanic hate crimes leads to persistent GPA declines among Hispanic students—by 0.07 points for school-based and 0.06 for neighborhood-based exposure. It also temporarily increases disciplinary referrals and lowers students' perceived safety. We find no significant effects of anti-Black hate crime exposure on Black students, and furthermore show no effects of hate crime exposure on non-targeted groups.

*We would like to thank Desmond Ang, Leah Boustan, Ellora Derenoncourt, Eric Edmonds, Raymond Fisman, Thomas Fujiwara, Joshua Goodman, Kevin Lang, Benjamin Marx, Dilip Mookherjee, Daniele Paserman for helpful comments. We acknowledge funding support provided by the Abdala Fieldwork Research Grant, Arnold Ventures Planning Grant, CISS Summer Mini-Research Grant, IED Student Research Grant, and Overdeck Education Innovation Fund. All errors are our own.

†Boston University (email: xuanli@bu.edu)

‡Syracuse University (email: yshi78@syr.edu)

§Princeton University (email: hantao.wu@princeton.edu)

1 Introduction

How does identity-based violence affect children and youths? A particularly stark manifestation is hate crimes — criminal offenses motivated by bias against a victim's race, ethnicity, religion, or other protected characteristics. Unlike other forms of violence, hate crimes target individuals on the basis of group identity and their impact can extend beyond direct victims to entire communities (Perry and Alvi, 2012). Exposure is increasingly common: one in four students report seeing hate symbols in schools, an increasingly common site of exposure alongside homes and public spaces (U.S. Government Accountability Office, 2021). Such exposure is likely to shape learning: hostile environments can undermine concentration and elevate stress (Wenger, Lantz and Gallardo, 2022; Hellyer and Gereke, 2024). Yet despite the prevalence of these experiences, systematic evidence on their academic and socio-emotional consequences remains limited.

This paper examines the impact of racially-motivated hate crimes on student outcomes.[1] The analysis draws on a newly linked dataset combining geocoded hate crime incidents from the Los Angeles County Commission on Human Relations with administrative student records from the Los Angeles Unified School District (2008–2024). We separately examine hate crimes occurring in students' schools and in their residential neighborhoods. For school-based exposure, we estimate a stacked difference-in-differences model comparing students in affected schools with peers in nearby control schools. For neighborhood-based exposure, we exploit hyperlocal variation in the timing and location of hate crimes, comparing students residing very close to incidents with those slightly farther away.

We find that exposure to racially-motivated hate crimes in schools can have adverse academic and behavioral consequences for targeted groups. Following anti-Hispanic hate crimes, Hispanic students' GPAs decline by an average of -0.074 points, with these negative effects persisting for up to four semesters after exposure. Non-targeted peers show no comparable effects. The lack of spillover effects on other racial groups is consistent with the racial identity-specific nature of the threat. In addition, after experiencing a hate crime on campus, Hispanic

[1] In 2017-18, nearly four-fifth of all documented hate crimes in the survey targeted a student's race or color (U.S. Government Accountability Office, 2021). The GAO report uses data from the School Crime Supplement of the National Crime Victimization Survey. Hate crime data from the FBI's Uniform Crime Reporting program (UCR) also show that racial or ethnic bias motivate the largest share of incidents (Federal Bureau of Investigation, 2024).

students exhibit a relative uptick in office referrals and a decrease on perceived safety on school grounds. These effects, however, dissipate after two to three semesters. Finally, after exposure, Hispanic students are more likely to be classified as "emotionally disturbed", a special education category defined by attributes such as the inability to learn, relationship problems, and unhappiness or depression. Notably, we do not detect comparable GPA or other behavioral, social or emotional changes for Black students who experience anti-Black hate at their schools.

The impact of neighborhood hate crime exposure is strikingly consistent with the school-based analysis. Hispanic students' GPA decrease following exposure to a nearby serious hate crime. The effect magnitude is slightly attenuated compared to the school analysis, but endures for several semesters in a similar fashion. There is a short-term decrease in perceived safety alongside increased likelihood of classification as "emotionally disturbed" that is apparent during the year of exposure, but not thereafter. As before, we do not observe similar impacts for Black students who experience anti-Black hate in their respective neighborhoods.

A key concern is whether our findings reflect bias-motivated violence specifically or exposure to violence more generally. To address this, we use incident-level data from the Los Angeles Police Department (LAPD) and the Los Angeles County Sheriff's Department (LASD). Focusing on non-hate crimes with minority victims, we find no comparable impacts, indicating that the adverse consequences we document stem from the identity-targeted threat unique to hate crimes.

Our paper builds on several strands of literature. The main contribution is to estimate the causal impact of hate crime exposure on schooling outcomes. Few studies provide evidence on the economic and social costs of hate crimes by linking its exposure to human capital accumulation and economic behavior. Existing studies have primarily relied on FBI hate-crime data to measure state- or county-level exposure and examine its association with various human capital outcomes, such as birth weight, college enrollment, and degree completion (Jaskiewicz, 2024; Baker and Britton, 2024; Petach and Rumbaugh, 2025). For example, Baker and Britton (2024) show that a one–standard-deviation increase in reported state-level hate crimes (approximately 114 incidents) is associated with a 17 to 22 percent rise in Black first-time student enrollment at historically Black colleges and universities (HBCUs). Increases in reported hate crime incidents also predict lower college completion for young Black individuals (Petach and

Rumbaugh, 2025). Instead of relying on the FBI UCR data and aggregating to commuting zone or state-level variation in hate crime exposure, our georeferenced records allow for more precise pinpointing of hate crime occurrence and hyperlocal variation in exposure. The link to administrative school records also facilitate an examination of effects on human capital accumulation for children and youths, which to our knowledge has not been studied in this context.

The paper also relates to a broader literature on the consequences of exposure to violence during childhood and adolescence. This literature typically exploits variation in violence across time and space to identify its causal impacts on academic achievement and later labor market outcomes. For instance, exposure to violent crime lowers test scores in Chicago (Sharkey, 2010), Colombia (Cristancho et al., 2024), and Mexico City (Chang and Padilla-Romo, 2023), with more pronounced decreases for reading compared to math in New York City and a large sample of U.S. school districts (Sharkey et al., 2014; Torrats-Espinosa, 2020). Similarly, exposure to other forms of violence–such as school shootings, police violence, sniper attacks, and terrorism–has been found to negatively affect educational achievement and attainment (Gershenson and Tekin, 2017; Ang, 2020; Bharadwaj et al., 2021; Shany, 2025; Cabral et al., 2025). In some cases the effects do not endure and fade out with temporal distance (Cristancho et al., 2024; Shany, 2025), while others lead to lasting changes in school quality and downstream postsecondary and labor market outcomes (Bharadwaj et al., 2021; Chang and Padilla-Romo, 2023; Cabral et al., 2025). Effects can reach far beyond academic achievement into mental health and general well-being. For example, surviving children in a Norwegian mass shooting exhibited increased utilization of health care services and more mental health diagnoses (Bharadwaj et al., 2021), and fatal school shootings increased youth antidepressant use (Rossin-Slater et al., 2020). Our results enables comparing the causal impact of hate crime exposure on a range of outcomes with that of other forms of violence.

In one view, hate crimes represent an extreme form of racial discrimination. In the classic economics of discrimination literature, racial disparities in the labor market are often decomposed into observable components—such as test scores and educational attainment—and the unexplained portion frequently ascribed to discrimination (Oaxaca, 1973; Blinder, 1973; Al-

tonji and Blank, 1999).[2] Some studies advance that the majority of racial differences in wages can be "explained" once an exhaustive set of human-capital measures is included—suggesting a relatively minor role for discrimination (Neal and Johnson, 1996; Heckman, 1998; Fryer Jr., 2011). Yet human capital formation can itself be endogenous to environments with discriminatory features (Thompson, 2025). Our findings support this dynamic view in which educational investments are impeded by exposure to racially motivated hate crimes, which underscores the importance of attenuating racial discrimination during the schooling years.

2 Definition of Hate Crimes

Hate crimes are criminal offenses that are motivated by bias against a victim's actual or perceived membership in a group, typically defined by race, ethnicity, religion, disability, sexual orientation, gender, or gender identity.[3] They represent a distinct category of crime in which the intended target is not only the victim, but also members belonging to the same group (National Institute of Justice, 2021). Because individuals are targeted for group identity rather than individual attributes, hate crime exposure can generate broader spillovers among others who share the targeted identity. There is some empirical support for elevated psychological distress and reduced well-being in affected communities, even if they are not individually targeted (Bell and Perry, 2015; Farrell and Lockwood, 2023).

A closely related but conceptually distinct notion is hate speech, which refers to any form of communication that employs pejorative or discriminatory language to threaten or demean groups based on their identity (United Nations, 2019). Hate speech constitutes a criminal offense only when it involves spoken or written threats of violence against individuals or groups; derogatory speech or epithets that do not involve such threats are not considered hate crimes (LA County Commission on Human Relations, 2022). Nevertheless, direct exposure to hate speech has been shown to elicit heightened fear, emotional distress, and physiological responses, reducing individuals' capacity to participate fully in society (Gelber and McNamara, 2016).

[2] A substantial empirical literature exists on racial inequality and discrimination in the labor market (see, for example, Lang and Lehmann (2012) and Bayer, Charles and Derenoncourt (2025)). While many studies account for differences in observable skills as described, we acknowledge a burgeoning and influential set of studies using audit and field experimental techniques to signal race. They provide compelling evidence for differential treatment by race in the labor market (Bertrand and Mullainathan, 2004; Bertrand and Duflo, 2017).

[3] This is the definition adopted by the Federal Bureau of Investigation for their data collection efforts. State statutes vary somewhat in their definitions.

We further distinguish hate crimes from other forms of student victimization, such as bullying and harassment. Bullying involves unwanted aggressive behavior that is either repeated or highly likely to be repeated (CDC, 2025). It can take multiple forms, including physical, psychological, social/relational, or online, with perceived or actual power imbalances being a distinguishing characteristic. Harassment refers to unwelcome conduct that can deny or limit a student's ability to participate in or benefit from a school's education program (U.S. Government Accountability Office, 2021). If either bullying or harassment incidents become motivated by bias against a protected characteristic and the act is sufficiently severe, they could also become hate crimes.

In the context of this project, hate crimes may occur on school property or in the neighborhoods surrounding students' homes, encompassing both residential and public spaces. Hate crimes that take place on school grounds may be perpetrated by peers, school staff, or individuals with no direct affiliation to the school.

Although the definition of a hate crime appears conceptually straightforward, systematically collecting reliable information on such incidents presents numerous practical challenges. The most commonly used dataset is the FBI's Uniform Crime Reporting (UCR) Program, yet this database is widely recognized to suffer from substantial underreporting and incomplete coverage (Kaplan, 2021). Moreover, it does not disclose the precise geographic locations of reported hate crimes, limiting its usefulness for fine-grained analyses. In Section B, we provide a detailed discussion of how the FBI classifies and collects hate-crime incidents and review several alternative data sources that have been used in the literature.

3 Data

Our analysis focuses on Los Angeles County—one of the most populous and ethnically diverse counties in the United States. According to the 2020 Census, the county is home to more than ten million residents, 48.0% of whom identify as Hispanic, 25.6% as White, 15.0% as Asian, and 7.9% as Black. The primary data sources for our analysis are two administrative datasets: an incident-level hate-crime dataset provided by the Los Angeles County Commission on Human Relations (LACCHR) and student administrative records from the Los Angeles Unified School District (LAUSD), the second-largest school district in the nation. We further supplement these data with administrative records on non–hate-crime incidents from the Los

Angeles Police Department (LAPD) and the Los Angeles County Sheriff's Department (LASD).

3.1 Hate crime data

Our hate crime data come from the Los Angeles County Commission on Human Relations (the Commission), which was established in the aftermath of the "Zoot Suit Riots" in 1944 and is one of the oldest and largest agencies of its kind in the United States. The Commission has maintained a database of hate crime data since 1980, which represents one of the longest-standing efforts to document hate crime incidents in the United States. While the FBI's Uniform Crime Reports compile data primarily from law enforcement agencies, the Commission collects information not only from local law enforcement but also from a much broader set of contributors—including more than 100 reporting entities such as K–12 schools, colleges and universities, and community-based organizations.[4] They screen each reported incident for conformity with state law and the definition of hate crimes as specified by the District Attorney office. California Penal Code § 422.55 stipulates that hate crime charges may filed for criminal acts committed, in whole or in part, based on the victim's actual or perceived race or ethnicity, disability, gender, nationality, religion, or sexual orientation. The Commission's classification of hate crimes depends on any bias, hatred, or prejudice based on perceived or actual protected groups representing a substantial factor in the commission of the offense. These biases or prejudices can be direct or circumstantial, and they can occur before, during, or after the offense (LA County Commission on Human Relations, 2001). Those that do not meet the evidentiary standard are excluded from statistical reporting. Reports from individuals and community organizations are cross-checked against cases from law enforcement agencies and any resulting duplicates are removed from the sample (LA County Commission on Human Relations, 2001). The consistency of the standards for inclusion in the hate crime dataset over at least two decades enables the comparability of records over time.[5]

The sourcing of hate-related incidents from a broader range of contributing agencies implies that the Commission's data should be more comprehensive than the FBI's Uniform Crime

[4]Since 2020, the Commission has also collected hate crime data directly via LA vs Hate, a platform that partners with community organizations enabling direct victim reports online or by email in a confidential manner. The platform distinguishes between hate crimes and hate incidents, with the latter encompassing hostile, bias-motivated expressions or actions that target actual or perceived identity(ies), but that does not rise to the level of a crime. The Commission has recently added a Hate Incident Report to capture this category of hate activity. Our dataset includes only hate crimes and does not include non-criminal hate incidents.

[5]More details on the Commission's classification of hate crimes are available in Appendix Section C, along with illustrative examples of hate crimes that occurred in school and neighborhood settings.

Reports for Los Angeles County. Figure A1 juxtaposes the number of hate crime incidents in 2008-2021 across these two data providers. Incident counts in the Commission data are consistently one-third to one-half more than the FBI data. Even with the improved coverage, we caution that the data are still likely subject to underreporting. However, for hate crimes that occur at schools, certain institutional reporting mechanisms help mitigate this concern. The Commission receives school-related hate crime data partly from local law enforcement departments and also directly from the Los Angeles Unified School District (LAUSD). The district maintains a mandatory reporting system requiring school administrators to document and escalate all suspected bias-motivated incidents. Given schools' highly supervised environments and trusted internal reporting channels, such incidents are more likely to be identified and recorded compared to hate crimes occurring elsewhere. To further mitigate concerns about under-reporting of hate crimes (Kena and Thompson, 2021), our neighborhood-based analysis focuses primarily on Part I offenses –serious crimes such as murder and aggravated assault–which, given their severity, are less prone to underreporting and more likely to have measurable community impacts.

[Insert Figure A1 here]

The hate crime dataset includes 6,040 incidents deemed by the Commission to meet the definition for this class of crimes from 2008-2021. Each record contains information on the exact date of offense, street address or intersection at which the crime occurred, reporting jurisdiction, bias motivation and targeted group, offense type, and information on the gender, race, and age of victims and suspects, where available.

[Insert Figure 1 here]

Panel A of Figure 1 maps the distribution of hate crimes across Los Angeles County between 2008 and 2021, aggregated at the ZIP code level. Hate crimes occur throughout the county but are relatively more concentrated in urbanized and demographically diverse areas—particularly in Central and West Los Angeles—whereas many high-crime neighborhoods in South and East Los Angeles report comparatively few hate-related incidents.

Panel B of Figure 1 plots the spatial distribution of overall crimes over the same period using administrative incident-level data from the Los Angeles County Sheriff's Department (LASD). General crimes cluster more strongly in South and East Los Angeles and in other economically disadvantaged neighborhoods, reflecting patterns of property and violent offenses.

While the two maps reveal a positive overall correlation—both types of incidents are more common in dense urban areas—their hotspots only partially overlap. This suggests that hate crimes are shaped not only by routine opportunity structures driving general crime but also by social diversity and intergroup exposure that vary across communities.

We focus on cases in which the reported bias motivation is based on race, ethnicity, or national origin, which together comprise more than 50 percent of all incidents in the dataset. Panel A of Figure 2 maps racially targeted hate crimes against Hispanic, Black, and Asian residents across Los Angeles County between 2008 and 2021. For comparison, Panel B uses 2020 Census data to illustrate the residential distribution of each racial group, measured as the population share within each ZIP code. Overall, incidents targeting Hispanic, Black, and Asian residents exhibit spatial clustering rather than random dispersion. Many of these incidents are concentrated in urbanized and demographically mixed neighborhoods where different groups interact frequently, rather than being confined to each group's own residential enclaves. Areas with larger Hispanic or Asian populations show somewhat higher levels of anti-Hispanic and anti-Asian hate crimes, but the relationship is far from proportional, and anti-Black hate crimes in particular are often reported outside predominantly Black neighborhoods. Taken together, these patterns suggest that intergroup exposure, rather than population size alone, plays an important role in the spatial distribution of racially motivated hate crimes.

Hate crimes that occur in schools likely differ in nature from those that take place elsewhere. Given our focus on the effects of hate crimes on students, Figure A1 compares the characteristics of school-based and non-school hate crimes in Los Angeles County, including incidents from all educational institutions rather than only those under LAUSD. Panel A shows differences in offense type. School-based hate crimes are overwhelmingly classified as Part II offenses (such as vandalism, simple assault, or intimidation), whereas non-school hate crimes include a larger share of Part I offenses (such as aggrevated assault or more serious violence). Vandalism is the most common form of hate crime on school grounds, accounting for more than half of all cases, while non-school hate crimes display a more even mix of vandalism, assault, and intimidation.

Although our main analysis focuses on racially targeted hate crimes, we provide descriptive statistics on the broader distribution of bias motivations for completeness. Panel B compares the targeted groups in school-based and non-school hate crimes. In both settings,

incidents motivated by race or ethnicity account for the majority of cases, but the specific patterns differ. School-based hate crimes disproportionately target Black students (35 percent of cases), whereas non-school incidents are more often directed against Hispanic or LGBTQ individuals. Notably, while Black victims constitute the largest share in both settings, Jewish students represent the second most frequently targeted group in school-based incidents—a pattern that differs from the broader population and warrants further investigation in future research. These contrasts suggest that school-based hate crimes are generally less violent but more racially focused, reflecting the social environment and patterns of interaction within schools.

3.2 Student administrative data

Our student records come from the Los Angeles Unified School District (LAUSD). We have data on all students from the 2007-2008 through 2023-2024 academic years, which contains information on student gender, year and month of birth, race and ethnicity, home language, special education status, English language learner status, economic disadvantage, parental education, and residential address.[6] During the sample period, LAUSD enrolled an average of 290,000 middle and high school students per year. Of these, 74.7% were Hispanic, 9.5% Black, 8.6% White, and 6.3% Asian. Moreover, more than 80% of students were classified as economically disadvantaged.

Our primary outcome is academic performance, measured by non-cumulative semester-level GPA. These data are available for all students in grades 6–12 and are recorded at a relatively high frequency.[7] Academic performance gaps across racial and ethnic groups are substantial: average GPA is 3.21 for Asian students, 3.00 for White students, 2.45 for Hispanic students, and 2.32 for Black students.

While academic measures are central to our analyses, we also examine a number of non-academic outcomes on student behavior and well-being. LAUSD provides data on all office referrals beginning from the 2013-2014 school year, regardless of whether a student received any

[6] A caveat is that the district only maintains the most recently reported address. This could potentially attenuate our estimates toward zero if student migration is substantial.

[7] While California administered standardized tests in math and reading throughout our sample period, the testing schemes changed over time, covered only certain grade levels, and were conducted annually. Students in grades 2–11 were assessed under the California Standards Tests (CST) between 2008 and 2013, after which the state transitioned to the Smarter Balanced Assessment (SBAC) system beginning in 2015, which tests only students in grades 3–8 and 11.

exclusionary discipline. This admits a broader range of behavioral patterns into the dataset and is less plagued by selection into a particular form of disciplinary consequence. Each incident is date-stamped and categorized by the final administrative decision. We use two additional sets of outcomes to gauge student behavioral and emotional well-being. One comes from special education data that documents whether each student ever met the criteria for inclusion under the 13 disability categories set forth by the Individuals with Disabilities Education Act (IDEA).[8] One of the categories is "Emotional Disturbance" (ED). A student can receive special education services under this category if their relational, psychological, and behavioral symptoms are deemed sufficiently severe, enduring, and adversely impact educational performance.[9] While this definition admits substantial discretion in diagnosis, there is a growing volume of research documenting a relationship between exposure to violence, trauma, and other adverse experiences and the manifestation of traumatic stress symptoms (Perfect et al., 2016). Students exhibiting emotional disabilities are at substantially elevated rates for previous exposure to adverse and traumatic experiences (Offerman et al., 2025). Finally, we make use of the district's School Experience Survey, administered annually to all students in grades 4-12.[10] One set of questions consistently ask about students' perceptions of safety at schools and in neighborhoods from 2014-2015 onwards, with responses measured on a Likert scale. We construct binary indicators for students agreeing to statements on feeling safe.

