Flight from Fights:
School shootings, mobility and neighborhood sorting

Marcos Demetry*

October 28, 2026

Click here for the latest version
[Preliminary work. Do not cite or circulate]

Abstract

Disadvantaged neighborhoods tend to remain disadvantaged. One reason is that it is not random which families respond to violence by moving. Selective exit along a socioeconomic gradient re-sorts families across neighborhoods. I estimate how shootings at U.S. public schools affect enrollment, the socioeconomic composition of the students who remain, and segregation between neighborhoods. Linking the K-12 School Shooting Database to the Common Core of Data for 1990 to 2022, I exploit the exogenous timing of shootings in a staggered Difference-in-Differences design and measure segregation at three nested scales. Enrollment falls by approximately 4 percent, driven by the exit of students ineligible for free or reduced-price lunch. Declines are larger after fatal shootings. Enrollment also falls at neighboring schools, pushing the nearest safe alternative farther out. Socioeconomic segregation between schools and between districts remains unchanged, but segregation between counties increases by roughly 40 percent above baseline. A single shooting therefore widens the socioeconomic gap between neighborhoods, and the families least able to move are left behind.

Keywords— crime, education, segregation, school shootings
JEL— I24, I31, J11, R23

*Linnaeus University and the Research Institute of Industrial Economics (IFN). marcos.demetry@lnu.se.
Thanks to Max Coveney, Jacob Enlund, Hans Grönqvist, Rucker Johnson, Sasha Kapoor, Nikita Koptyug, Olivier Marie, Vahid Moghani, Erik Prawitz, Steven P Raphael, Roman Rivera, and Jonathan Thebes, and participants at the Erasmus University of Rotterdam PhD seminar, SOLE-JOLE 2026, EALE 2026, EEA-ESEM 2026, the Ratio Institute, SWEGPEC Workshop 2026, Umeå University NEWER Workshop, and Linneaus University for valuable feedback.

1 Introduction

Why do disadvantaged neighborhoods remain disadvantaged? The persistence of spatial inequality is among the most enduring problems of modern society. Roughly 70 percent of families living in concentrated-poverty neighborhoods in 1972 remained in similarly poor neighborhoods four decades later (Sharkey 2013). A well-established literature documents that childhood exposure to neighborhood disadvantage persistently reduces educational attainment and long-run earnings (Chetty and Hendren 2018; Chetty et al. 2016; Sampson 2012; Sharkey 2013). However, less is known about the dynamic mechanisms that sustain and deepen this localized disadvantage over time. One channel is that adverse neighborhood shocks degrade not only the immediate environment, but potentially also the socioeconomic composition of the residents who remain. Since high-income households face fewer mobility restrictions and have a higher marginal willingness to pay for amenities, localized shocks may trigger endogenous residential sorting that systematically traps low-income households in declining areas (Sharkey 2010, 2018).

This paper explores this sorting mechanism by exploiting school shootings as an exogenous shock to a central local amenity: the safety of the public school environment. I formalize this framing—safety as a priced neighborhood attribute, and a shooting as a sudden, unanticipated degradation of its value—in a spatial equilibrium model of residential choice. Standard residential sorting is driven by slow-moving, correlated disamenities such as gradual economic decline or rising local crime rates. A school shooting, by contrast, provides a highly salient and publicized shock whose timing is largely orthogonal to underlying neighborhood trends (Rossin-Slater et al. 2020; Yousaf 2021). This combination of exogeneity, salience, and precise temporal measurement provides a rare quasi-experimental setting to observe how a community's socioeconomic composition shifts in the immediate aftermath of an amenity shock.