[8]Identification of students with special needs follows a multi-step process, commencing when a parent, teacher, other school personnel, or community member makes a written request for assessment to the LAUSD school principal. An assessment plan is then provided to the parent/guardian, who has 15 days to approve. From there, the district has 60 days to complete the assessment, which can include individual testing, child observations, records review, and interviews. After the assessments, the Individualized Education Program (IEP) team meets to determine eligibility for special education services.

[9]A student is eligible under emotional disturbance if they exhibit one or more of the following characteristics, over a long period of time and to a marked degree, which adversely affects educational performance:

1. An inability to learn which cannot be explained by intellectual, sensory, or health factors
2. An inability to build or maintain satisfactory interpersonal relationships with peers and teachers
3. Inappropriate types of behavior or feelings under normal circumstances exhibited in several situations
4. A general pervasive mood of unhappiness or depression; and
5. A tendency to develop physical symptoms or fears associated with personal or school problems.

[10]Survey administration has historically been in the fall, between late October and early December. The most recent year in our sample was administered in spring 2024 (February-March) rather than fall 2023.

3.3 Sample construction

We begin by identifying racially biased hate crimes that occurred on school campuses using the Commission's database of hate crimes, and then geocode each incident by latitude and longitude to link it to the LAUSD school directory. The LAUSD dataset contains only latitude and longitude coordinates (reported to three decimal places) for each school, omitting school names and street addresses to protect privacy. We then calculate the distance between the coordinates of each hate crime and those of LAUSD schools. Some hate crimes that occur in educational settings may take place in private schools, universities, or other types of institutions. To reduce the risk of misclassifying incidents that occurred at non-LAUSD schools, we treat an incident as a relevant school-based case if it falls within 0.1 miles of an LAUSD campus coordinate.[11] There are 91 school-based incidents, with 24/60/7 involving anti-Hispanic, anti-Black, and anti-Asian hate crimes.[12]

Among anti-Hispanic and anti-Black hate crimes in our sample schools, over 90% involve less serious offenses (i.e., Part II crimes), such as vandalism (43%), simple assault (29%), intimidation (11%), and disorderly conduct (10%). The more serious offenses are limited to aggravated assault, which accounts for 7% of cases. Overall, 74% of suspects are male, and 83% are likely students, as proxied by being under the age of 18. Among anti-Black hate crimes, 76% of suspects are Hispanic, while among anti-Hispanic hate crimes, 61% of suspects are Black. Not all hate crimes entail identifiable individual victims; vandalism, for instance, typically does not. When victim information is available, 48% are male, and 70% are likely students based on age.

Approximately 26% of schools in our sample experienced more than one hate crime during the study period. To balance the number of observed events and to limit the influence of potential cumulative effects, we restrict the analysis to incidents separated by at least one year. In cases where multiple hate crimes occurred within a single year at the same school, only the earliest incident is retained.[13]

[11] The mean distance is 0.06 miles.

[12] Because of the relatively small number of anti-Asian hate crimes in our sample and the relatively small share of Asian students in schools, we lack sufficient statistical power to conduct meaningful analyses. Therefore, our analysis focuses primarily on anti-Hispanic and anti-Black hate crimes.

[13] In our robustness checks, we also consider two alternative restrictions: keeping only the earliest-occurring incident for each school, and focusing exclusively on schools with a single hate crime. The results are similar.

4 Within-School Exposure to Hate Crimes

4.1 Empirical strategy

4.1.1 Control schools

Our objective is to estimate the causal effect of hate crime exposure on student outcomes. A central empirical challenge is that schools or neighborhoods in which hate crimes occur may differ systematically from those without such incidents. To address this concern, the school-level analysis employs a stacked difference-in-differences design with student fixed effects, comparing changes in outcomes for students in schools that experienced a hate crime with those for students in nearby, otherwise similar schools that did not, before and after the incident. To ensure comparability between treated and control schools, we first restrict the set of potential control schools to institutions serving the same grade span, thereby avoiding comparisons across fundamentally different student bodies (e.g., middle versus high schools). We then further limit the sample to schools located within the same neighborhood, defined by zip code, so that they are exposed to similar local environments, with the key distinction being whether a hate crime occurred at a given point in time.

Table A1 juxtaposes baseline attributes across the analytic sample of 192 treated and control schools. Column (3) shows differences in means while column (4) reports associated p-values. While these are unconditional values, we then evaluate for any significant differences when conditioning on zip code fixed effects to ensure comparability within smaller geographic units (columns (5)-(6)).

The balance tests show that observable characteristics appear similarly distributed across treatment and control schools, as we observe no significant differences across all covariates. In addition to pairwise t-tests, the associated F-statistic shows no evidence that the covariates jointly have a significant impact on treatment assignment.

4.1.2 Empirical specification

We employ a stacked difference-in-differences design to estimate the causal effect of hate crime exposure, comparing within-student changes in outcomes between treated and control schools before and after the incident. To implement this approach, we primarily rely on an event study specification that traces out the impact of hate crimes over time. Below we detail the model as well as particular specification choices:

$$Y_{is(n)t} = \sum_{t \in T, t \neq -1} \beta_t HCSchool_{st} + \alpha_i + \theta_{nt} + \epsilon_{it} \quad (1)$$

$Y_{is(n)t}$ is the outcome for student i enrolled in school s, during time period t (e.g. academic year or semester). $HCSchool_{st}$ takes on a value of 1 for schools in neighborhood n that experienced a racially-motivated hate crime starting from the period of exposure.[14] We include individual fixed effects α_i, which identify effects from within-student changes over time by comparing each student's outcomes before and after exposure to a hate crime. Zip code-by-period fixed effects θ_{nt} account for shocks common to students in each granular geographic unit and can flexibly accommodate time trends specific to each semester. Standard errors are clustered at the school level.

The coefficient of interest is β_t, which captures the relative period-by-period differences in within-student changes among those enrolled in schools experiencing a hate crime, compared to those enrolled at control schools. Identification hinges on the plausibility of parallel trends. In this instance, the assumption is that outcomes would have evolved similarly across these two groups of students in the absence of any hate crime exposure. We carefully examine pre-trends to assess the plausibility of this assumption. We also conduct a number of robustness checks in further support of validity.

It is important to note that when implementing this strategy, we estimate the models separately for each targeted group (e.g., anti-Hispanic and anti-Black hate crimes), restricting the sample in each case to the corresponding racial or ethnic group. We estimate the effects of each category of hate crimes separately by targeted group because the exposure are inherently group-specific. For instance, anti-Black hate crimes are expected to directly affect Black students, while anti-Hispanic incidents may primarily influence Hispanic students. Pooling across groups could conflate these distinct exposures and obscure heterogeneous responses. We then supplement these analyses by examining outcomes for non-targeted students to assess potential spillover effects, such as whether hate crimes targeting one group influence the broader student population.

[14] We index time by t, which is centered on period 0 corresponding to the semester of the hate crime occurrence; negative and positive values of t indicate semesters before and after the incident, respectively.

4.2 Effects on GPA

We begin by estimating the effects of hate crime exposure on academic outcomes, focusing on non-cumulative semester GPA. We estimate equation 1 separately for students of different racial and ethnic groups, and display the coefficients corresponding to β_t in Figure 3.

[Insert Figure 3 here]

Pre-period treatment coefficients for the anti-Hispanic hate crime sample cluster tightly around 0, and accordingly there is no evidence of significant differences across treatment and control groups leading up to a hate crime occurrence. After hate crime exposure, Hispanic students in treated schools show a relative decline in GPA. The effects are relatively modest the semester of the crime, but persist until the GPA decrease reaches 0.1 points four semesters later. In contrast, Black students exposed to hate crimes in their schools show a GPA trajectory that is statistically indistinguishable from the semester immediately prior to the event, relative to peers in control schools.[15]

To assess whether the observed effects are concentrated among the targeted group, we next estimate the impacts of anti-Hispanic and anti-Black hate crimes on non-targeted students. If exposure operates through identity-specific perceived threats, effects should primarily arise among members of the targeted group. Figure 4 presents the estimated effects for non-targeted students. The treatment coefficients are statistically indistinguishable from zero across all post-treatment periods, suggesting that hate crime exposure does not have measurable effects on students who do not belong to the targeted group.

[Insert Figure 4 here]

To contextualize the magnitude of the GPA effects, the decline grows to 0.10 points by $t = 2$. This is comparable in scale to the 0.07–0.08 point decline observed by Ang (2020) following exposure to police killings near students' homes in LAUSD. Although not directly comparable, our results are of a similar order of magnitude to prior evidence linking exposure to gun violence and other adverse events to reduced academic achievement (Sharkey et al., 2014; Gershenson and Tekin, 2017). Yet, unlike these broader shocks, the effects we document are group-specific, arising only among students belonging to the targeted identity group.

[15] We note that while coefficients for the two pre-periods t=-2 and t=-3 in the Black student sample are statistically indistinguishable from zero, there is a slight pretrend for the period t=-4. Even so, coefficients for the pre-treatment periods are jointly insignificant.

We undertake a number of additional analyses to ensure that our results are robust to specification choices. Table 1 presents coefficients corresponding to the interaction of $HCSchool$ denoting treated schools with an indicator that collapses post-exposure periods ($Post$). Column (1) estimates the effect of hate crimes on student GPA, separately for crimes motivated by anti-Hispanic and anti-Black bias. The effect is -0.074 points for Hispanic students and statistically distinguishable from zero. In contrast, we do not observe a discernible shift in Black students' GPAs when exposed to anti-Black hate crimes.

[Insert Table 1 here]

Column (2) augments the baseline specification by adding school-by-grade fixed effects, which control for grading standards and other school-level practices that remain stable over the study period. We also test for robustness when making different sample restrictions. Column (3) focuses on the first hate-crime incident in each school to ensure that results are not driven by repeated exposures. Column (4) excludes schools in which we observe multiple treatments, which comprise approximately one-quarter of the sample. The resulting coefficient for the Hispanic sample attenuates somewhat to -0.041, suggesting that effects may be pronounced in schools with multiple anti-Hispanic hate crimes and that impacts could vary by treatment dosage or accumulated exposure. The next specification excludes transfer students, finding that the results are largely unchanged. Column (6) includes only Part II hate crimes, ensuring that our results are not driven by the small subset of more severe incidents. Column (7) focuses on hate crimes occurring no later than 2019, since COVID-19 could have altered pedagogical modalities alongside other academic and social factors that could concurrently influence GPA. Doing so yields a coefficient of -0.078 for the Hispanic sample, which is close to the original point estimate.

Finally, we need to verify that our setting of a difference-in-differences model under staggered treatment timing is not introducing biases via the inclusion of comparisons between already-treated and yet-to-be-treated schools. We implement event study weights following Goodman-Bacon (2021). The resulting reduction of -0.159 points shows a larger treatment effect than our baseline specification, suggesting that we could be underestimating the adverse impact of hate crime exposure among Hispanic students.

4.2.1 Effect heterogeneity

The next set of analyses examines whether the impact of hate crimes is uniform across individuals or varies along key demographic and contextual dimensions. For instance, students from economically disadvantaged backgrounds may experience differential impacts compared with their less economically vulnerable peers. However, the effects may also be uniform across groups if hate crimes function as a pervasive identity-based threat. Panel A of Figure 5 explores this possibility by studying the effects of anti-Hispanic hate crime exposure on targeted Hispanic students' GPA, split by gender, grade level, home language, parental education, and economic disadvantage. There are no apparent differences in responses among male and female Hispanic students, or across students in middle vs. high school grades. Point estimates are also similar across students whose home language is English versus another language, those from different family backgrounds as measured by economic disadvantage, and by prior academic performance. Lastly, we investigate whether prior exposure to hate crimes shapes students' responses. Students in schools located in neighborhoods that experienced severe anti-Hispanic hate crimes in the preceding year display larger estimated effects, although overlapping confidence intervals indicate that we cannot statistically rule out similar impacts across groups.

[Insert Figure 5 here]

Alternatively, a student may be more impacted if they share socio-demographic attributes with the victim or suspected perpetrator. Panel B examines heterogeneity by the gender and age of victims and suspects. We deem someone as a student victim or suspect if they are of school-going age (below 18). Point estimates are similar across both victim and suspect gender groups. The effect is -0.05 points when the victim is under 18, compared to -0.02 (and statistically indistinguishable from zero) for older victims. More heterogeneity is apparent across suspect age bins. The magnitude of GPA decline when the hate crime is perpetuated by younger suspects is -0.20 points compared to -0.06 points for older suspects, although the difference is still imprecisely measured.

4.3 Effects on discipline, perceived safety, and special education classification

The next set of results go beyond academic indicators to focus on students' behavioral, social, and emotional responses. These dimensions can mediate the causal effect of hate crime

exposure on academic achievement, but they also independently merit attention given that they could influence human capital and labor market outcomes and are of economic importance in their own right. We begin with the frequency of disciplinary incidents. Hispanic students are involved in 0.1 additional office referrals during the semester in which they were exposed to an anti-Hispanic hate crime in school (Figure 6). This effect persists for two additional semesters, and then converges back to the same within-student changes as Hispanic peers in non-affected schools. As before, we observe no significant coefficients on pre-periods, so the divergence in disciplinary behavior only begins after the semester of exposure. The uptick in the number of disciplinary incidents is notable but not unprecedented, as previous research has documented an increase in misbehavior even from indirect exposure to peers facing adverse circumstances such as domestic violence at home (Carrell and Hoekstra, 2010).

[Insert Figure 6 here]

At the same time that exposed Hispanic students are increasing their involvement in disciplinary incidents, they are also reporting a decline in perceived safety in school. For this outcome, we rely on the annually administered School Experience Survey, which asks students in grades 4-12 about the extent to which they agree with feeling safe at their school. Figure 7 shows that the decrease in the likelihood of feeling safe begins in the year of exposure, and is at its highest point one year later. The adverse impact on perceived safety dissipates two years after the initial onset.

[Insert Figure 7 here]

Finally, we examine the effects of hate crime exposure on the propensity of being classified under the "emotional disturbance" category defined under IDEA. Classification of this condition is based on five core characteristics: inability to learn, relationship problems, inappropriate behavior, unhappiness or depression, and/or physical symptoms or fears. Figure 8 shows that exposure does not meaningfully change the rate of classification in the year of exposure, but the likelihood of being labeled as "emotionally disturbed" increases steadily in the three ensuing years up to 0.002 percentage points. The lack of effects during the year of exposure may reflect the relatively long lead-up to special education classification, usually involving a parent or staff member who initiates the assessment request and months before the team can make a determination on eligiblity. As a point of comparison, these effects are about three to four times the magnitudes of those reported in the case of exposure to police killings

(Ang, 2020).

[Insert Figure 8 here]

5 Neighborhood Exposure to Hate Crimes

5.1 Empirical Strategy

Establishing causal effects of neighborhood hate crime exposure on student outcomes presents significant empirical challenges. Neighborhoods experiencing hate crimes may differ systematically from those that do not, reflecting underlying socioeconomic conditions, residential sorting patterns, or pre-existing trends in student well-being. To address these identification concerns, we exploit hyperlocal variation in exposure to hate crimes within neighborhoods, following the approach developed by Ang (2020). Our analysis compares changes in outcomes over time among students who reside very close to hate crime locations with those who live slightly further away yet still within the same neighborhood. Both groups are therefore likely to share similar socioeconomic backgrounds and experience comparable local conditions, with the key difference being their exposure to nearby hate crime incidents.

This identification strategy builds on the observation that most hate crimes in our sample attract little media attention, suggesting that their impacts are likely to be highly localized. The localized nature of these incidents allows us to exploit fine-grained spatial variation to identify the causal effect of exposure. Moreover, because hate crimes target specific racial and ethnic groups, we estimate effects separately for each targeted group and then examine outcomes for non-targeted students in the same areas. In doing so, we can assess potential spillovers and determine whether the observed effects are driven by broader local factors rather than group-specific exposure.

Our analytical sample includes 438 serious racially motivated hate crimes classified as Part I offenses that occurred between 2008 and 2021. Part I offenses include more salient crimes such as aggravated assault and murder. These incidents are less likely to be underreported and are expected to have stronger impacts.[16] Among these incidents, 136 targeted Hispanic individuals and 302 targeted Black individuals. For each hate crime, we identify all students

[16]Part I crimes constitute 7.2 percent of all racially motivated hate crimes. Including Part II offenses would exceed our computational feasibility. Focusing on severe offenses also complements our school-level analysis, where most incidents are classified as Part II crimes.

residing within 3 miles of the incident location using their home addresses. We observe these students across multiple semesters both before and after exposure, tracking their academic trajectories, behavioral patterns, and psychological well-being. When students experience multiple hate crime exposures during the sample period, we retain only the first incident to ensure clean identification of treatment effects.

5.1.1 Graphical Evidence

We begin our empirical analysis by documenting the spatial gradient of hate crime effects using graphical evidence. If neighborhood exposure to hate crimes affects student outcomes, we would expect the magnitude of the effects to decline with distance from the incident location, even though the precise spatial range of influence is not known a priori. To examine this spatial heterogeneity, we estimate a difference-in-differences specification that disaggregates treatment effects across distance bands.

$$Y_{i(n)t} = \boxed{\sum_{d} \beta_d \cdot \text{Post}_t \cdot \text{Dist}_{i(d)}} + \varphi_i + \zeta_{nt} + \epsilon_{int} \tag{2}$$

where $Y_{i(n)t}$ represents the outcome for student i, residing in neighborhood n, during semester t. The term $\sum_d \text{Dist}_{i(d)}$ represents a set of mutually exclusive indicators capturing the distance from each student's home to the hate crime location, measured in 0.1-mile intervals. The variable Post_t denotes semesters following hate crime occurrence. The sample includes only students from the targeted racial or ethnic group.

The coefficients of interest, β_d, measure the change in outcomes for students at distance d from the hate crime in the post-exposure period, relative to students residing beyond 1.5 miles. This specification provides a nonparametric approach to tracing the spatial decay of hate crime effects. We include student fixed effects, φ_i, to control for time-invariant characteristics and zip code-by-semester fixed effects, ζ_{nt}, to ensure comparisons occur among students within the same local area experiencing common temporal shocks. Standard errors are clustered at the zip code level.

[Insert Figure 9 here]

Figure 9 presents estimated coefficients β_d plotted against distance from hate crime locations, separately for anti-Hispanic and anti-Black hate crimes. Each point estimate corre-

sponds to a 0.1-mile distance bin. For Hispanic students exposed to anti-Hispanic hate crimes, we observe significant negative effects on GPA concentrated among those residing within 0.1 miles of incident locations. The point estimate indicates a GPA decline of approximately 0.07 points. Coefficients quickly converge toward zero and become statistically insignificant beyond 0.1 miles, indicating that the impacts of hate crimes are highly localized. In comparison, Ang (2020) document effects of police violence extending up to 0.5 miles, possibly reflecting the greater salience and wider dissemination of information surrounding police-related incidents.

In contrast, Black students exposed to anti-Black hate crimes exhibit more variable patterns across the distance gradient. The attenuation is less systematic, and no distance bins achieve conventional statistical significance, consistent with our school-level findings of muted academic responses among Black students.

These spatial patterns yield several observations. First, the sharp distance decay for Hispanic students confirms that hate crime effects operate through highly localized mechanisms rather than diffusing broadly across neighborhoods. Second, this pattern provides data-driven guidance for defining treatment proximity and reinforces confidence in our identification, as students just beyond the exposure radius live nearby and therefore share comparable neighborhood and socioeconomic characteristics with treated students. Moreover, this data-driven approach to defining treatment proximity avoids imposing arbitrary distance thresholds. By first estimating the full spatial gradient nonparametrically, we allow the data to reveal the geographic scale at which hate crimes exert meaningful impacts.[17]

5.1.2 Estimating Specification

Having established through graphical evidence that hate crime effects concentrate within 0.1 miles of incident locations, we proceed to our main empirical specification. Based on the spatial gradient documented in Figure 9, we define the treatment group as students residing within 0.1 miles of hate crime locations—the distance band exhibiting the largest effects. The comparison group comprises students from the same targeted racial group living 0.1 to 3 miles away, sufficiently distant to avoid direct exposure yet proximate enough to share neighbor-

[17] The specific location of hate crime incidents may differ from victims' residential addresses, which are unobserved in our data. To the extent that hate crimes exert greater influence within victims' home communities, the estimated effects should be interpreted as a lower bound of the true impact.

hood characteristics.

We estimate the following event study specification to trace out the dynamic effects of hate crime exposure:

$$Y_{i(n)t} = \sum_{t \in T, t \neq -1} \beta_t \cdot \text{Nearby}_i \times \text{Post}_t + \varphi_i + \zeta_{nt} + \epsilon_{int} \tag{3}$$

where $Y_{i(n)t}$ represents the outcome for student i enrolled in school s, residing in neighborhood n, during semester t. The indicator Nearby_i equals one if the student's home address lies within 0.1 miles of a hate crime location, and zero if the student resides 0.1 to 3 miles away. The variable Post_t denotes relative time periods, where t indexes semesters before and after hate crime occurrence, normalized such that $t = 0$ corresponds to the semester of the crime. We omit the period immediately preceding the crime ($t = -1$) as the reference category.

The coefficients of interest, β_t, capture the differential change in outcomes for nearby students relative to comparison students in each period, measured against the baseline difference in period $t = -1$. This event study framework serves two purposes. First, it assesses parallel trends by examining whether pre-treatment coefficients cluster around zero, providing evidence that nearby and comparison students would have followed similar trajectories absent hate crime exposure. Second, it captures the temporal dynamics of treatment effects, allowing us to distinguish immediate reactions from longer-term impacts.

The specification includes student fixed effects (φ_i) to absorb time-invariant characteristics that may correlate with both residential location and outcomes, including family background, baseline ability, and unobserved preferences. Zip code-by-semester fixed effects (ζ_{nt}) flexibly control for time-varying neighborhood-level shocks common to all students in the same area, such as changes in local economic conditions, school quality, or community resources. This combination ensures that identification relies on within-student variation over time, comparing each student's outcomes before and after hate crime exposure while netting out aggregate trends affecting their neighborhood. Standard errors are clustered at the zip code level.

We implement this specification using a stacked event study design to address potential bias from treatment effect heterogeneity (Goodman-Bacon, 2021). For each hate crime incident, we construct a separate dataset comprising all students from the targeted racial group

living within 3 miles who had not previously experienced a nearby hate crime. We then stack these incident-specific datasets and estimate Equation 3 on the pooled sample, assigning equal weight to each event. This ensures treated students (within 0.1 miles) are compared only to never-treated students within the same event cohort, avoiding problematic comparisons to already-treated or differently-timed units. The stacked design is particularly important in our setting, as students may experience multiple hate crimes over the study period, with timing varying substantially across incidents.

Our identification strategy assumes students living very close to hate crimes (within 0.1 miles) and those living slightly farther away (0.1 to 3 miles) would have experienced parallel trends absent the crimes. Several empirical patterns support this assumption. First, the spatial decay is sharp and localized—students beyond 0.1 miles exhibit no discernible response, suggesting distance affects exposure intensity rather than reflecting pre-existing trajectory differences. Second, we formally test for differential pre-trends by examining whether β_t coefficients for $t < -1$ differ statistically from zero. Third, we assess balance in observable student characteristics across treatment and comparison groups to evaluate whether residential sorting generates systematic compositional differences. As shown in Figure 10, there are no systematic differences in demographic characteristics—including gender, parental education, poverty status, language spoken at home, school type, or age—by distance from the crime location. This suggests that students living near and slightly farther from hate crime sites are comparable along observed dimensions.

[Insert Figure 10 here]

5.2 Effects on GPA

We now examine the causal effect of neighborhood hate crime exposure on academic achievement, as measured by semester GPA.

Figure 11 presents event study estimates of Equation 3, plotting coefficients β_t for each semester relative to hate crime occurrence. Panel A reports results for Hispanic students exposed to anti-Hispanic hate crimes; Panel B presents estimates for Black students exposed to anti-Black hate crimes.

[Insert Figure 11 here]

For Hispanic students, pre-treatment coefficients cluster tightly around zero with no dis-

cernible trend, supporting the parallel trends assumption. In the four semesters preceding exposure, point estimates hover near zero and none approach conventional significance levels. Hispanic students living within 0.1 miles of future hate crime locations experienced similar GPA trajectories as comparison students residing 0.1 to 3 miles away prior to the incidents.

Following exposure, we observe a substantial and persistent decline in GPA among nearby Hispanic students. The effect emerges in the semester of the crime ($t = 0$), with a point estimate of -0.06 points, though the confidence interval includes zero. The impact intensifies in subsequent semesters, reaching -0.10 points by $t = 1$ and stabilizing near that magnitude through $t = 4$. Hate crime exposure generates lasting academic consequences that neither dissipate nor further deteriorate beyond one semester post-exposure. Effect magnitudes and the sustained nature of impact mirror our findings on school-based hate crime exposure. One distinction worth noting is that we focus on more severe crimes (Part I offenses) in the neighborhood analysis, in contrast to the preponderence of less serious offenses (Part II) in the school analysis.

Black students exposed to anti-Black hate crimes demonstrate no statistically significant response across the event window. Pre-treatment coefficients are more variable than in the Hispanic sample, with the estimate at $t = -3$ reaching -0.10 points, though statistically indistinguishable from zero. Post-treatment effects similarly fail to achieve significance, with point estimates oscillating between -0.03 and 0.01 points. The absence of detectable impacts aligns with our school-level findings and may reflect factors such as differential exposure patterns or variation in community support structures.

Table 2 summarizes the average treatment effect on the treated, collapsing the post-period into a single indicator. Column (1) presents baseline estimates with student fixed effects and zip code-by-semester fixed effects. For Hispanic students (Panel A), the average post-treatment effect is -0.064 points (SE = 0.030), significant at the 5 percent level. This represents approximately 2.5 percent of the standard deviation in semester GPA, which is very similar to the 3.0 percent standard deviation decrease for exposure to police killings (Ang, 2023). Our estimates suggest that hate crimes generate academic consequences of similar magnitude.

[Insert Table 2 here]

We assess robustness across specifications in columns (2)-(7). Column (2) incorporates grade-by-semester fixed effects to account for differential grade distribution between treat-

ment and comparison students, which could generate spurious effects if grade-specific trends coincide with hate crime timing. The point estimate remains virtually unchanged at -0.061 points. Column (3) adds event-by-semester fixed effects, exploiting variation in hate crime timing within incident-specific cohorts, yielding a stable estimate of -0.061. Column (4) controls for local violence, defined as non-hate-motivated Part I crimes targeting the same minority group within a 0.5-mile radius. This addresses the concern that hate crimes may coincide with broader violence increases that independently affect outcomes. The treatment effect persists at -0.061, indicating impacts reflect hate crime exposure specifically rather than general violence.

Column (5) restricts the sample to zipcodes experiencing a single hate crime during the study period, eliminating potential confounding from cumulative exposure. The estimated effect of -0.064 remains consistent. Column (6) excludes hate crimes occurring after 2019 to ensure results are not driven by COVID-19 disruptions, which fundamentally altered educational environments and may have interacted with hate crime exposure. The point estimate remain unchanged at -0.064. Column (7) implements event-study weights following Goodman-Bacon (2021) to address potential bias from treatment effect heterogeneity. The weighted estimate of -0.049 (SE = 0.019) is somewhat smaller but remains statistically and economically meaningful.

For Black students (Panel B), the average treatment effect ranges from -0.003 to -0.032 across specifications, with standard errors between 0.044 and 0.085. No specification yields a statistically significant estimate, corroborating the event study evidence of muted academic responses. The imprecision partly reflects lower density of Black students in affected neighborhoods, limiting power to detect moderate effects.

5.3 Heterogeneity Analysis

Average treatment effects may mask important variation across student subgroups and crime characteristics. Understanding this heterogeneity illuminates underlying mechanisms and identifies populations requiring targeted support following hate crime incidents.

Figure 12 presents treatment effect heterogeneity for anti-Hispanic hate crimes along two dimensions: student demographics (Panel A) and crime characteristics (Panel B). Each point estimate represents the average post-treatment effect for a specific subgroup. All specifications include student fixed effects and zip code-by-semester fixed effects, with standard errors

clustered at the zip code level.

[Insert Figure 12 here]

There are no apparent differences in responses among male and female Hispanic students, since results show largely overlapping coefficients and confidence intervals. We similarly do not find evidence of effect heterogeneity by grade level (e.g. middle or high school). The point estimate is -0.164 among grades 6-8, compared to 0.031 for high schoolers. While magnitudes suggest larger effects among younger students in our sample, the relatively sizable standard errors preclude the conclusion of differential impacts. The same applies to the home language of students: we cannot reject the equivalence of coefficients for students who have English as a home language, although the coefficients are more precisely estimated for students who do not have English as a home language.

Next we turn to socioeconomic measures such as parental education and economic disadvantage. Hispanic students whose parents have less than a college education and those deemed as economically disadvantage based on family income have effect sizes around -0.08, compared to better-off students for whome hate crime exposure yields a positive coefficient of 0.06. This suggests potentially larger effects among more vulnerable, disadvantaged students.

Panel B examines heterogeneity by the gender and age of both victims and suspects. We classify an individual as a young victim or suspect if they are of school-going age, defined as under 18 years old. Across all attributes, point estimates on the impact of hate crime exposure are negative and sufficiently close to each other that we cannot reject their equivalence. The lack of disparate impact in this context may reflect data limitations or the particular circumstances of school-based hate crime exposure in this study.

These patterns suggest that family background as a mediating channel, with the effects concentrated among economically disadvantaged students and those whose parents have not completed college. Two caveats merit emphasis: stratified estimates involve smaller samples with reduced statistical power, and the heterogeneity of crime characteristics may partly reflect neighbourhood selection rather than causal differences in the impact of crime.

5.3.1 Spillover Effects on Non-targeted Students

The preceding analyses focus on students from targeted racial groups—those sharing the identity singled out by hate crimes. However, hate crimes may generate broader effects ex-

tending beyond direct victims and their co-ethnic peers. Non-targeted students residing in the same neighborhoods could experience spillover effects through heightened awareness of intergroup tensions. Alternatively, the absence of effects among non-targeted students would strengthen the interpretation that observed impacts operate primarily through identity-based threat perceptions rather than generalized environmental deterioration.

Figure 13 examines whether hate crimes affect non-targeted students residing near incidents. Panel A presents event study estimates for non-Hispanic students exposed to anti-Hispanic hate crimes, comparing those within 0.1 miles to those 0.1 to 3 miles away. Pre-treatment coefficients fluctuate around zero (-0.10 to 0.00 points) without discernible trends. Post-treatment estimates similarly cluster near zero (-0.05 to 0.05), with no coefficient achieving statistical significance. Panel B examines non-Black students exposed to anti-Black hate crimes, revealing identical patterns with point estimates ranging from -0.05 to 0.08 points and no systematic post-exposure divergence.

[Insert Figure 13 here]

The absence of spillovers carries three implications. First, it validates our identification strategy by demonstrating that proximity to hate crimes does not mechanically correlate with academic disruption independent of targeted group membership. If unobserved neighborhood shocks were driving results, all nearby students would exhibit similar effects regardless of racial identity. The null findings thus function as a placebo test, confirming effects concentrate among students sharing the targeted identity. Second, results suggest hate crimes operate through identity-based psychological mechanisms rather than generalized environmental channels, as non-targeted students share the same schools and neighborhoods yet exhibit no detectable response. Third, minimal spillovers mitigate general equilibrium concerns that could complicate interpretation by contaminating control groups with indirect exposure.

We acknowledge power limitations—our analysis rules out spillovers exceeding approximately 0.10 points but cannot exclude smaller effects. Nevertheless, tight clustering of point estimates around zero throughout the event window, contrasting sharply with significant negative effects among targeted students, provides evidence that spillovers are either absent or economically small.

5.4 Effects on Mental Health

Beyond academic performance, hate crime exposure may generate psychological consequences manifesting in classifications of emotional disturbance, which renders students eligible for special education services. Under IDEA, emotional disturbance encompasses persistent patterns of relationship difficulties, inappropriate behavior, pervasive unhappiness or depression, and/or physical symptoms that adversely affect educational performance.

We examine whether neighborhood hate crime exposure increases emotional disturbance classification using administrative special education records. Because classification occurs annually, we estimate event study specifications at the academic year level, tracking students from three years before to three years after exposure. The specification follows Equation 3 with annual outcomes and zipcode-by-year fixed effects.

[Insert Figure 14 here]

Figure 14 presents event study estimates. Panel A shows results for Hispanic students exposed to anti-Hispanic hate crimes. Pre-treatment coefficients cluster near zero, supporting parallel trends. Following exposure, emotional disturbance diagnoses increase sharply in year $t = 0$, but the effect quickly attenuates and becomes statistically indistinguishable from zero in subsequent years. Panel B presents results for Black students exposed to anti-Black hate crimes. Consistent with null GPA effects, we observe no statistically significant changes in emotional disturbance classification throughout the event window.

[Insert Table 3 here]

Table 3 summarizes the magnitude of these effects through a parametric event study specification. Panel A reports dynamic treatment effects for each post-exposure year separately, while Panel B presents the average treatment effect pooling all post-periods. For Hispanic students (Column 1), the year 0 effect is 0.011 percentage points ($p < 0.10$), corresponding to approximately one additional classification per 91 exposed students. The coefficients for years 1–3 are near zero and statistically insignificant, confirming the transitory nature of the impact. The average treatment effect of 0.002 percentage points is not statistically significant, illustrating why examining period-specific effects is crucial for detecting short-lived impacts that would otherwise be masked by averaging across all post-periods. For Black students (Column 2), all coefficients are small in magnitude and statistically insignificant, consistent with

the graphical evidence.

The concentration of effects in the immediate aftermath suggests hate crime exposure triggers acute psychological distress. Given that formal classification requires sustained, severe symptoms warranting special education placement, our estimates likely understate the full scope of psychological impacts affecting students with subclinical distress. The rapid attenuation after year zero indicates that while initial distress is severe enough to warrant clinical recognition, formal classifications return to baseline levels within one year.

5.5 Effects on Safety Perception

The academic and mental health consequences may operate through students' subjective assessments of safety. We examine effects on self-reported safety using LAUSD's School Experience Survey, administered annually to students in grades 4–12. We construct binary indicators for feeling safe at school and in their neighborhood, estimating annual event study specifications with student fixed effects and zipcode-by-year fixed effects.

[Insert Figure 15 here]

Figure 15 presents event study estimates for Hispanic students exposed to anti-Hispanic hate crimes. Panel A examines school safety perceptions. Pre-treatment coefficients fluctuate modestly around zero, supporting parallel trends. Following exposure, perceived school safety declines sharply in year $t = 0$, but the effect attenuates in subsequent years. Panel B examines neighborhood safety perceptions, showing similar patterns with a decline concentrated in the year of exposure.

[Insert Table 4 here]

Table 4 summarizes these effects through a parametric event study specification. Panel A reports dynamic treatment effects for each post-exposure year, while Panel B presents the average treatment effect. For school safety perceptions (Column 1), the year 0 effect is -0.073 percentage points ($p < 0.05$), representing a 7.6 percent decline relative to the baseline mean. This effect is concentrated in the immediate aftermath, with coefficients for years 1–2 near zero. The year 3 coefficient of -0.024 percentage points is not statistically significant. For neighborhood safety perceptions (Column 2), the year 0 effect is -0.049 percentage points, though not statistically significant at conventional levels. The year 3 effect of -0.035 percentage points ($p < 0.10$) suggests some delayed impact. The average treatment effects are -0.016

and -0.017 percentage points for school and neighborhood safety, respectively, with the latter achieving marginal significance.

The spillover from neighborhood to school contexts merits emphasis. Students exposed to hate crimes near their homes report feeling less safe at school, suggesting threat perceptions generalize across environments or collective anxiety diffuses through peer discussions. Both safety perceptions and emotional disturbance classifications exhibit largely transient effects—concentrated in year zero—while GPA declines persist for four to five semesters. This temporal pattern suggests that students may fall behind academically during an initial period of heightened distress, with these learning deficits persisting even as measured indicators of safety perception and mental health return to baseline.

5.6 Long-Run Educational Outcomes

While the preceding analysis documents persistent declines in semester-level GPA following hate crime exposure, a natural question concerns whether these academic disruptions translate into longer-term educational consequences. We now examine three cumulative outcomes measured when students exit LAUSD: cumulative unweighted GPA, high school graduation, and college enrollment.

5.6.1 Empirical Specification

Unlike semester-level GPA, these outcomes are measured once per student and do not vary over time within individuals, precluding the standard event study framework. We address this challenge by exploiting variation in exposure timing across different cohorts of students from the same neighborhoods. Older students who had already graduated at the time of a hate crime were never exposed during their schooling, while younger students from the same neighborhoods who were still enrolled faced direct exposure.

We construct a cohort-based sample as follows. For each hate crime incident, we identify all students residing within 3 miles of the crime location, regardless of their enrollment year. We then calculate each student's expected grade at the time of the hate crime:

$$\text{Expected Grade} = \text{Actual Grade} + (\text{Hate Crime Year} - \text{School Year}) \tag{4}$$

This measure indicates what grade a student would be in when the hate crime occurred.

For example, a student who was in 12th grade during the 2020–2021 school year would have an expected grade of 13 for a hate crime that occurred in 2021–2022—that is, two years after graduation. We restrict the sample to students with expected grades between 6 and 14, where grades 6–12 represent students still enrolled (exposed cohort) and grades 13–14 represent students who had already graduated (comparison cohort). When students experience multiple hate crimes, we retain each unique student-by-expected-grade observation, keeping only the nearest incident when multiple crimes occur at the same expected grade. The resulting dataset is uniquely identified by student i and expected grade g.

We estimate a difference-in-differences specification:

$$Y_{ig} = \beta \left(\text{Nearby}_i \times \text{Treated Grade}_g \right) + \alpha_z + \gamma_g + \mathbf{X}_i' \theta + \epsilon_{ig} \tag{5}$$

where Y_{ig} represents the outcome for student i at expected grade g. Nearby_i is an indicator variable equal to one for students residing within 0.1 miles of a hate crime location, with those living 0.1–3 miles away serving as the comparison group. The indicator Treated Grade_g equals one for students in expected grades 6–12 and zero for those in grades 13–14. The coefficient β captures the differential effect of hate crime exposure between students who were still enrolled and those who had already graduated.

The specification includes student residential zipcode fixed effects (α_z), expected grade fixed effects (γ_g), and individual-level controls $\mathbf{X}_i$ comprising gender, poverty status, and parental education. Standard errors are clustered at the zipcode level.

5.6.2 Results

Table 5 presents estimates of long-run educational outcomes for Hispanic students exposed to neighborhood hate crimes. Column (1) shows that proximity to hate crimes during schooling reduces cumulative unweighted GPA by 0.095 points ($p < 0.01$), representing approximately 4.1 percent of the control mean. Column (2) indicates that high school graduation declines by 2.4 percentage points ($p < 0.05$), a 5.6 percent reduction relative to the control mean. Column (3) shows that college enrollment decreases by 3.1 percentage points ($p < 0.01$), corresponding to a 10.6 percent reduction. These effects are statistically significant across all three outcomes.

[Insert Table 5 here]

The magnitude of these effects is substantial. The GPA decline of 0.095 points cumulates over students' high school careers, consistent with the persistent semester-level effects documented earlier. The graduation and college enrollment effects translate to meaningful changes in educational attainment: for every 100 Hispanic students exposed to nearby hate crimes during high school, approximately 2 to 3 fewer graduate and 3 fewer enroll in college compared to unexposed students from similar neighborhoods. Given that college enrollment serves as a key pathway to economic mobility, these reductions suggest hate crimes may perpetuate racial inequality by impeding human capital accumulation during critical developmental periods.

The cohort-based design compares students from the same neighborhoods who differ in whether they were still enrolled when hate crimes occurred. The identifying assumption is that absent hate crimes, educational outcomes would have evolved similarly across cohorts within neighborhoods. While we cannot fully rule out cohort-specific trends correlated with hate crime timing, the inclusion of expected grade fixed effects mitigates this concern. The results are consistent with hate crime exposure shaping educational trajectories that extend beyond high school, affecting not only academic performance but also attainment milestones that determine postsecondary opportunities.