A central methodological challenge in the neighborhood effects literature is that exposure to violence is endogenous. Standard violent crime is non-random, chronically underreported, and concentrated in areas undergoing simultaneous economic decline (Ang 2021; Caudillo and Torche 2014). As previously noted, school shootings provide a methodological exception. The rising incidence of school shootings is well documented (Reeping et al. 2022; Wippell et al. 2026). The 1990s saw roughly 30 shootings per year while the early 2020s had over 200 per year, more than six times the 1990s average. The K-12 School Shooting Database (Riedman 2024) records more than 2,300 incidents since the 1960s, roughly a quarter of them fatal. The scale of this deterioration is starkest at the extremes. The safest year in the past three decades saw 19 shootings (1990), while the worst saw 308 (2022), a more than fifteenfold increase, as illustrated in Figure 1. Crucially, while crime-related gun discharges are concentrated in disadvantaged and urban environments, indiscriminate school shootings occur disproportionately across whiter, more affluent, rural, and suburban districts (Baird et al. 2017; Levine and McKnight 2020). While certain school and district characteristics predict the latent risk of *where* a shooting occurs (Baird et al. 2017; Levine and McKnight 2020), the timing of an incident is fundamentally unpredictable by observable characteristics (García-Montoya et al. 2022; Rossin-Slater et al. 2020; Yousaf 2021). It is this conditional exogeneity in timing that I exploit as a quasi-natural experiment, utilizing a staggered difference-in-differences design to estimate the causal effect of school shootings on the long-run socioeconomic composition of affected schools and neighborhood sorting.

When a shooting happens, parents say they will leave. For example, Figure A.1 shows news clippings from a shooting in Oakland in 2022 that left one person killed and five others injured, writing that "parents are threatening to pull their students out of school" (NBC Bay Area 2022). Similarly, following the Uvalde shooting, there was news reporting of parents seeking out new schools in nearby districts (Scribner 2022). For families, this shock to the perceived safety and amenity value of the local public school (Towers et al. 2015) triggers a residential response that is highly asymmetric, governed by two complementary mechanisms. First, lower-income households face credit constraints and mobility frictions such as transaction costs and liquidity limits that make relocating to a safer district or substituting toward private education infeasible, while high-SES households face lower mobility frictions (Abouk and Adams 2013; Boustan 2013). Second, under the single-crossing property typical of spatial equilibrium models, the marginal willingness to pay for neighborhood safety rises monotonically with income (Epple and Sieg 1999; Epple and Romano 2003). Imposing these properties implies that the community sorts at a single income threshold, with every household above the cutoff having both the means and the motive to leave while those below it stay. The benefit of this simplification of the sorting structure is that it yields signed predictions. One can model richer monotonic relationships but single-crossing is enough for the model to generate a testable hypothesis that exit concentrates among higher-income households. Wealthier households substitute toward safer districts, while lower-income households remain structurally trapped. Consequently, a shooting does not merely reduce aggregate enrollment, it mechanically increases the concentration of low-income students (Yang and Gopalan 2023). Furthermore, this initial compositional shift may trigger a self-reinforcing amenity cascade. As affluent households exit, they contract the local tax base, depress peer effects, and erode the aggregate enrollment required to sustain school operations (Smith and Zimmer 2022). In this way, a discrete violent event acts as a catalyst for long-term spatial inequality.

The empirical evidence on the mobility response to localized violence remains mixed. At the individual level, exposed students show no increase in school switching or exit (Cabral et al. 2026), and over the longer run they move less and are less likely to reach higher-income areas (Sezer 2022); at the school level, enrollment falls markedly (Beland and Kim 2016; Muñoz-Morales and Singh 2023; Yang and Gopalan 2023). The two sets of results need not conflict, because exit and composition are different objects. A school can lose students because families stop arriving, not only because the families already there leave, and an empirical design that follows exposed students will not capture the first margin. What the evidence leaves open is the geographic scale at which the compositional effects emerge. Since spatial mobility frictions are strictly increasing in distance, households will theoretically re-sort at the minimum geographic radius required to secure a safer educational amenity within their budget constraint. Therefore, estimating mobility responses at an arbitrary or singular geographic scale may yield spurious zero effects. Furthermore, while the macro-trend of rising residential socioeconomic segregation is well-documented (Owens et al. 2022), the micro-spatial dynamics driving this phenomenon remain underexplored. Although recent work suggests that localized shocks can alter student composition at the school level (Burdick-Will et al. 2024; Burdick-Will et al. 2021; Caudillo and Torche 2014; Hailey 2025), building on canonical evidence of violence-induced residential flight (Morenoff and Sampson 1997), the aggregate geographic scale at which these shocks ultimately generate spatial segregation is not known ex ante. Identifying this relevant spatial scale of equilibrium sorting remains a central empirical challenge.