6 Non–Hate Crime Exposure as a Comparison

Our findings thus far demonstrate that exposure to hate crimes generates significant academic and socio-emotional consequences for students from targeted racial groups, while non-targeted students exhibit no comparable effects. A critical question remains: do these impacts reflect the bias-motivated nature of hate crimes specifically, or would any violent incident involving minority victims produce similar effects regardless of perpetrator intent?

To address this question, we examine whether non-hate crimes with minority victims generate comparable impacts on student outcomes. If the observed effects operate primarily through identity-based threat perceptions—whereby students perceive hate crimes as symbolic attacks on their entire group—then incidents lacking explicit bias motivation should not produce the same pattern of results, even when victims share students' racial identity. Conversely, if the effects reflect general exposure to violence in students' environments, we would expect similar impacts regardless of whether crimes are classified as hate-motivated.

6.1 Data and Empirical Approach

We obtained comprehensive crime incident data from two law enforcement agencies covering the period 2009–2021. For school-based incidents, we use data from the Los Angeles Police Department (LAPD), which explicitly records whether crimes occurred on school premises—enabling us to cleanly identify on-campus incidents and directly replicate the school-based identification strategy from Section 4. For neighborhood-based incidents, we rely on the Los Angeles County Sheriff's Department (LASD), which has jurisdiction over a broader geographic area encompassing the residential neighborhoods where LAUSD students live throughout the county, including areas outside LAPD jurisdiction where some students reside. Both datasets include geocoded records of regular crimes along with demographic information on victims.[18]

Given the substantially larger volume of non-hate crimes compared to hate crimes, we draw a 1 percent random sample from each dataset to maintain computational tractability. For school-based analysis, this yields 168 crime incidents occurring at middle and high schools where the victim was Hispanic. For neighborhood-based analysis, we obtain 545 part I incidents with Hispanic victims. We then reconstruct stacked event study samples following the same procedures described in Sections 4 and 5, ensuring comparability in sample construction.

The specifications mirror those used in our main analyses. For school-based incidents, we compare Hispanic students in schools experiencing crimes with Hispanic victims to students in nearby control schools, estimating Equation 1 with student fixed effects and zipcode-by-semester fixed effects. For neighborhood-based incidents, we compare Hispanic students residing within 0.1 miles of crime locations to those living 0.1 to 3 miles away, estimating Equation 3 with the same fixed effects structure. Standard errors are clustered at the school level for school-based analysis and at the zipcode level for neighborhood-based analysis.

6.2 Results

Figure 16 presents event study estimates for both school-based (Panel A) and neighborhood-based (Panel B) exposure to non-hate crimes with Hispanic victims. Across both settings,

[18]The publicly available incident-level LASD crime data lack demographic details on victims. We obtained these variables through a FOIA request; however, the dataset does not include information on whether incidents occurred on school premises. Consequently, for analyses of within-school exposure to crime, we rely on LAPD data, which record the location type of each incident.

we observe no systematic relationship between crime exposure and Hispanic students' GPA. Pre-treatment coefficients cluster around zero with no discernible trends, confirming comparability between treatment and comparison groups prior to incidents. More importantly, post-treatment estimates remain statistically indistinguishable from zero throughout the event window. Point estimates oscillate between approximately -0.05 and 0.05 in both panels, with confidence intervals consistently encompassing zero. This pattern contrasts sharply with our main findings, where hate crime exposure generated persistent GPA declines of 0.06 to 0.10 points lasting four or more semesters.

[Insert Figure 16 here]

The null results hold despite the fact that these crimes involve actual violence against minority victims and occur in the same geographic contexts as hate crimes—either on school campuses or in residential neighborhoods where students live. The absence of effects cannot be attributed to insufficient statistical power, as our sample sizes (168 school incidents and 545 neighborhood incidents) provide reasonable precision, yielding confidence intervals that would detect moderate effects if present.

6.3 Interpretation and Implications

These findings carry several important implications for interpreting our main results. First, they demonstrate that the academic consequences documented in Sections 4 and 5 are not simply mechanical artifacts of exposure to violence in students' schools or neighborhoods. If proximity to crime per se disrupted learning—through mechanisms such as generalized fear, environmental stress, or community disorder—we would observe comparable effects here. The contrast between null effects for non-hate crimes and substantial negative effects for hate crimes suggests that bias motivation constitutes the critical distinction.

Second, the results support an identity-based threat interpretation of hate crime impacts. Students appear to respond not merely to individual acts of violence against co-ethnic victims, but specifically to incidents signaling that their group is targeted for who they are. This aligns with theoretical frameworks emphasizing hate crimes' unique capacity to generate collective threat perceptions extending beyond direct victims to entire communities sharing the targeted identity.

Finally, the placebo analysis addresses concerns about potential confounding from cor-

related neighborhood characteristics. One might worry that schools or neighborhoods experiencing hate crimes differ systematically in ways that independently affect student trajectories. However, these same locations also experience non-hate crimes with minority victims—incidents that occur in identical geographic contexts yet produce no effects. This within-location comparison strengthens causal inference by demonstrating that something specific to hate crimes, rather than broader neighborhood attributes, drives our main findings.

Taken together, the placebo results reinforce our core conclusion that hate crimes exert distinctive impacts on targeted students' human capital formation. The effects we document reflect not general exposure to violence, but rather the unique psychological and social consequences of bias-motivated attacks that threaten students' sense of safety and belonging based on immutable group characteristics.

7 Conclusion

This paper provides new evidence on how hate crime exposure can affect student outcomes. We link comprehensive, georeferenced hate crime with administrative student records from Los Angeles County to more precisely measure exposure and clearly identify causal effects than previously possible. Specifically, we implement complementary identification strategies that allow for studying both school- and neighborhood-based exposure. The first compares students in treated schools to control schools in a stacked difference-in-differences with student fixed effects. The second compares students living within 0.1 miles of neighborhood hate crimes to peers who are slightly more distant.

We show that exposure to hate crimes meaningfully impacts students' academic performance and socio-emotional well-being, with findings consistent across both school- and neighborhood-based designs. Exposure to anti-Hispanic hate crimes leads to persistent GPA declines for Hispanic students, while Black students exposed to anti-Black hate crimes show no analogous effects. Moreover, we find no statistically significant spillover effects on students outside the targeted group. This is consistent with hate-crime effects, if present, operating through identity-specific perceived threats rather than generalized disruptions affecting all students.

Beyond academic achievement, hate crime exposure increases disciplinary incidents, reduces perceived safety at school and in neighborhoods, and raises the likelihood that students

are classified as "emotionally disturbed." The durability of academic effects on GPA, compared to the more transitory nature of behavioral and safety responses, suggests that learning losses may accumulate even after safety perceptions begin to recover. These findings contribute to our understanding of how exposure to racially-motivated hate crimes – an extreme manifestation of identity-based discrimination – can impede educational achievement during critical developmental periods. This complements existing work documenting discrimination in labor market returns to human capital and underscores the importance of understanding how discrimination shapes the formation of human capital itself. The adverse impacts on behavioral and socioemotional outcomes further suggest that hate crimes affect not only cognitive learning but also the development of non-cognitive skills and socioemotional capacities. Taken together, these findings emphasize the need for comprehensive, school- and community-level approaches to address the growing prevalence and consequences of such adverse exposures.

References

Altonji, Joseph G., and Rebecca M. Blank. 1999. "Race and gender in the labor market." 3143–3259.

Ang, Desmond. 2020. "The Effects of Police Violence on Inner-City Students*." *The Quarterly Journal of Economics*, 136(1): 115–168.

Ang, Desmond. 2023. "*The Birth of a Nation:* Media and Racial Hate." *American Economic Review*, 113(6): 1424–1460.

Baker, Dominique J., and Tolani Britton. 2024. "Hate Crimes and Black College Student Enrollment." *Education Finance and Policy*, 19(2): 187–217.

Bayer, Patrick, Kerwin Kofi Charles, and Ellora Derenoncourt. 2025. "Racial Inequality in the Labor Market." *NBER Working Paper No. 33372*.

Bell, James G., and Barbara Perry. 2015. "Outside Looking In: The Community Impacts of Anti-Lesbian, Gay, and Bisexual Hate Crime." *Journal of Homosexuality*, 62(1): 98–120.

Bertrand, M., and E. Duflo. 2017. "Field Experiments on Discrimination." In *Handbook of Economic Field Experiments*. Vol. 1, 309–393. Elsevier.

Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review*, 94(4): 991–1013.

Bharadwaj, Prashant, Manudeep Bhuller, Katrine V. Løken, and Mirjam Wentzel. 2021. "Surviving a mass shooting." *Journal of Public Economics*, 201: 104469.

Blinder, Alan S. 1973. "Wage Discrimination: Reduced Form and Structural Estimates." *The Journal of Human Resources*, 8(4): 436–455.

Cabral, Marika, Bokyung Kim, Maya Rossin-Slater, Molly Schnell, and Hannes Schwandt. 2025. "Trauma at School: The Impacts of Shootings on Students' Human Capital and Economic Outcomes." *The Review of Economic Studies*, rdaf027.

Carrell, Scott E., and Mark L. Hoekstra. 2010. "Externalities in the Classroom: How Children Exposed to Domestic Violence Affect Everyone's Kids." *American Economic Journal: Applied Economics*, 2(1): 211–228.

CDC. 2025. "Bullying."

Chang, Eunsik, and María Padilla-Romo. 2023. "When Crime Comes to the Neighborhood: Short-Term Shocks to Student Cognition and Secondary Consequences." *Journal of Labor Economics*, 41(4): 997–1039.

Cristancho, Juan C., Drew H. Bailey, Greg J. Duncan, Andres Molano, Arturo Harker, and Ervyn Norza. 2024. "Effects of homicide timing on test scores: Quasi-experimental evidence from two cities in Colombia." *Child Development*, 95(4): 1124–1141.

deGiere, Gregory. 2004. "Protecting Californians from Hate Crimes." California Senate Office of Research.

Farrell, Amy, and Sarah Lockwood. 2023. "Addressing Hate Crime in the 21st Century: Trends, Threats, and Opportunities for Intervention." *Annual Review of Criminology*, 6(Volume 6, 2023): 107–130.

Federal Bureau of Investigation. 2024. "Reported Hate Crime at Schools: 2018-2022." U.S. Department of Justice.

Fryer Jr., Roland G. 2011. "Racial inequality in the 21st century: the declining significance of discrimination." In *Handbook of Labor Economics*. Vol. 4, , ed. David Card and Orley Ashenfelter, 855–971. Elsevier.

Gelber, Katharine, and Luke McNamara. 2016. "Evidencing the harms of hate speech." *Social Identities*, 22(3): 324–341.

Gershenson, Seth, and Erdal Tekin. 2017. "The Effect of Community Traumatic Events on Student Achievement: Evidence from the Beltway Sniper Attacks." *Education Finance and Policy*.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*, 225(2): 254–277.

Heckman, James J. 1998. "Detecting Discrimination." *Journal of Economic Perspectives*, 12(2): 101–116.

Hellyer, Joshua, and Johanna Gereke. 2024. "The shadow of fear: hate crime victimization and stress after Charlottesville." *Frontiers in Psychology*, 15: 1384470.

Jaskiewicz, Anna. 2024. "Anti-Black Hate Crimes and Black-White Health Gap."

Kaplan, Jacob. 2021. "Uniform Crime Reporting (UCR) program data: A practitioner's guide." *CrimRxiv*.

Kena, Grace, and Alexandra Thompson. 2021. "Hate crime victimization, 2005–2019." *Report NCJ*, 300954.

LA County Commission on Human Relations. 2001. "Hate Crime Report 2000."

LA County Commission on Human Relations. 2022. "Hate Crime Report 2021."

Lang, Kevin, and Jee-Yeon K. Lehmann. 2012. "Racial Discrimination in the Labor Market: Theory and Empirics." *Journal of Economic Literature*, 50(4): 953–958.

Li, Xuan, and Hantao Wu. 2024. "Political Responses to Hate Crimes." *Available at SSRN 4969664*.

Masucci, Madeline, and Lynn Langton. 2017. "Hate Crime Victimization, 2004-2015." Bureau of Justice Statistics.

National Institute of Justice. 2021. "Overview of Hate Crime." Department of Justice.

Neal, Derek A., and William R. Johnson. 1996. "The Role of Premarket Factors in Black-White Wage Differences." *Journal of Political Economy*, 104(5): 869–895.

Oaxaca, Ronald. 1973. "Male-Female Wage Differentials in Urban Labor Markets." *International Economic Review*, 14(3): 693–709.

Offerman, Evelyne C. P., Michiel W. Asselman, Floor Bolling, Geert-Jan J. M. Stams, Rianne J. Bosman, Ramón J. L. Lindauer, and Machteld Hoeve. 2025. "Prevalence of adverse childhood experiences among students with emotional and behavioral difficulties or disorders - A multi-level meta-analysis." *Journal of Applied Developmental Psychology*, 101: 101871.

Perfect, Michelle M., Matt R. Turley, John S. Carlson, Justina Yohanna, and Marla Pfenninger Saint Gilles. 2016. "School-Related Outcomes of Traumatic Event Exposure and Traumatic Stress Symptoms in Students: A Systematic Review of Research from 1990 to 2015." *School Mental Health*, 8(1): 7–43.

Perry, Barbara, and Shahid Alvi. 2012. "'We are all vulnerable': The in terrorem effects of hate crimes." *International Review of Victimology*, 18(1): 57–71.

Petach, Luke, and Dustin Rumbaugh. 2025. "Hate Crime and Human Capital." *The Review of Black Political Economy*, 52(2): 216–235.

Rossin-Slater, Maya, Molly Schnell, Hannes Schwandt, Sam Trejo, and Lindsey Uniat. 2020. "Local exposure to school shootings and youth antidepressant use." *Proceedings of the National Academy of Sciences*, 117(38): 23484–23489.

Shany, Adi. 2025. "Too Scared for School?: Effects of Terrorism on Students' Achievement." *Journal of Human Resources*, 60(3): 977–1018.

Sharkey, Patrick. 2010. "The acute effect of local homicides on children's cognitive performance." *Proceedings of the National Academy of Sciences*, 107(26): 11733–11738.

Sharkey, Patrick, Amy Ellen Schwartz, Ingrid Gould Ellen, and Johanna Lacoe. 2014. "High Stakes in the Classroom, High Stakes on the Street: The Effects of Community Violence on Student's Standardized Test Performance." *Sociological Science*, 1: 199–220.

Thompson, Owen. 2025. "Human Capital and Racial Inequality in the US Labor Market." National Bureau of Economic Research w34251.

Torrats-Espinosa, Gerard. 2020. "Crime and Inequality in Academic Achievement Across School Districts in the United States." *Demography*, 57(1): 123–145.

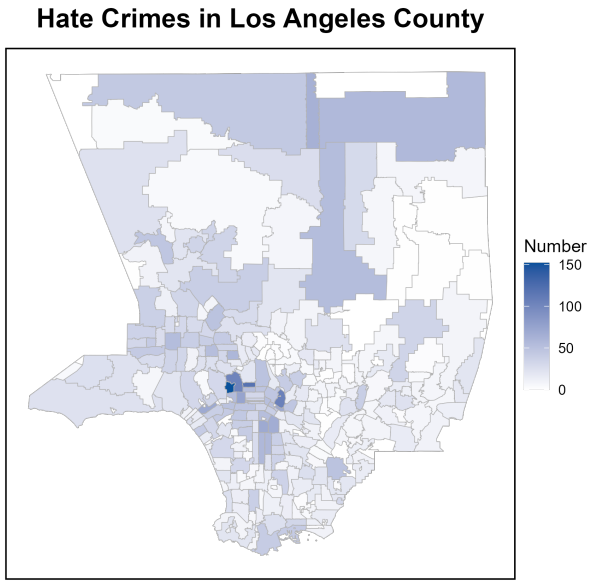
United Nations. 2019. "Strategy and Plan of Action on Hate Speech."

U.S. Government Accountability Office. 2021. "Students' Experiences with Bullying, Hate Speech, Hate Crimes, and Victimization in Schools." GAO-22-104341.

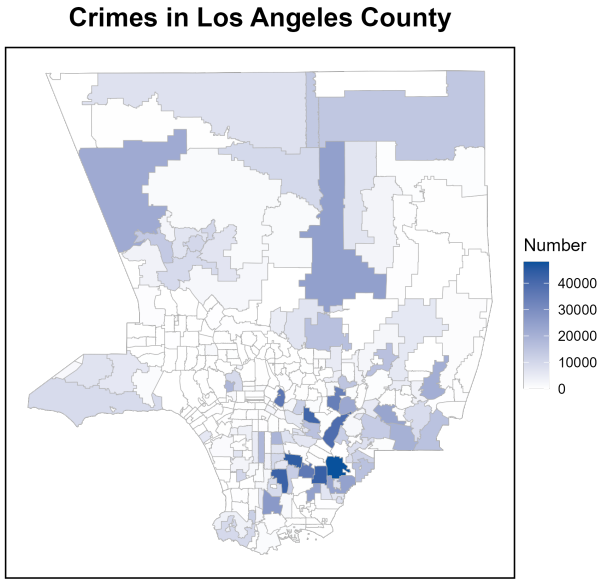
Wenger, Marin R., Brendan Lantz, and Gabriella Gallardo. 2022. "The Role of Hate Crime Victimization, Fear of Victimization, and Vicarious Victimization in COVID-19-Related Depression." *Criminal Justice and Behavior*, 49(12): 1746–1762.

Figures and Tables

Figure 1: Hate Crimes and Crimes in Los Angeles County (2008-2021)



Panel A: Hate Crimes in Los Angeles County



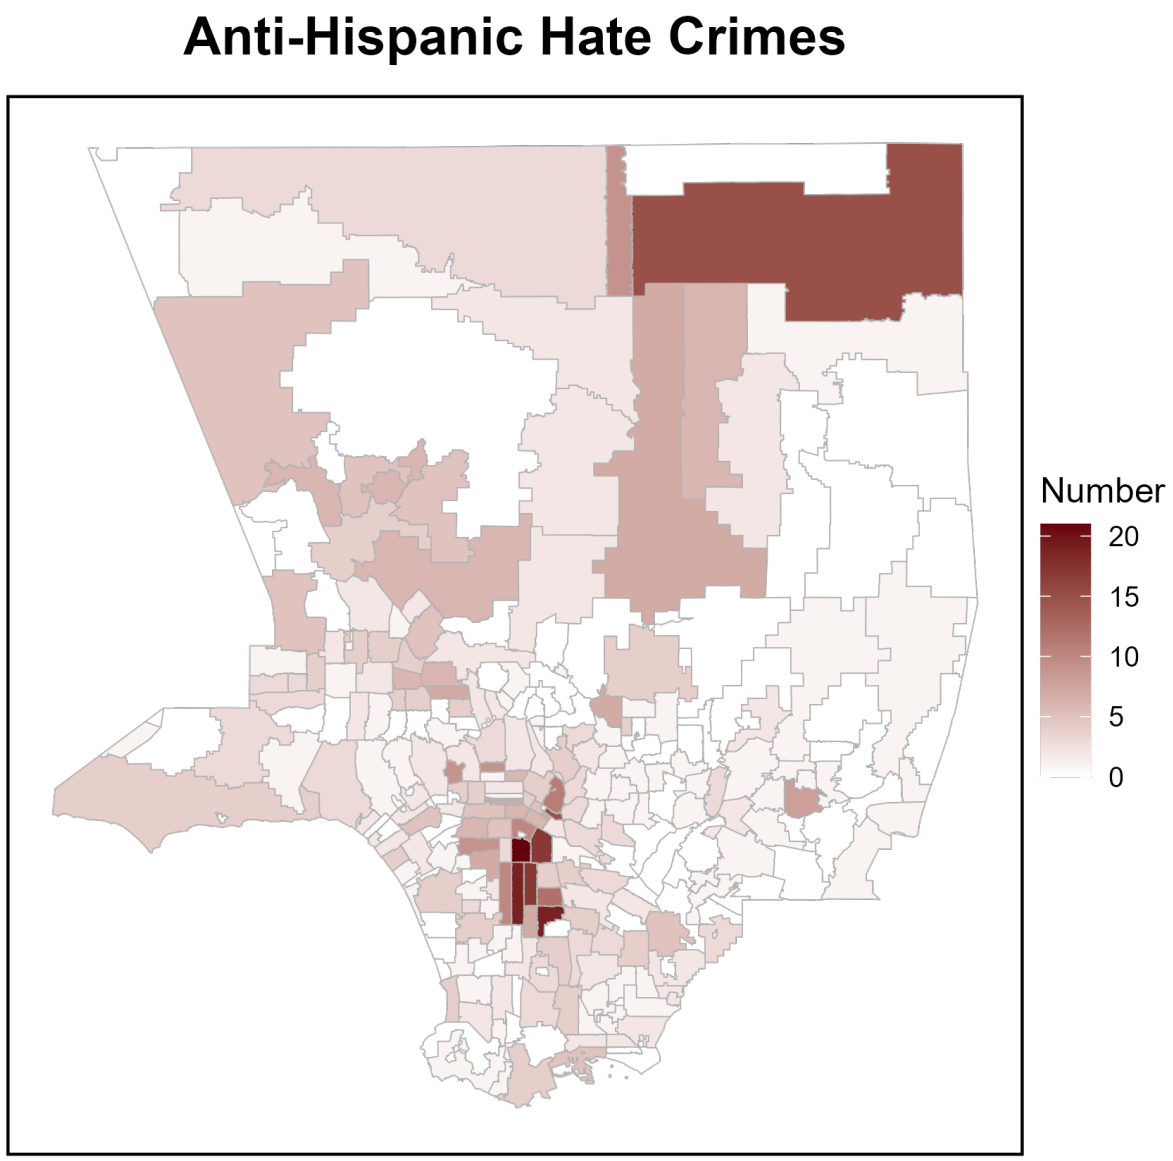
Panel B: Crimes in Los Angeles County

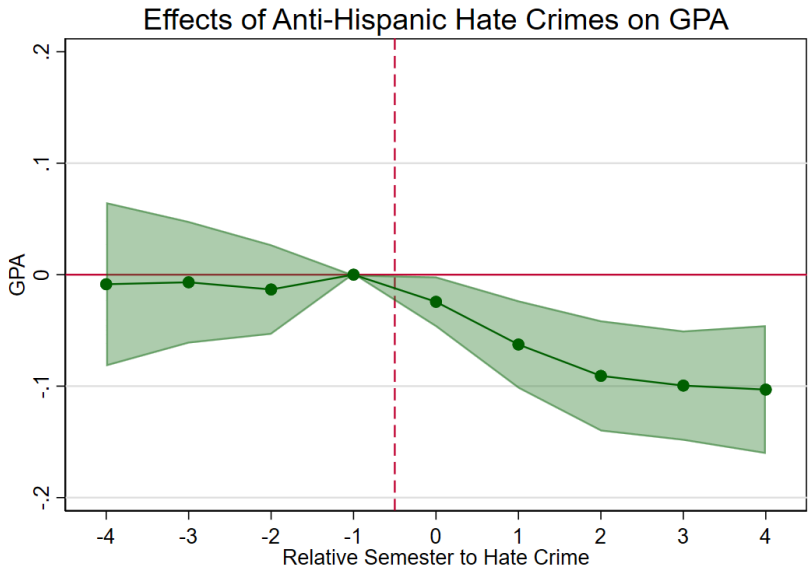
Notes: Panel A shows the spatial distribution of hate crimes in Los Angeles County between 2008 and 2021, using data from the Los Angeles County Commission on Human Relations (LACCHR). Panel B presents the spatial distribution of overall crimes during the same period, based on incident-level data from the Los Angeles County Sheriff's Department (LASD). Both datasets are aggregated to the ZIP-code level to facilitate comparison across geographic areas.

[Back to Page 8]

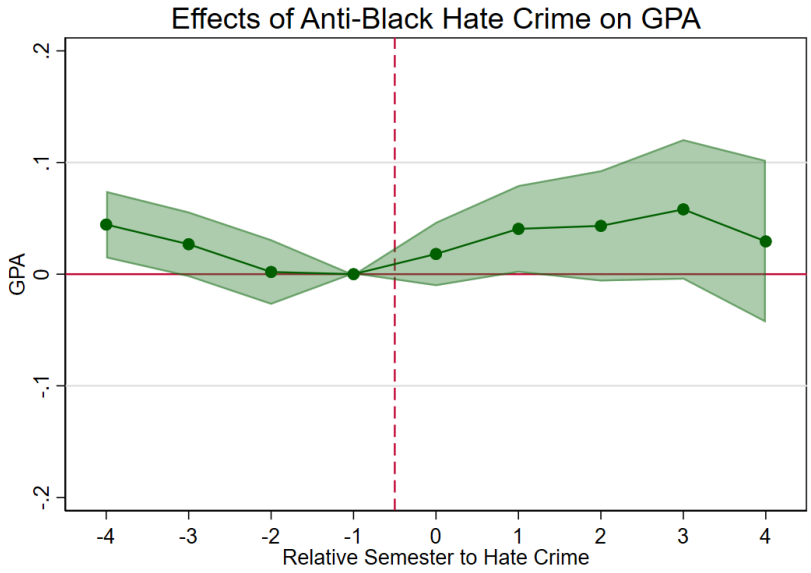
Figure 2: Racially-Targeted Hate Crimes and Racial Demographics in Los Angeles County

0.45
(a) t





Panel A: Anti-Hispanic Hate Crimes

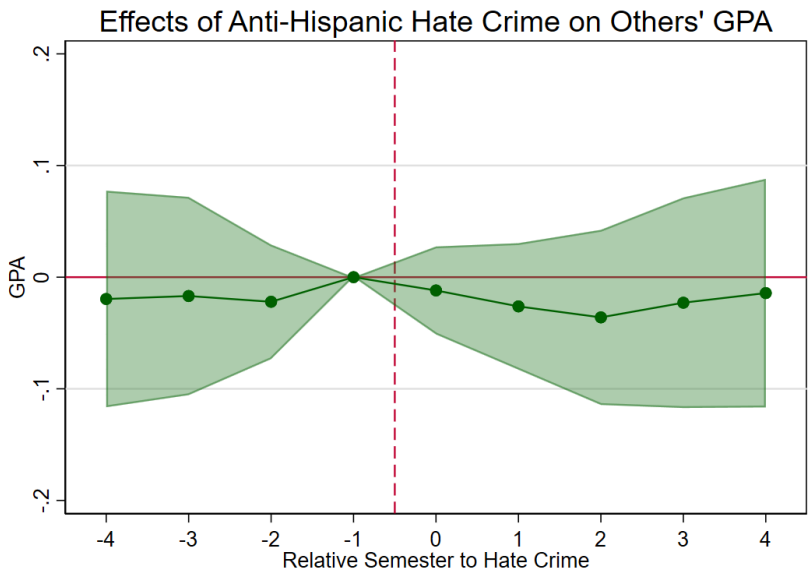


Panel B: Anti-Black Hate Crimes

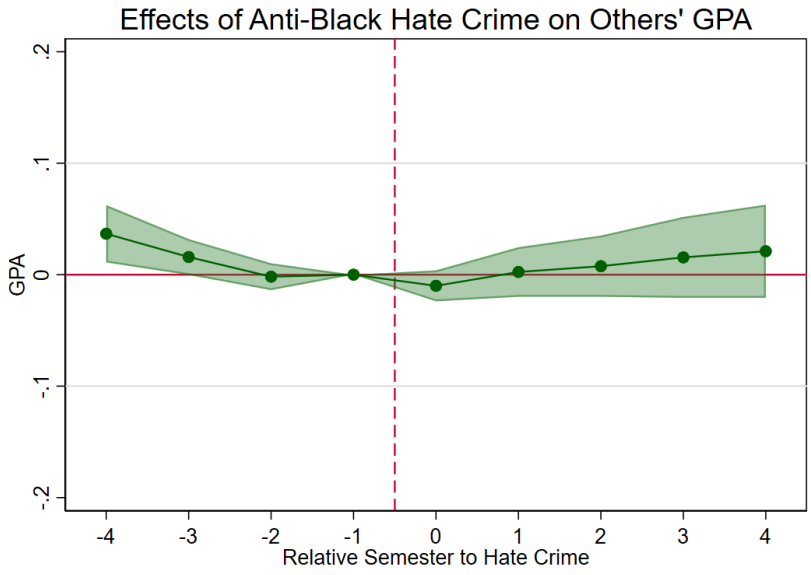
Figure 3: The Impact of Hate Crimes in School on Targeted Students' GPA

Notes: Figures plot stacked event-study estimates of the effects of school hate crime incidents on semester GPA for target students. Panel A reports estimates for anti-Hispanic hate crimes; Panel B reports estimates for anti-Black hate crimes. All specifications follow Equation (2) and include student fixed effects and zipcode-by-semester fixed effects. Standard errors are clustered at the school level. Vertical dashed lines denote the semester of the incident; shaded areas represent 90% confidence intervals.

[Back to Page 15]



Panel A: Anti-Hispanic Hate Crimes

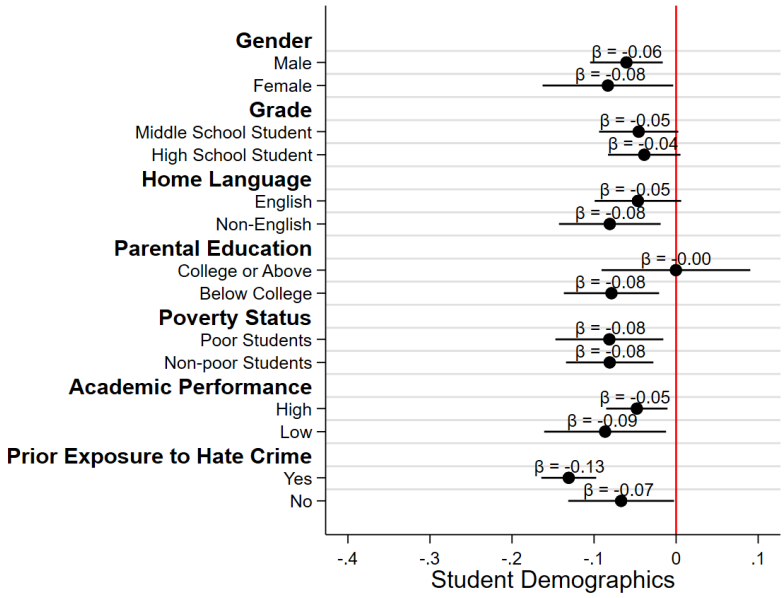


Panel B: Anti-Black Hate Crimes

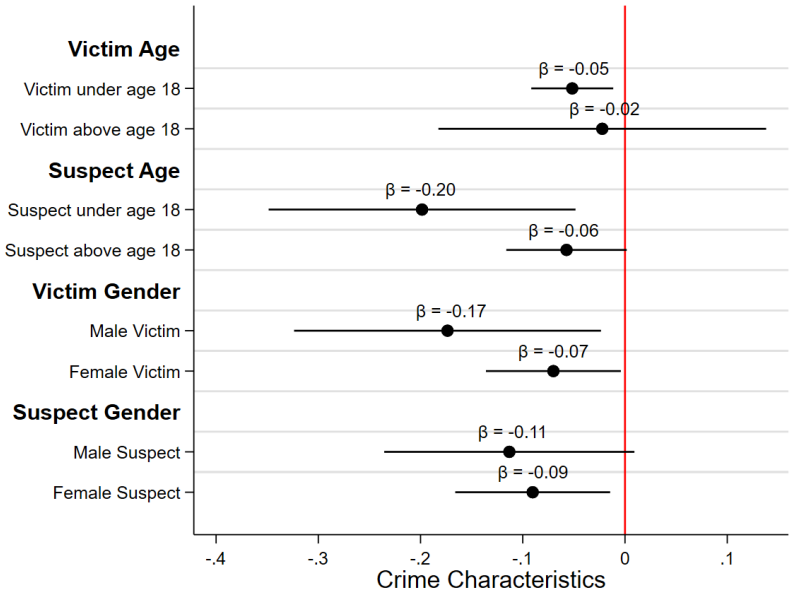
Figure 4: The Impact of Hate Crimes in School on Non-Targeted Students' GPA

Notes: Figures plot stacked event-study estimates of the effects of school hate crime incidents on GPA for non-target students. Panel A reports estimates for anti-Hispanic hate crimes; Panel B reports estimates for anti-Black hate crimes. All specifications follow Equation (2) and include student fixed effects and zipcode-by-semester fixed effects. Standard errors are clustered at the school level. Vertical dashed lines denote the semester of the incident; shaded areas represent 90% confidence intervals.

[Back to Page 15]



Panel A: Heterogeneity by Student Demographics



Panel B: Heterogeneity by Crime Characteristics

Figure 5: Heterogeneous Effects of Anti-Hispanic Hate Crimes in School on GPA

Notes: Figures report heterogeneity in the effects of anti-Hispanic hate crime incidents on GPA for Hispanic students. Panel A examines differences by student demographics (gender, grade level, home language, parental education, poverty status, academic performance, and prior exposure to anti-Hispanic hate crimes). Panel B examines heterogeneity by crime characteristics, including victim and suspect gender as well as student status (proxied by being under age 18). Estimates are based on the specification in Equation (2), including student fixed effects and zipcode-by-semester fixed effects. Standard errors are clustered at the school level. Horizontal bars indicate 90% confidence intervals.

[Back to Page 17]

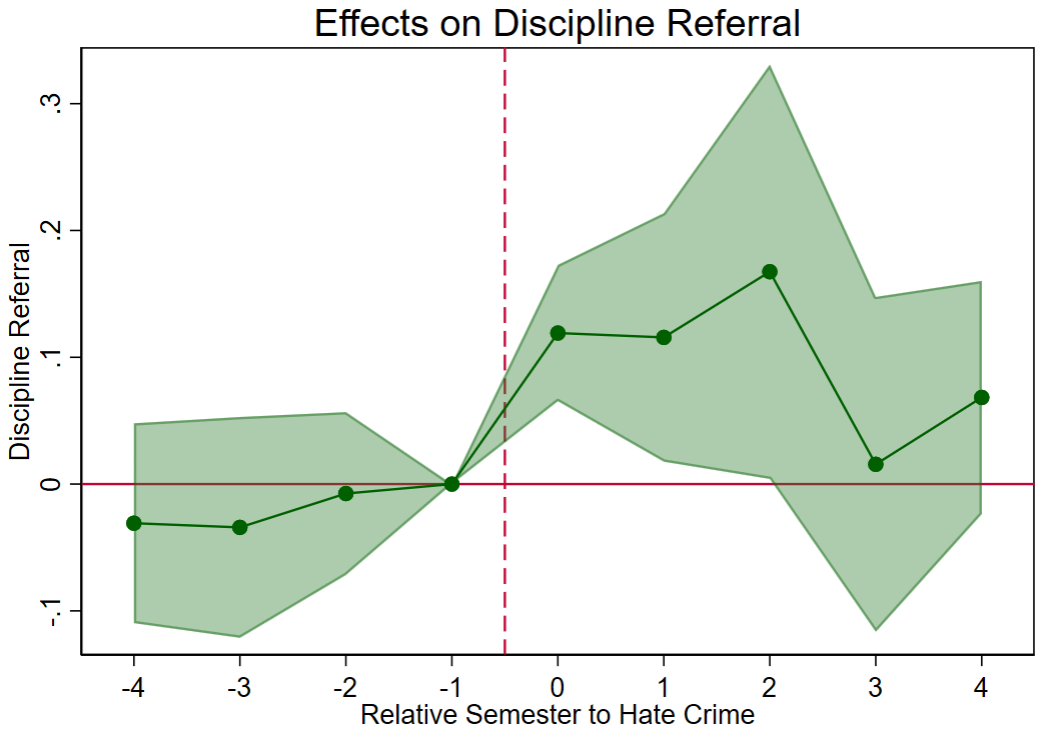


Figure 6: The Impact of Anti-Hispanic Hate Crimes in School on Discipline Referral
Notes: Figures plot stacked event-study estimates of the effects of anti-Hispanic hate crime incidents in school on the number of discipline referrals for Hispanic students. All specifications follow Equation (2) and include student fixed effects and zipcode-by-semester fixed effects. Standard errors are clustered at the school level. Vertical dashed lines denote the semester of the incident; shaded areas represent 90% confidence intervals.

[Back to Page 18]

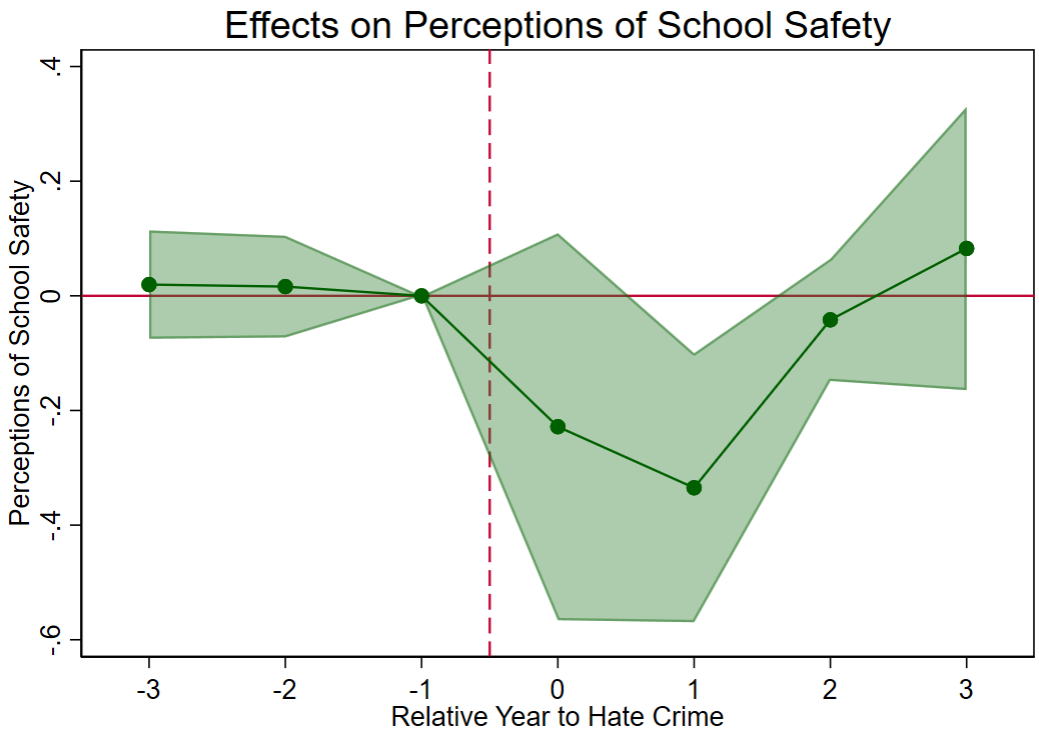


Figure 7: The Impact of Anti-Hispanic Hate Crimes in School on Perceptions of School Safety

Notes: Figures plot stacked event-study estimates of the effects of anti-Hispanic hate crime incidents in school on the self-reported perception of school safety for Hispanic students. All specifications follow Equation (2) and include student fixed effects and zipcode-by-year fixed effects. Standard errors are clustered at the school level. Vertical dashed lines denote the academic year of the incident; shaded areas represent 90% confidence intervals.