To resolve this empirical ambiguity, I use the Common Core of Data (a panel covering the universe of U.S. public K-12 schools) from 1990 to 2022 and link it to the K-12 School Shooting

Database (Riedman 2024). By exploiting the precise geolocation of each incident, I map each shooting to its focal school and its respective district and county boundaries. This allows me to trace compositional shifts through three nested geographies, rather than relying on an arbitrary level of aggregation. Since each relocation margin (within-district, inter-district, and inter-county) implies a distinct spatial equilibrium, I calculate segregation measures at all three margins. Which of these margins binds depends on the geographic scale at which a genuinely safe educational substitute becomes available. Empirically, I document that a shooting reduces enrollment not only at the affected school but also at neighboring schools within an approximate 10-kilometer radius. This indicates that intra-district and nearby schools are contaminated by the same localized shock via spatial spillovers (Deb and Gangaram 2024; Levine and McKnight 2020). Consequently, the nearest genuinely safer alternative typically lies further geographically, perhaps across a county line. This dynamic is consistent with endogenous re-sorting occurring primarily between counties, resulting in null compositional shifts at localized within-district and within-county margins. To detect this re-sorting, I measure shifts in socioeconomic composition. Since the single-crossing property implies that flight is selected on income, I partition students at the federal Free or Reduced-Price Lunch (FRPL) eligibility threshold (185% of the federal poverty level), separating lower- from higher-income households. Yousaf (2021) demonstrates that, conditional on population, counties experiencing mass shootings are statistically indistinguishable from those that do not across a rich set of political, demographic, economic, and crime characteristics; therefore, their occurrence cannot be predicted using local covariates. Consistent with evidence that school shootings are similarly unpredictable based on observable school characteristics (Rossin-Slater et al. 2020), I treat the exact timing of an incident conditional on time-invariant location characteristics absorbed by unit fixed effects as the source of identifying variation, and I estimate its causal effect using a staggered difference-in-differences framework.

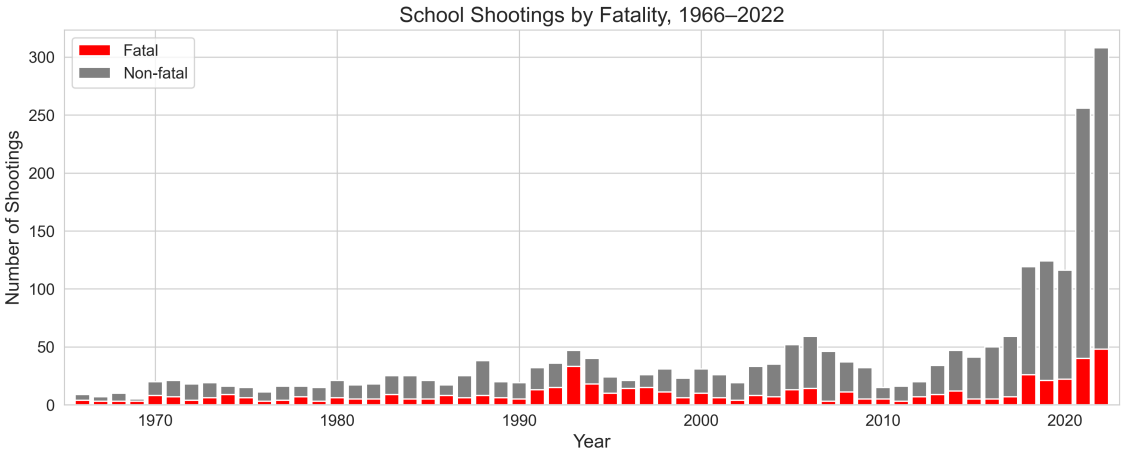
I find evidence of macro-spatial residential relocation rather than intra-district school switching. Consistent with recent literature, I observe that school-level enrollment declines by approximately 4 percent following a shooting, a reduction driven by the exit of students ineligible for free or reduced-price lunch (FRPL). Consequently, the concentration of low-income students in affected schools rises. However, I find no evidence of sorting by socioeconomic status at the within-district or between-districts margins. Instead, segregation manifests exclusively between counties: the normalized exposure index for FRPL versus non-FRPL students rises by approximately 0.022 within states (or 40% above baseline). This spatial sorting effect accrues gradually over three to five years, a dynamic consistent with the mobility frictions and transaction costs associated with cross-county residential relocation. Although the data restrict us to observing aggregate compositional shifts rather than individual household migration trajectories, the emergence of socioeconomic segregation exclusively at the inter-county scale implies that the observed enrollment declines are driven by residential sorting.

This paper makes three primary contributions to the literatures on the exposure to violence and neighborhood effects. First, I advance the empirical literature on localized violence. While existing research documents aggregate enrollment declines following school shootings (Beland and Kim 2016; Yang and Gopalan 2023), I identify the precise geographic scale at which this flight translates into socioeconomic segregation. Several recent studies find no evidence of out-migration following a shooting, whether in individual-level school-switching and out-of-state migration records (Cabral et al. 2026) or in aggregate county- and state-level population flows (Brodeur and Yousaf 2025; Deb and Gangaram 2024). By estimating effects across nested spatial margins, I show that these apparent null effects are an artifact of geographic scale: tracking compositional

shifts between free or reduced-price lunch (FRPL) and non-FRPL students across nested boundaries reveals that residential sorting does occur, but across county lines rather than within districts. Second, I offer a methodological contribution to the broader literature on residential sorting. A fundamental empirical challenge in spatial equilibrium models is that household preferences and neighborhood amenities are simultaneously determined (Bayer et al. 2025; Bayer and McMillan 2012). Since standard sorting is driven by endogenous, slow-moving disamenities, isolating causal sorting behavior from underlying neighborhood decay is difficult. By exploiting school shootings as exogenous, unanticipated shocks to a single public amenity—school safety—we bypass this endogeneity, and are able to identify the sorting mechanisms that are typically obscured by the endogeneity of choice and place. My main analysis uses all shootings. One might worry that this broad definition inherits the underreporting that plagues crime data more generally, since minor incidents are recorded unevenly across places. To address this, I re-estimate the analysis on fatal shootings alone which do not suffer from underreporting and find the same results qualitatively. Finally, I contribute a theoretical synthesis of how acute, localized shocks generate long-run spatial inequality. To better understand the mechanisms, I introduce a spatial sorting model governed by a single income threshold—the point at which a household's marginal willingness to pay for school safety intersects with asymmetric mobility frictions. This mechanism is rooted in the single-crossing property of standard sorting models (Epple and Sieg 1999; Epple and Romano 2003), and serves as an organizing device rather than a structural estimator. Its comparative statics explain the aggregate enrollment decline, its concentration among non-FRPL students, and the heightened response to fatal events. However, the county-level scale of the segregation is not derived from the threshold alone; it follows only once I add the assumption that neighboring schools are contaminated by the same shock through spatial spillovers, pushing the nearest genuinely safe substitute further away. If there were no spillovers, we may observe sorting on smaller geographic scales. Ultimately, by combining this theoretical intuition with my nested-geography design, I provide evidence for how discrete amenity shocks interact with households' budget constraints to permanently trap low-income households in declining neighborhoods.