[Back to Page 18]

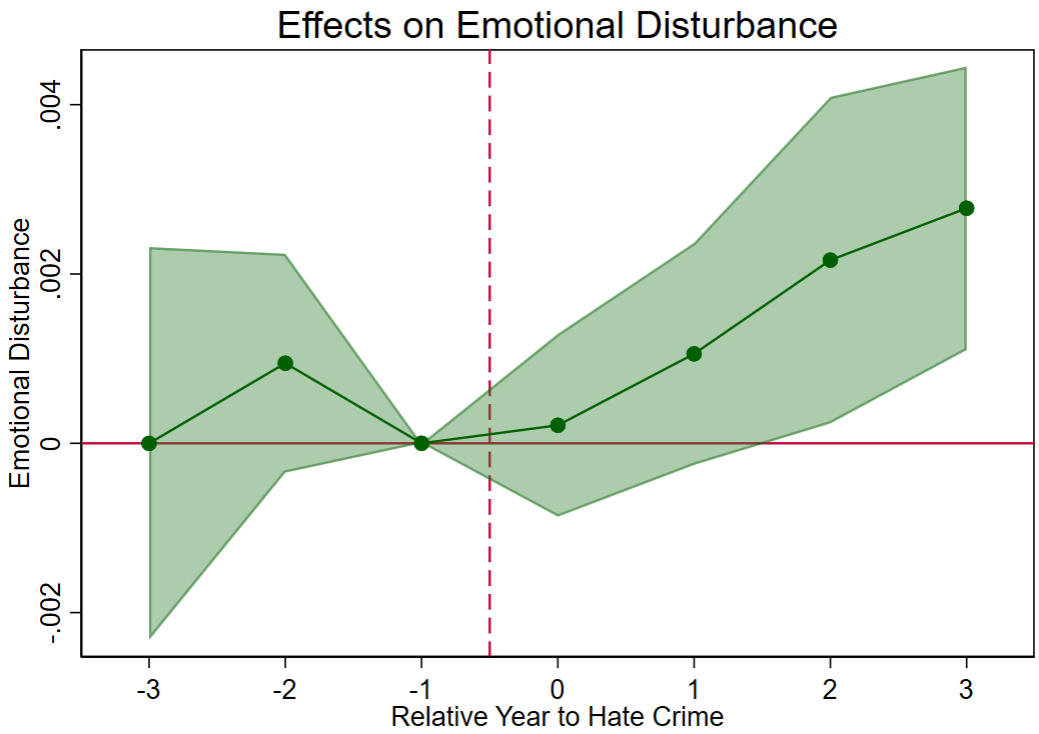
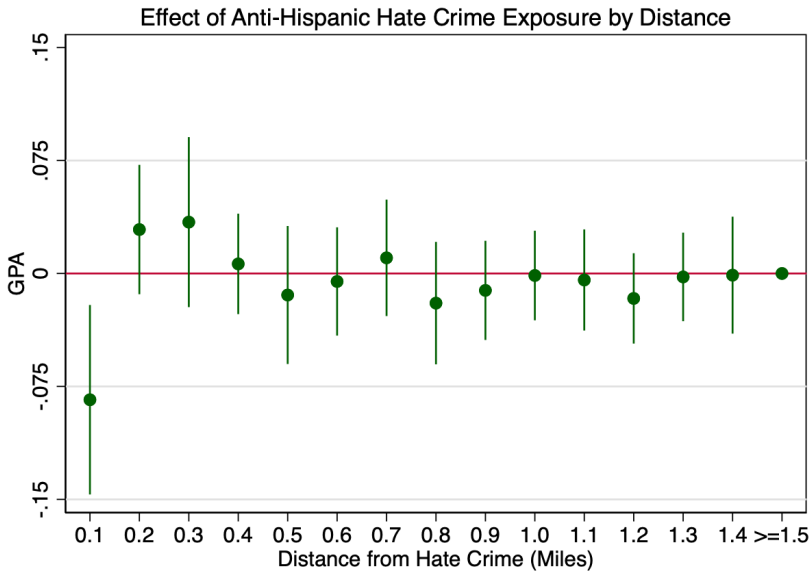


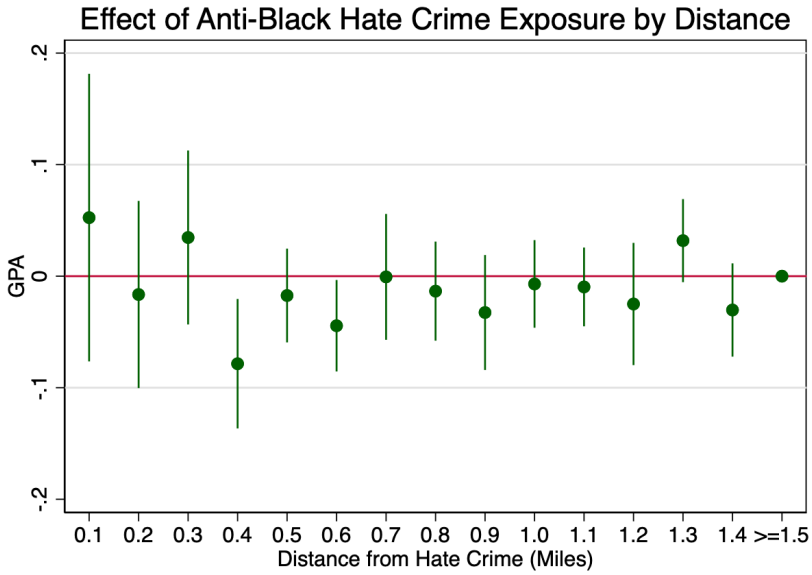
Figure 8: The Impact of Anti-Hispanic Hate Crimes in School on Emotional Disturbance

Notes: Figures plot stacked event-study estimates of the effects of anti-Hispanic hate crime incidents in school on an indicator for Hispanic students' emotional disturbance. All specifications follow Equation (2) and include student fixed effects and zipcode-by-year fixed effects. Standard errors are clustered at the school level. Vertical dashed lines denote the academic year of the incident; shaded areas represent 90% confidence intervals.

[Back to Page 19]



Panel A: Hate Crimes Against Hispanic



Panel B: Hate Crimes Against Black

Figure 9: Effects of on Student GPA by Distance

Notes: This figure reports stacked event-study estimates of hate crime effects on student GPA based on Equation 2. Panel A displays effects of anti-Hispanic hate crimes on Hispanic students by distance; Panel B displays effects of anti-Black hate crimes on Black students by distance. Each point estimate corresponds to a 0.1-mile distance bin, measuring proximity to the nearest hate crime incident from student home. The reference category comprises students whose nearest incident occurred between 1.5 and 3 miles away. All specifications include student and zipcode-by-semester fixed effects. Vertical bars represent 90% confidence intervals; standard errors are clustered at the zipcode level.

[Back to Page 20]

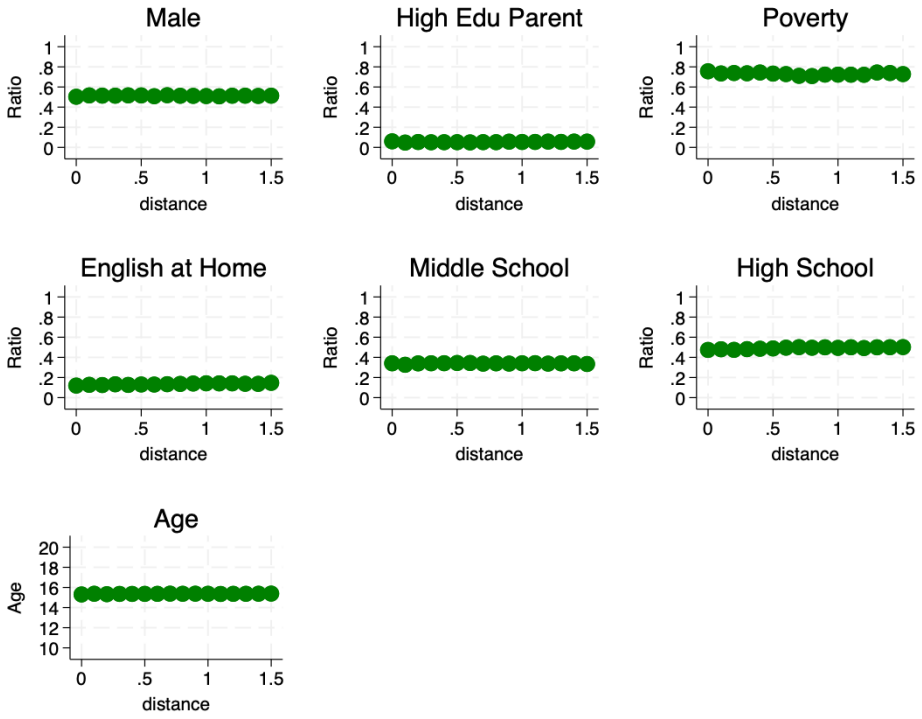
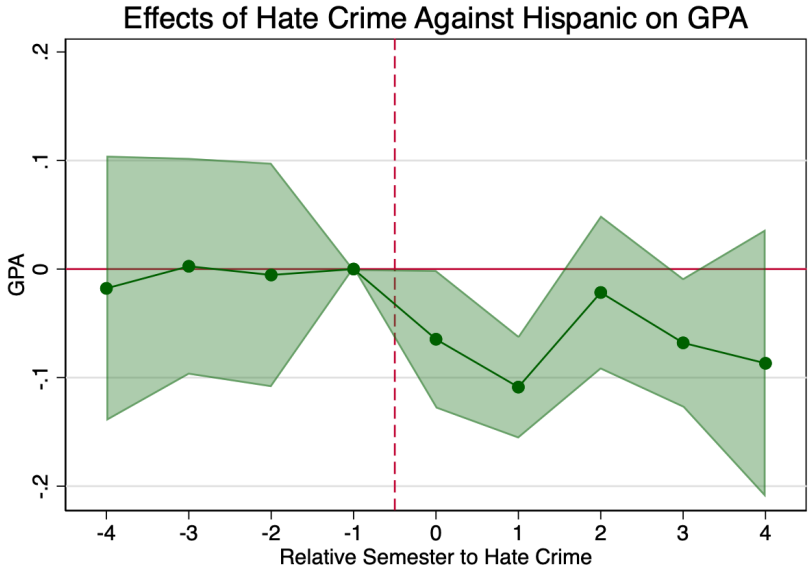


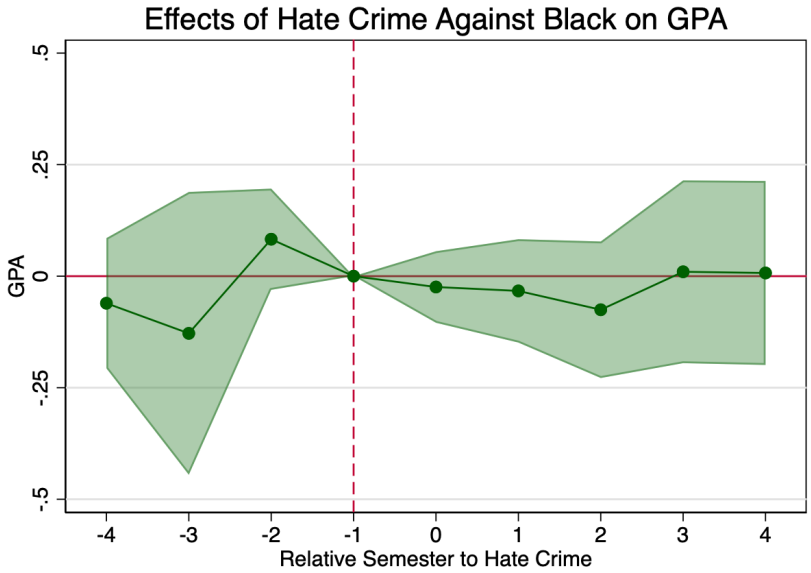
Figure 10: Demographic Characteristics of Students by Distance

Notes: This figure plots mean student demographic characteristics by distance from hate crime locations, measured in 0.1-mile bins. "Male" denotes the proportion of male students. "High Education Parent" indicates the share of students whose parents have a high level of education. "Poverty" is the share of students classified as economically disadvantaged. "English at Home" captures the proportion of students who speak English at home. "Middle School" and "High School" denote the proportion of students attending each school type, respectively. "Age" reports the average age of students. Across all characteristics, there are no discernible gradients by distance, indicating that demographic composition is balanced across treatment and comparison groups.

[Back to Page 23]



Panel A: Hispanic

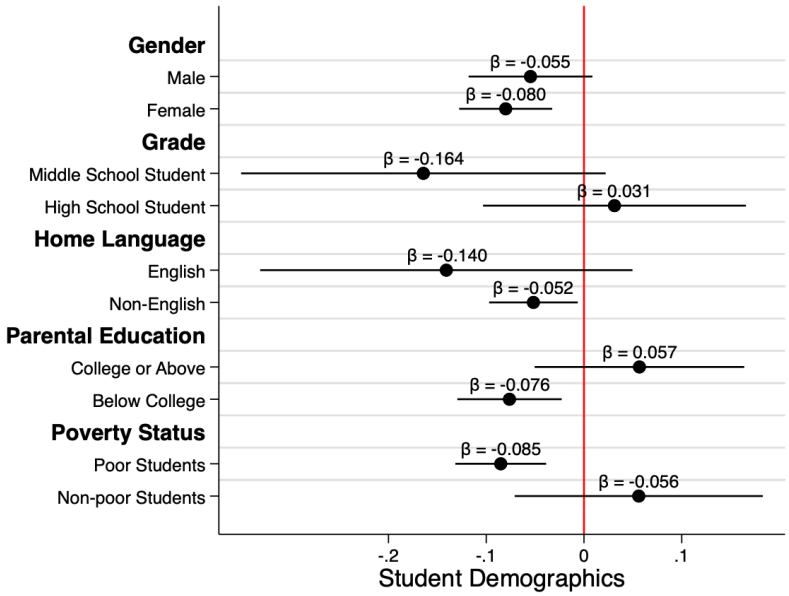


Panel B: Black

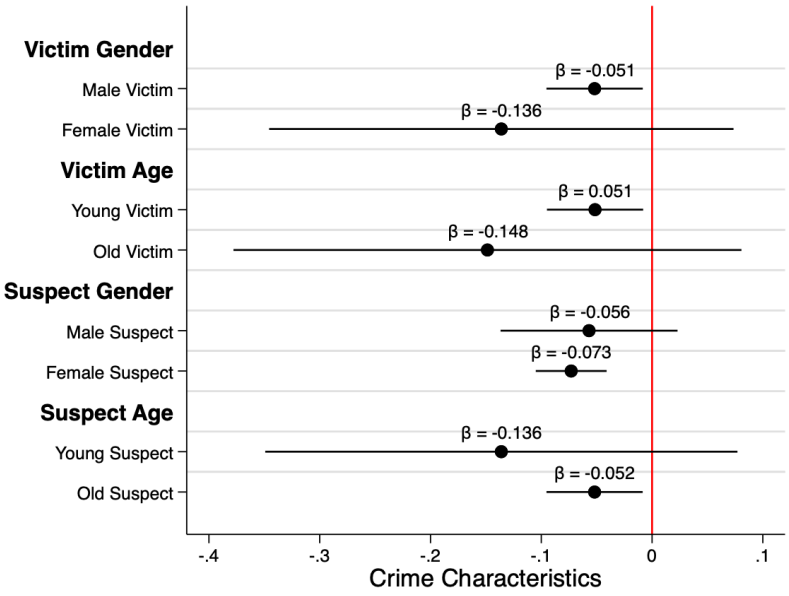
Figure 11: Dynamic Effects of Neighborhood Hate Crimes on Student GPA

Notes: This figure presents event-study estimates of hate crime effects on standardized GPA for targeted students. Treatment students reside within 0.10 miles of a hate crime during high school; control students' hate crime occurred between 0.10 and 3 miles away. Panel A reports anti-Hispanic hate crime effects on Hispanic students; Panel B reports anti-Black hate crime effects on Black students. All specifications include student and zipcode-by-semester fixed effects. Standard errors are clustered at the zipcode level. Vertical dashed lines denote the incident semester; shaded areas represent 90% confidence intervals.

[Back to Page 23]



Panel A: Heterogeneity by Student Demographics

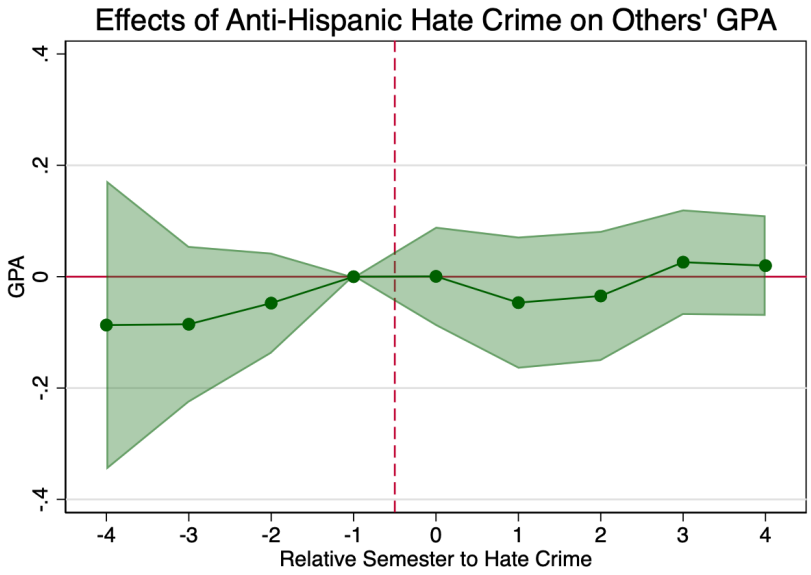


Panel B: Heterogeneity by Crime Characteristics

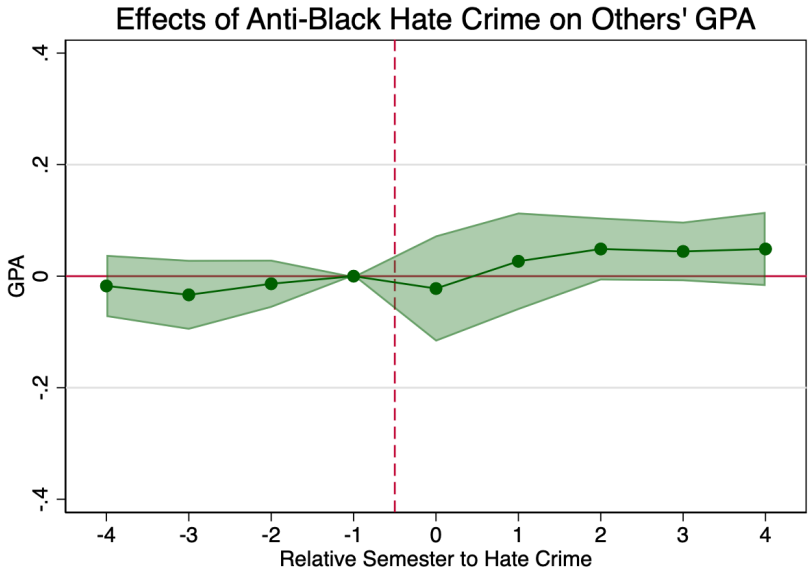
Figure 12: Heterogeneous Effects of Neighborhood Hate Crimes on GPA

Notes: Figures report heterogeneity in the effects of neighborhood anti-Hispanic hate crime incidents on GPA. Panel A examines differences by student demographics (gender, grade level, home language, parental education, and poverty status). Panel B examines heterogeneity by crime characteristics (victim and suspect gender and age). Estimates include student fixed effects and zipcode-by-semester fixed effects. Standard errors are clustered at the zipcode level. Horizontal bars indicate 90% confidence intervals.

[Back to Page 26]



Panel A: Hispanic

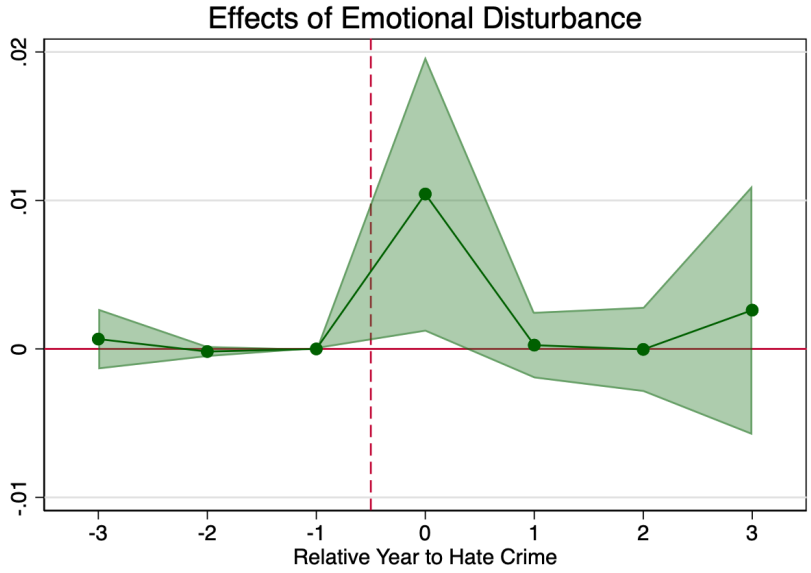


Panel B: Black

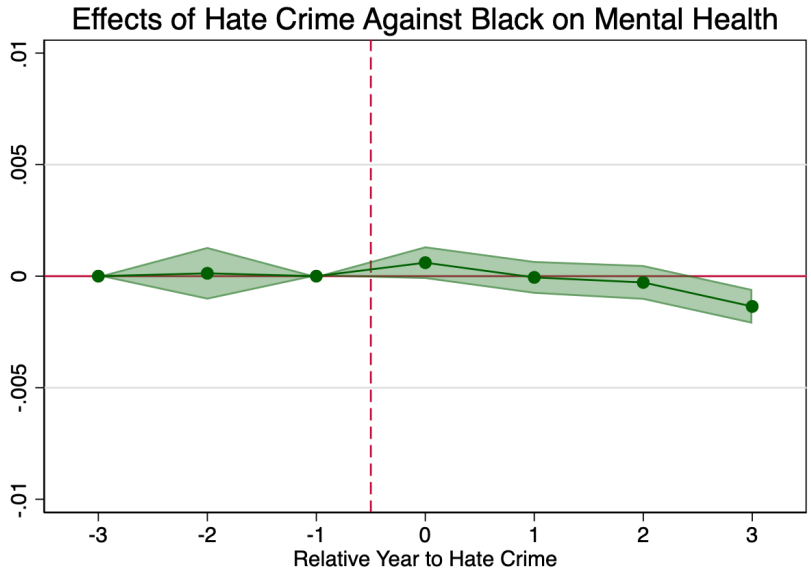
Figure 13: The Impact of Neighborhood Hate Crimes on Non-Targeted Students' GPA

Notes: This figure plots stacked event-study estimates of hate crime effects on GPA for non-target students. Treatment students reside within 0.10 miles of a hate crime during high school; control students' nearest hate crime occurred between 0.10 and 3 miles away. Panel A reports anti-Hispanic hate crime effects on non-Hispanic students; Panel B reports anti-Black hate crime effects on non-Black students. All specifications follow Equation 3 and include student and zipcode-by-semester fixed effects. Standard errors are clustered at the zipcode level. Vertical dashed lines denote the incident semester; shaded areas represent 90% confidence intervals.