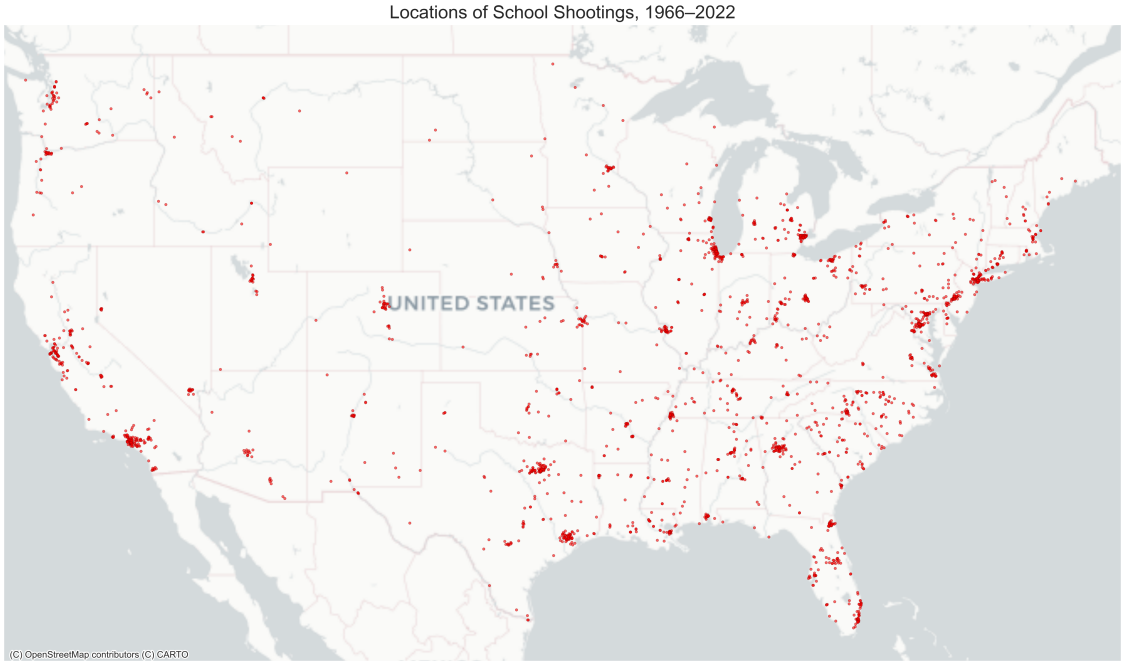
The remainder of the paper is organized as follows. Section 2 describes the U.S. school system and open-enrollment regimes, the school shootings setting, and reviews the literature on shootings and mobility. Section 3 details the data on school shootings, student enrollment, and the construction of the nested segregation indices. Section 4 develops a simple spatial sorting model governed by a single income threshold that organizes the empirical predictions. Section 5 outlines the staggered difference-in-differences identification strategy. Section 6 presents the main results: the school-level enrollment decline, the shift in school socioeconomic composition, and the socioeconomic sorting that emerges across geographic scales. Section 7 examines mechanisms and treatment-effect heterogeneity, including spatial spillovers, the fatality of the shooting and the underlying level of violence, and the ease of switching schools. Section 8 reports robustness checks, and Section 9 concludes.

Figure 1. School Shootings per Year by Fatality



Notes: Source: K-12 School Shootings Database (Riedman 2024). Red bars indicate fatal shootings.

Figure 2. Shootings Heatmap



Notes: Shootings heatmap for contiguous United States, 1970–2023, based on data from (Riedman 2024).

2 Background

School and Residential Choice in the U.S.

How families choose schools depends on many factors. Broadly, there are three schooling alternatives: public schools, private schools, and homeschooling. Within public schools are both traditional and charter or magnet schools. Magnet schools offer specialized programs, often for gifted students. Students are assigned to public schools based on where they live, or the catchment area that their home resides in. The overwhelming majority, around 70% of U.S. students, attend their assigned neighborhood school (Wang et al. 2019). For example, Figure 3 shows school district boundaries in thick purple outlines and catchment areas within the districts in thinner lines. The remaining 30% are distributed between private schools, public schools outside of their neighborhood, or homeschooling. Non-sectarian private schools are expensive (Broughman et al. 2021) and can cost up to two to three times more than the average Catholic school. Magnet schools are highly sought after, and admissions are often determined by lottery. Therefore, choice of school is highly interlinked with the choice of residential location. Reaching a different public school normally means moving house, and moving is costly. The geographic scale at which sorting surfaces after a shooting therefore reveals who could afford to leave.