[Back to Page 27]



Panel A: Hispanic

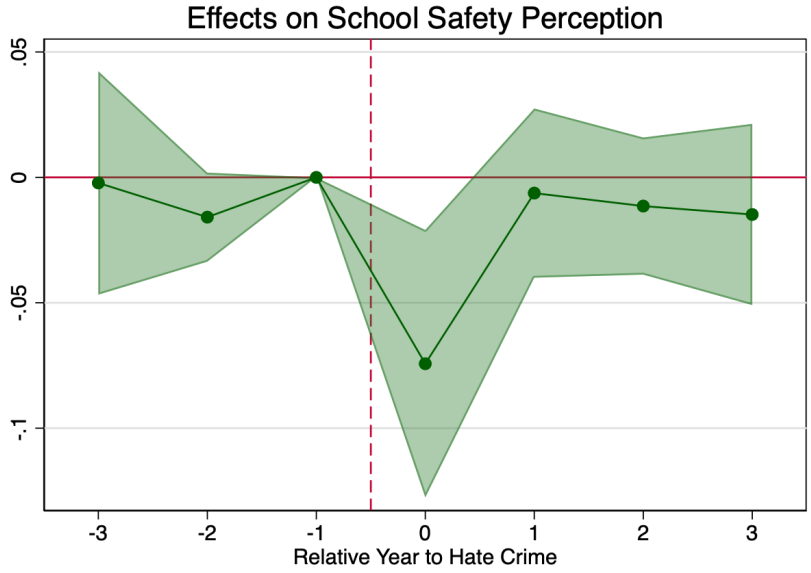


Panel B: Black

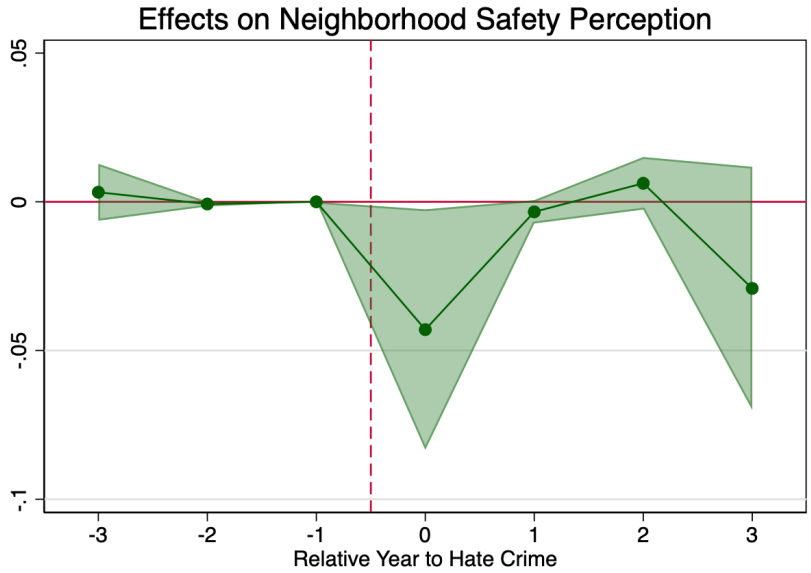
Figure 14: Effects of Neighborhood Hate Crimes on Emotional Disturbence

Notes: Figures plot stacked event-study estimates of the effects of hate crime incidents on emotional disturbence for target students. Panel A reports estimates of anti-Hispanic hate crime among Non-hispanic students; Panel B reports estimates of anti-Black hate crime for Non-Black students. Mental health is measured by an indicator for whether the student was diagnosed with Emotional Disturbance (ED) in a given year. All specifications follow Equation 3 and include student fixed effects and zipcode-by-year fixed effect. Standard errors are clustered at the zipcode level. Vertical dashed lines denote the year of the incident; shaded areas represent 90% confidence intervals.

[Back to Page 28]



Panel A: School Safety Perception

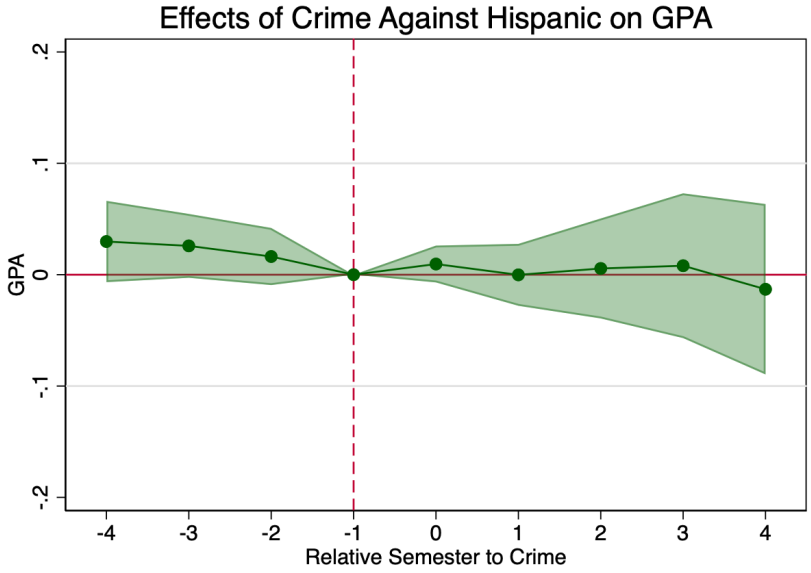


Panel B: Neighborhood Safety Perception

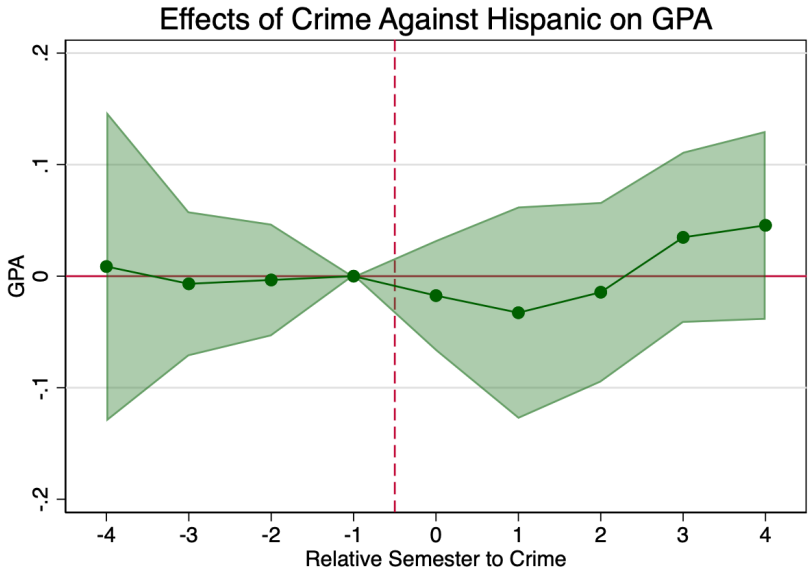
Figure 15: Effects of Neighborhood Hate Crimes on Safety Disturbence

Notes: Figures plot stacked event-study estimates of the effects of neighborhood anti-Hispanic hate crime incidents on students' perception of safety. Panel A reports estimates of anti-Hispanic hate crime on school safety perception; Panel B reports estimates of anti-Black hate crime for Non-Black students. Safety perception is measured by an indicator for whether the student feel safe in school or neighborhood. All specifications follow Equation 3 and include student fixed effects and zipcode-by-year fixed effect. Standard errors are clustered at the zipcode level. Vertical dashed lines denote the year of the incident; shaded areas represent 90% confidence intervals.

[Back to Page 29]



Panel A: School-Based Exposure to Non-Hate Crimes



Panel B: Neighborhood-Based Exposure to Non-Hate Crimes

Figure 16: Effects of Non-Hate Crimes with Hispanic Victims on Hispanic Students' GPA

Notes: This figure presents event study estimates examining whether non-hate crimes with Hispanic victims affect Hispanic students' GPA. Panel A reports results for school-based incidents, comparing Hispanic students in schools experiencing crimes with Hispanic victims (drawn from LAPD data) to students in nearby control schools. Panel B reports results for neighborhood-based incidents, comparing Hispanic students residing within 0.1 miles of crimes with Hispanic victims (drawn from LASD data) to those living 0.1 to 3 miles away. Both analyses use 1 percent random samples of non-hate crime incidents (N=168 for schools, N=545 for neighborhoods) and follow the stacked event study design and specifications described in Sections 4 and 5. All specifications include student fixed effects and zipcode-by-semester fixed effects. Standard errors are clustered at the school level (Panel A) or zipcode level (Panel B). Vertical dashed lines denote the semester of the incident; shaded areas represent 90% confidence intervals.

[Back to Page 34]

Table 1: The Impact of Hate Crimes in School on Targeted Students' GPA

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				GPA				
	Base	Alt. Controls			Alt. Sample			Weight
Panel A: Anti-Hispanic Hate Crimes								
Average Treatment Effect	-0.074**	-0.075**	-0.064**	-0.041*	-0.061**	-0.069*	-0.078**	-0.113*
	(0.035)	(0.035)	(0.032)	(0.021)	(0.030)	(0.035)	(0.035)	(0.061)
Obs.	220192	220192	212810	205696	174807	211179	218685	220192
Panel B: Anti-Black Hate Crimes								
Average Treatment Effect	0.030	0.028	0.019	0.021	0.004	0.032	0.033	0.025
	(0.028)	(0.032)	(0.040)	(0.047)	(0.019)	(0.029)	(0.030)	(0.039)
Obs.	91071	91071	75026	60798	65815	88158	88909	91071
Student FEs	✓	✓	✓	✓	✓	✓	✓	✓
Zipcode×Semester FEs	✓	✓	✓	✓	✓	✓	✓	✓
School×Grade FEs		✓						
Event Weights								✓
Sample Exclusion			Non-First Cases	Multi-Treated Sch.	Transfer Stud.	Part I	Post-2019	

Notes: Each column reports the coefficient on the hate crime school×post interaction from a difference-in-differences regression estimating the effect of school hate crimes on the semester GPA of targeted students. Panel A reports results for anti-Hispanic hate crimes and Panel B for anti-Black hate crimes. All regressions include student fixed effects and zipcode-by-semester fixed effects. Column (1) shows baseline results. Column (2) adds school-by-grade fixed effects. Columns (3)–(7) restrict the sample by excluding, respectively, non-first hate crime cases, schools with multiple hate crime incidents, transfer students, Part I hate crimes, and hate crime incidents occurring after 2019. Column (8) implements event weights for the stacked design. Standard errors clustered at the school level.

[Back to Page 16]

Table 2: The Impacts of Neighborhood Hate Crime on GPA

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	GPA						
	Base	Alt. Controls			Alt. Sample		Weight
Panel A: Anti-Hispanic Hate Crimes *Sample Mean: 2.496*							
Average Treatment Effect	-0.064** (0.030)	-0.061** (0.025)	-0.061** (0.025)	-0.061** (0.025)	-0.064** (0.030)	-0.064** (0.030)	-0.049** (0.019)
Obs.	674,104	674,104	674,104	674,104	674,104	674,104	674,104
Panel B: Anti-Black Hate Crimes *Sample Mean: 2.317*							
Average Treatment Effect	-0.015 (0.085)	-0.018 (0.079)	-0.018 (0.079)	-0.018 (0.079)	-0.015 (0.085)	-0.003 (0.084)	-0.032 (0.044)
Obs.	131,439	131,439	131,439	131,439	131,439	131,439	131,439
Student FEs	✓	✓	✓	✓	✓	✓	✓
Zipcode×Semester FEs	✓	✓	✓	✓	✓	✓	✓
Grade×Semester FEs		✓	✓	✓			
File FEs			✓	✓			
Violence Rate				✓			
Exclusion					Multi-T	post-2019	
Weights							✓

Notes: This table reports the coefficient on the post × treatment group interaction from a difference-in-differences regression of GPA. Panel A reports results for Hispanic students; Panel B reports results for Black students. The treatment group comprises students residing within 0.10 miles of a hate crime incident during high school; the control group comprises students whose nearest incident occurred between 0.10 and 3 miles away. All regressions include student fixed effects and zipcode-by-semester fixed effects, as well as separate controls for event time and treated proximity indicators. Column (1) presents baseline estimates. Columns (2)–(3) incorporate additional fixed effects (grade-by-semester and file-by-semester). Column (4) adds controls for non-hate crime incidents targeting the same minority group. Columns (5)–(6) restrict the sample by excluding, respectively, zipcodes with multiple hate crime incidents and incidents occurring after 2019. Column (7) implements event-study weights for the stacked design. Standard errors are clustered at the zipcode level unless otherwise indicated.

[Back to Page 24]

Table 3: Effects of Neighborhood Hate Crimes on Emotional Disturbance

Variable	(1) Hispanic	(2) Black
Baseline Mean	0.005	0.005
Panel A: Dynamic Treatment Effects		
Year 0	0.011*	-0.001
	(0.006)	(0.001)
Year 1	-0.001	-0.001
	(0.002)	(0.001)
Year 2	-0.000	0.002
	(0.002)	(0.002)
Year 3	-0.002	0.001
	(0.002)	(0.001)
Panel B: Average Treatment Effects		
	0.002	-0.000
	(0.002)	(0.001)
Observations	748,207	118,344
R-squared	0.248	0.306
Student FEs	✓	✓
Year × Zipcode FEs	✓	✓
Grade FEs	✓	✓

Notes: This table presents parametric event study estimates of hate crime effects on emotional disturbance classification. Panel A shows dynamic treatment effects for years 0–3 after hate crime exposure, with all pre-exposure years serving as the reference group. Panel B shows the average treatment effect combining all post-exposure periods. The dependent variable is an indicator for emotional disturbance classification in a given year. Column (1) reports results for Hispanic students exposed to anti-Hispanic hate crimes; Column (2) reports results for Black students exposed to anti-Black hate crimes. Treatment intensity is defined as the inverse distance to hate crimes for students residing within 0.1 miles, with those living 0.1–3 miles serving as the control group. All specifications include student fixed effects, year-by-zipcode fixed effects, and grade-level fixed effects. Standard errors clustered at the zipcode level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Effects of Neighborhood Hate Crimes on Safety Perceptions

Variable	(1) School Safety	(2) Neighborhood Safety
Baseline Mean	0.963	0.956
Panel A: Dynamic Treatment Effects		
Year 0	-0.073**	-0.049
	(0.031)	(0.031)
Year 1	-0.001	-0.015
	(0.019)	(0.011)
Year 2	-0.007	-0.002
	(0.015)	(0.009)
Year 3	-0.024	-0.035*
	(0.028)	(0.021)
Panel B: Average Treatment Effects		
	-0.016	-0.017*
	(0.014)	(0.009)
Observations	410,572	410,709
R-squared	0.732	0.730
Student FEs	✓	✓
Year × Zipcode FEs	✓	✓
Grade FEs	✓	✓

Notes: This table presents parametric event study estimates of anti-hispanic hate crime effects on safety perceptions. Panel A shows dynamic treatment effects for years 0–3 after hate crime exposure, with all pre-exposure years serving as the reference group. Panel B shows the average treatment effect combining all post-exposure periods. The dependent variables are binary indicators for feeling safe at school (Column 1) and in the neighborhood (Column 2), measured annually through LAUSD's School Experience Survey. Results are for Hispanic students exposed to anti-Hispanic hate crimes. Treatment intensity is defined as the inverse distance to hate crimes for students residing within 0.1 miles, with those living 0.1–3 miles serving as the control group. All specifications include student fixed effects, year-by-zipcode fixed effects, and grade-level fixed effects. Standard errors clustered at the zipcode level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Long-Run Educational Outcomes: Cohort-Based Estimates for Hispanic Students

	(1) Cumulative GPA	(2) HS Graduation	(3) College Enrollment
Treatment Dist × Treated Grade	-0.095***	-0.024**	-0.031***
	(0.022)	(0.013)	(0.008)
Observations	879,314	944,693	944,693
Sample Mean	2.342	0.432	0.292
Student Zipcode FEs	✓	✓	✓
Expected Grade FEs	✓	✓	✓
Individual Controls	✓	✓	✓

Notes: This table reports cohort-based difference-in-differences estimates of long-run educational outcomes for Hispanic students exposed to anti-Hispanic hate crimes. The dependent variables are cumulative unweighted GPA (Column 1), an indicator for high school graduation (Column 2), and an indicator for college enrollment (Column 3), all measured when students exit LAUSD. The sample includes students with expected grades 6–16, where expected grade indicates what grade a student would be in at the time of the hate crime. ‘reatment Dist is the inverse distance to hate crimes for students residing within 0.1 miles (zero for those 0.1–3 miles away). Treated Grade equals one for students in expected grades 6–12 (still enrolled) and zero for those in grades 13–16 (already graduated). All specifications include student residential zipcode fixed effects, expected grade fixed effects, and individual-level controls (gender, poverty status, parental education). Standard errors clustered at the zipcode level are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix