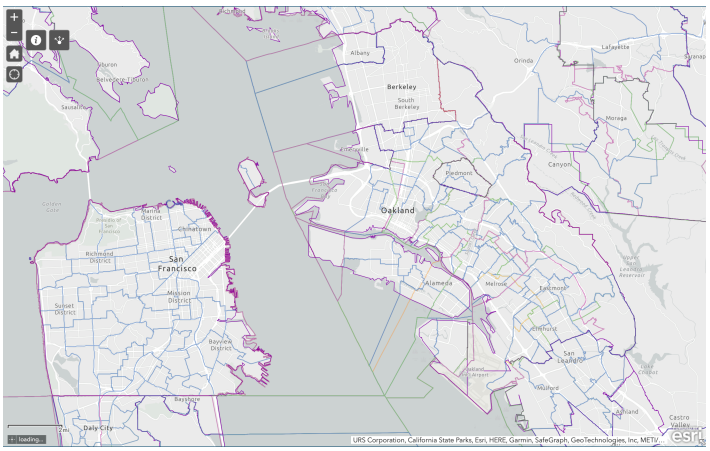
School choice options differ between states. Open enrollment is a form of school choice where students can switch to other traditional public schools outside of their catchment areas. Depending on state legislation, students may be permitted to switch schools within the district or between districts. In 21 of 50 states, within-district transfers are permitted. The majority of states, 43 of 50, permit between-district transfers. Even in states with strict open enrollment rules, there are often exemptions for certain groups such as children with disabilities or with a parent serving in the military (Erwin et al. 2022).

A school district is an organizing body that governs the schools under its control. A district can contain one or several schools of different grades, a county can contain one or several districts, and there are several counties in a state. In 2015–16, the U.S. had roughly 98,000 public schools in about 18,000 districts and 3,000 counties. A geographic school district controls all the schools within its border. There are also administrative school districts, often in control of charter or magnet schools outside of their own geographic borders. In this paper, we focus on geographic school districts as they are most closely linked to the residential relocation channel. The number of same-level schools that are inside a district reflect the number of choices that a family has in order to switch schools without leaving the district. The median district runs three schools, and the mean runs six schools.

School shootings

When does a shooting count as a school shooting? The prototypical image of a school shooting is a non-targeted or indiscriminate mass shooting involving a semi-automatic rifle. However, school shootings often have broader definitions. In this paper, following Riedman (2024), I define a school shooting as any event where a gun was brandished or fired, or bullets hit school property, irrespective of time or day. The details of this definition will be covered in Section 3 on Data below. Most shootings occur outside of the school building, for example in the parking lot. They are often

Figure 3. School Districts and Catchment Areas in the Upper Bay Area



Notes: School district boundaries for 2021–2022 from NCES, overlaid with public-school catchment boundaries (light blue) for 2015–2016 from the NCES Education Demographic and Geographic Estimate program.

the result of escalated disputes and are very rarely a targeted murder or assassination. They occur most likely during morning classes. In 95% of the cases, the perpetrator is male. The highest recorded number of school shootings are in Chicago, Illinois. Other hot spots include Detroit Michigan, Baltimore Maryland, Los Angeles California, Philadelphia Pennsylvania and Houston Texas. Shootings are most likely to occur in densely populated urban areas. School shootings are also more likely to occur in bigger schools, often high schools, but result in greater fatalities in elementary schools.