A Figures and Tables

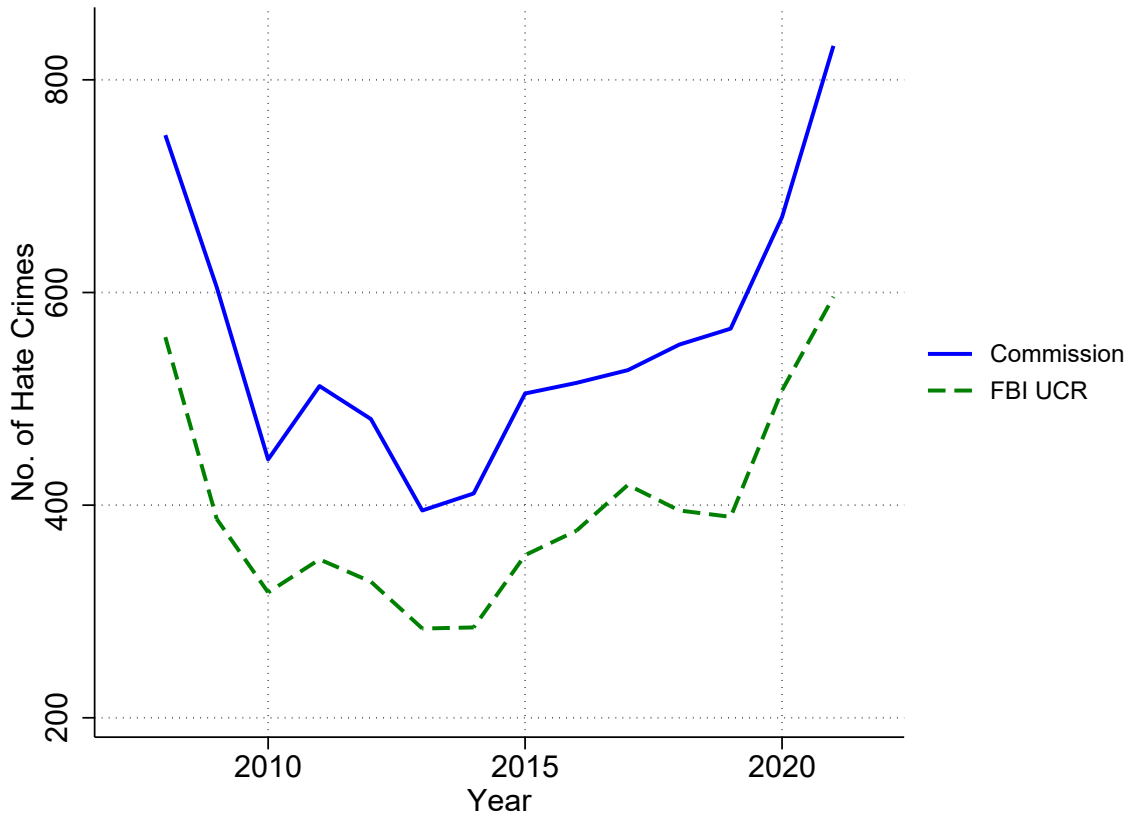


Figure A1: Comparing hate crime data by source

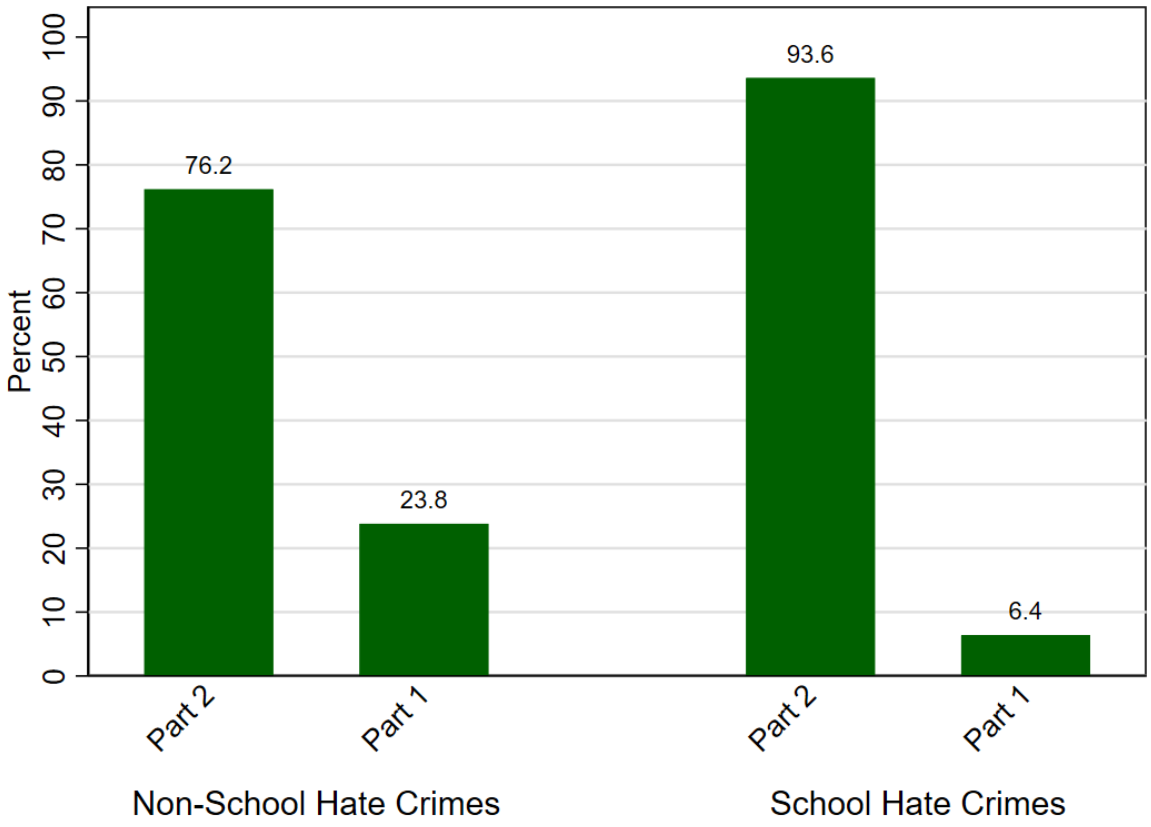
Notes: This figure compares annual hate crime counts in Los Angeles County from 2008 to 2021 using two data sources: the Los Angeles County Commission on Human Relations (solid blue line) and the FBI's Uniform Crime Reports (UCR; green dashed line).

[Back to Page 8]

Figure A1: Characteristics of School-Based versus Non-School Hate Crimes

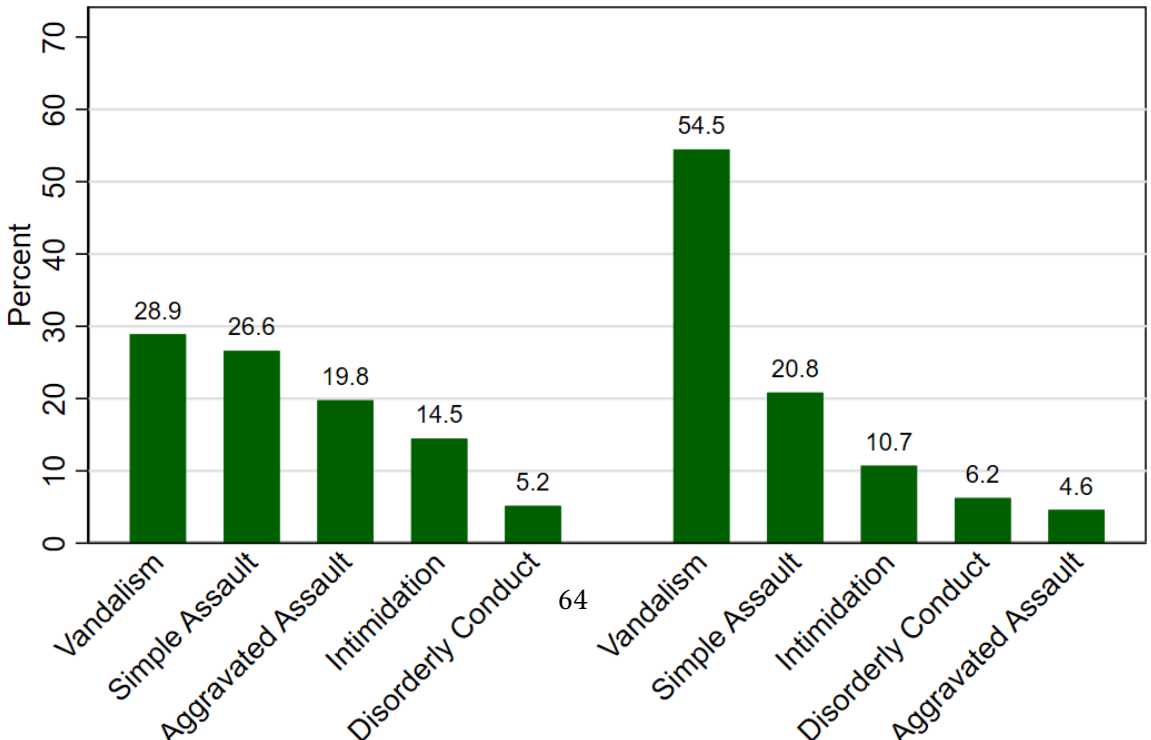
0.45

(a) t



0.45

(b) t



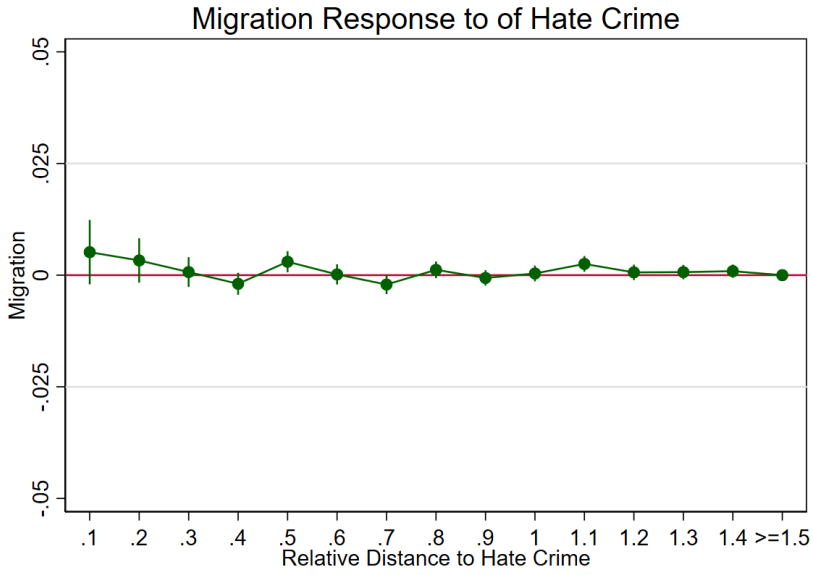


Figure A2: Migration Response to Anti-Hispanic Hate Crimes by Distance

Notes: This figure presents difference-in-differences estimates of hate crime effects on residential migration by distance from incident locations, using L2 voter file data spanning 2014–2022. Each point estimate corresponds to a 0.1-mile distance bin, measuring proximity from registered voters' residential addresses to the nearest anti-Hispanic hate crime location. The outcome is an indicator equal to one if the individual's residential zip code changes between consecutive years. The reference category comprises individuals whose nearest hate crime occurred beyond 1.5 miles. All specifications include individual fixed effects and zipcode-by-year fixed effects. Standard errors are clustered at the zipcode level.

[Back to Page **??**]

Table A1: Balance Between Treated and Control Schools

Variable	Treated (1)	Control (2)	Difference (3)	P-value (4)	Conditional (5)	Cond. P-value (6)
Male Share	0.523	0.533	-0.010	0.459	-0.008	0.470
	(0.027)	(0.073)				
Black Share	0.109	0.068	0.041	0.179	-0.001	0.923
	(0.139)	(0.116)				
Hispanic Share	0.706	0.793	-0.087	0.212	0.032	0.175
	(0.211)	(0.220)				
White Share	0.108	0.073	0.035	0.342	-0.025	0.151
	(0.144)	(0.133)				
Asian Share	0.071	0.085	0.014	0.491	-0.004	0.617
	(0.057)	(0.067)				
Middle School Share	0.324	0.288	0.035	0.768	-0.090	0.207
	(0.471)	(0.424)				
English Home Lang. Share	0.320	0.284	0.036	0.393	-0.029	0.101
	(0.176)	(0.166)				
High Parental Educ. Share	0.141	0.113	0.028	0.497	-0.019	0.236
	(0.121)	(0.147)				
Poverty Share	0.748	0.808	-0.061	0.259	-0.009	0.771
	(0.222)	(0.217)				

Notes: Data collapsed to the school level. Treated schools are defined as those experiencing a hate crime incident, while control schools are schools located in the same zipcode that never experienced a hate crime and are of the same school type. To assess ex-ante comparability, student characteristics are measured in the year prior to the occurrence of a hate crime. Reported variables are school-level shares (averages of individual indicators): gender, race/ethnicity (Black, Hispanic, White, Asian), middle school enrollment, nativity (home language English), parental education (college or above), and poverty status. Columns (1)–(2) report means with standard deviations in parentheses. Columns (3)–(4) report unconditional differences and corresponding p-values. Columns (5)–(6) report differences conditional on zip code fixed effects. All regressions are weighted by the number of students per school, and robust standard errors are reported.

[Back to Page 13]

B Commonly used hate-crime data

Accurately measuring hate crimes presents substantial empirical challenges. A commonly used source for this category of offenses in the United States is the FBI's Uniform Crime Reports (UCR) Hate Crime Statistics. While the UCR provides a long time series and standardized definitions, the reliability of these data depends on a two-tiered process of classification and reporting that may introduce several layers of measurement error.

At the first stage, local law enforcement officers determine whether an incident qualifies as a hate crime based on evidence of "bias motivation." The FBI's guidelines define a hate crime as one in which there is evidence that the offender was motivated by bias against a protected characteristic (e.g., race, religion, ethnicity, sexual orientation, or disability). Determining whether such motivation exists depends heavily on investigative practices and officer discretion. This process requires specialized training and the collection of supporting evidence such as suspect statements, symbols, or victim testimony. However, training standards and implementation vary considerably across states and agencies (Kaplan, 2021). As a result, what constitutes sufficient evidence of bias may differ from one jurisdiction to another.

At the second stage, participation in the UCR program is voluntary. Local law enforcement agencies are responsible for submitting classified incidents to the FBI, and participation varies widely across jurisdictions (Farrell and Lockwood, 2023). Some agencies—particularly in large metropolitan areas—report zero hate crimes for multiple consecutive years, which is implausible given population size and diversity. Such gaps reflect not only differences in administrative capacity but also variation in incentives and community relations that affect whether agencies choose to report hate-related offenses. Together, these factors imply that observed hate crime counts reflect not only true incidence but also institutional capacity and local enforcement culture, complicating empirical inference from this administrative data.

A complementary data source is the Bureau of Justice Statistics' National Crime Victimization Survey (NCVS), which records self-reported bias-motivated incidents regardless of whether they were reported to police. The NCVS consistently reports far higher rates of hate-crime victimization than the FBI's Uniform Crime Reports (UCR), suggesting that the UCR series likely understates the true incidence of hate crimes due to underreporting and variation in classification practices. However, the NCVS identifies hate crimes based on victims' perceptions of bias motivation—such as the offender's perceived identity, the use of hate language or symbols, or other contextual cues—which introduces the potential for over-inclusion of incidents that may not meet the evidentiary standards used in official investigations. On balance, the NCVS provides strong evidence that hate crimes are substantially undercounted in administrative records, even if its estimates may sometimes overstate true incidence.[19]

Neither the UCR nor the NCVS data disclose the precise locations of hate-crime incidents. Yet, as shown by Li and Wu (2024), even severe and violent hate crimes are rarely covered

[19]Victims may not report incidents to local law enforcement agencies because of fear of retaliation, mistrust of authorities, uncertainty about the seriousness of the incident, or the perception that reporting would not lead to meaningful action (Masucci and Langton, 2017).

by the media, implying that their social and psychological effects are likely to be highly localized. If such impacts are indeed geographically concentrated, access to accurate location information becomes crucial for identifying the most directly affected communities and assessing localized spillovers. Taken together, these data limitations underscore the inherent difficulty of constructing reliable and spatially comparable measures of hate-crime incidence across jurisdictions, thereby posing challenges for causal identification in empirical analysis.

[Back to Page 6]

C Details on Hate Crime Data Classification

C.1 Determining What Constitutes a Hate Crime

The Commission undertakes a thorough and consistent assessment to determine whether reported incidents meet the legal definition of hate crimes under California Penal Code § 422.55. According to this statute, hate crime charges may be filed for criminal acts committed, in whole or in part, because of the victim's actual or perceived characteristics, including race, ethnicity, disability, gender, nationality, religion, or sexual orientation. Many of California's statutory definitions require that the perpetrator's intent must be to commit the crime because of, by reason of, or motivated by the victim's real or perceived protected characteristic. One definition states that "because of" implies that the bias motivation must be a cause in committing the offense, although other causes can also exist. In the latter case, the bias motivation must be a substantial factor in realizing the particular result (deGiere, 2004).

Following the statutory guidelines, the Commission's classification process requires that bias, hatred, or prejudice based on these protected characteristics constitute a substantial factor in the commission of the offense. Importantly, these biases or prejudices can manifest before, during, or after the criminal act, and evidence may be either direct or circumstantial (LA County Commission on Human Relations, 2001).

A critical challenge in hate crime classification is establishing that bias motivation played a substantial role in the offense. The Commission considers several types of evidence in making this determination. They include explicit bias-related statements, slurs, or epithets. Evidence can also include the presence of hate symbols, graffiti, or imagery associated with bias, such as swatiskas, nooses, and graffiti containing bias-related slurs and epithets. The circumstances surrounding the incident may suggest bias motivation even without explicit words, symbols, or imagery. The Commission also takes contextual evidence into account, such as the location and timing of the crime and targeting of victims when engaged in activities promoting their identity.

LA County Commission on Human Relations (2001) provides more specific guidance on hate crime classification:

- When bias evidence is based on speech alone, the speech must have threatened violence against a specific person or a group of persons. It must also be clear that the perpetrator had the apparent ability to carry out the threat.
- When derogatory words or epithets are directed against a member of a protected class but no violence is threatened and there is no apparent ability to carry out the threat, it is considered a hate incident. These incidents do not arise to the level of criminal offenses and are protected under free speech laws.
- Graffiti is a hate crime when it is racial, ethnic, religious, homophobic, or sexist in nature. Common signs include use of hate group symbols, slogans, or epithets. To be a hate crime, the graffiti must be directed at a particular person or group of persons.

- Vandalism to a house of worship or ethnic, religious, or LGBTQ organization is generally considered a hate crime.

C.2 Illustrative Examples

Below we provide a few cases demonstrating offenses classified as hate crimes, as presented in the LA County Commission on Human Relations Hate Crime Reports. Given the topical focus of our paper, we only include racially and ethnically motivated hate crimes.

- At a high school a Latino student called a learning disabled black student, a "retarded nigger." The suspect punched the victim in the mouth, face, head and stomach. After the victim was knocked to the ground the suspect continued hitting him and called him "nigger" several more times. Other students pulled the suspect away from the victim until school officials arrived. The suspect was arrested and booked. (2016)
- A black male employee was locking the parking lot gate of an elementary school at night. A young Latino male suspect pulled a hand gun from his waist and told him, "Fuck niggers! You're in the wrong neighborhood! Get out of here before something happens to you, nigger!" The victim fled in his car. (2017)
- Five Latino/as were moving furniture out of an apartment when a black male suspect exited his unit and started arguing with one of them. He yelled, "Too fucking loud, fucking Mexicans" before punching the first victim multiple times. The other victims tried to intervene but the suspect pushed and punched them. He then shouted, "I'll kill you motherfuckers... fuckin' Mexicans! I'm gonna get my gun and shoot all you Mexicans. (2021)
- An Asian male victim was waiting at the train station. Unprovoked, a Black male suspect approached him and began yelling. Despite the victim's efforts to defuse the situation, the suspect became increasingly agitated and eventually spat towards the victim, which struck him on the leg. The suspect continued to yell, "chink," "go back to China," "all Chinese are shit," and "the US should go to war with China." The suspect then said, "I will kill you," and "I'm gonna shoot you, I have a gun." (2023)

[Back to Page 7]

C.3 Addressing Selective Migration in the Neighborhood-Based Analysis

A potential concern for our neighborhood-based analysis is selective migration. Since the LAUSD dataset records only students' most recently reported addresses rather than tracking residential moves over time, families might relocate to safer neighborhoods following hate crime exposure. If students living very close to hate crimes are more likely to move away, our estimates could be biased by differential attrition.

We address this concern through three complementary analyses. First, to more directly assess residential mobility responses, we analyze individual-level migration patterns using L2 voter file data spanning 2014–2022. These administrative records provide annual snapshots of registered voters' residential addresses, allowing us to construct a migration indicator based on whether individuals change zip codes between years. We implement a distance-based difference-in-differences design analogous to our main spatial analysis. We estimate migration responses across 0.1-mile distance bins, where each coefficient β_d captures the differential change in migration rates for individuals residing at distance d from hate crime locations in the post-exposure period, relative to those living beyond 1.5 miles. Figure A2 displays the spatial gradient of migration responses by distance from hate crime locations. Across all distance bins, point estimates cluster tightly around zero with no evidence of differential migration between those living very close to hate crimes and those slightly farther away. While voter files capture only the citizen population and therefore may not fully represent our student sample, this analysis provides direct micro-level evidence that hate crimes do not trigger selective residential mobility at the hyperlocal scale relevant to our identification strategy.

Second, we examine aggregate mobility patterns using American Community Survey data from 2016–2020. For each census block group, we calculate the tenure rate—defined as the share of residents living in the same house as one year prior. We compare tenure rates in block groups where anti-Hispanic hate crimes occurred in the five years prior to those in block groups with no hate crime exposure over the same period. 91.8% of those in treated block groups reported the same residence during the prior year, compared to 91.0% among non-treated block groups ($p = 0.612$). This aggregate-level analysis complements our individual-level evidence by capturing mobility for the full population, including non-citizens who do not appear in voter files.

Taken together, the analyses based on individual voter migration and aggregate census mobility provide consistent evidence that hate crime exposure is unlikely to induce differential residential mobility between the treatment and comparison groups. This null result may reflect several factors: residential moves are costly and difficult to execute quickly, many families face constraints on where they can relocate, or individual hate crime incidents may not be salient enough to trigger immediate relocation decisions. Nevertheless, we cannot entirely rule out selective migration if moves occur with sufficient delay to fall outside our observation window. Subject to this caveat, the convergent evidence from multiple data sources supports the interpretation that our estimated effects reflect causal impacts of hate crime exposure rather than compositional changes from selective attrition.