The label "school shooting" pools two phenomena. The first is the rarer but more publicized indiscriminate rampage shootings which fall disproportionately on whiter, more affluent, rural, and suburban districts (Baird et al. 2017; Levine and McKnight 2020) such as Columbine (1999), Sandy Hook (2012), Parkland (2018), Santa Fe (2018), and Uvalde (2022). The second is everyday gun violence at schools, and is common in urban areas and concentrated in disadvantaged communities. A reader familiar with the mass shooting literature may expect treated schools to be of certain characteristics. However, a broad definition includes both of these types of school shootings. I pool the two because a shooting of either kind degrades a priced local amenity, the safety of the school, and a family weighs the loss against the cost of leaving. The broad definition spans the full severity distribution, so the later split into fatal and non-fatal events isolates the severe subset and proxies the size of the shock, the parameter σ of the model in Section 4.

Previous literature on the consequences of school shootings

Previous literature establishes that school shootings have severe and lasting effects on the students exposed to them. Beland and Kim (2016), comparing a shooting school to other schools in the same district in a difference-in-differences design, find a decline in average test scores following a fatal shooting. Levine and McKnight (2024) in event-study case studies of the three deadliest school shootings (Columbine, Sandy Hook, and Parkland) also find negative effects on test scores. Keeping in mind that households of higher socioeconomic status value school quality more (Abdulkadiroğlu et al. 2020; Bjerre-Nielsen and Gandil 2024; Burgess et al. 2015), we can expect higher socioeconomic status families to be more likely to move or switch schools following a

shooting. Cabral et al. (2026), matching exposed Texas students to students in nearest-neighbor matched control schools in other districts, find higher chronic absenteeism, more grade repetition, and lower high-school graduation. Youth antidepressant use and stress-related emergency-room visits rise within five miles of a fatal shooting, relative to ten to fifteen miles away (Gujral et al. 2023; Rossin-Slater et al. 2020). The damage concentrates in fatal and indiscriminate events. Non-fatal shootings leave antidepressant use unchanged (Rossin-Slater et al. 2020), and suicidal shootings leave enrollment and test scores unchanged (Beland and Kim 2016).

The effects of shootings on enrollment are mixed, depending on the level of analysis and on which type of exit one looks at. Public-school enrollment at the shooting school falls. Beland and Kim (2016) find that ninth-grade enrollment drops about six percent after a fatal shooting. The decline is concentrated in open-enrollment schools and in larger schools. Yang and Gopalan (2023), in an event study with state-by-year fixed effects on district and school panels, find a decline in enrollment of around four to five percent. Muñoz-Morales and Singh (2023) find a similar decline in enrollment using a stacked design that compares households living just inside a shooting school's attendance boundary to households living just outside.

At the individual level, the mobility results are mainly null. Cabral et al. (2026) find no statistically significant increase in school switching or exit from the Texas public system. Deb and Gangaram (2024), using survey respondents linked to county-year exposure in non-linear cohort models, find no differential out-migration in Census data. Sezer (2022), following cohorts in the PSID with an age-based difference-in-differences, finds that students exposed to a shooting leave their neighborhood less at age 30, and are less likely to move to higher-income districts than unexposed peers. Interestingly, Sezer (2022) finds that children to exposed individuals also have lower upward mobility.

The effects on the private-school margin also differs by level of analysis. Abouk and Adams (2013), in a state-year panel, find private high-school enrollment increases by about ten percent after a shooting while public high school enrollment decreased marginally. Measured locally, the sign reverses: near a shooting, private schools lose enrollment along with public ones (Muñoz-Morales and Singh 2023; Yang and Gopalan 2023).

How mobility responses sort students along socioeconomic lines has not been studied. Yang and Gopalan (2023) hypothesise that differential mobility could widen socioeconomic segregation but do not test the claim directly. They find enrollment falls in the shooting district and that the students who remain are more likely to be eligible for free or reduced-price lunch, and they infer that socioeconomic segregation across schools rises.

Most previous estimates come from a single event, like Columbine, Sandy Hook or Parkland in Levine and McKnight (2024), or from a matched sample within a single state (Cabral et al. 2026). Other studies pool many shooting events but estimate with two-way fixed effects, which can be biased under staggered timing with heterogeneous effects. Several of these studies also suffer from spillover bias. Beland and Kim (2016), for example, compare a shooting school to other schools in the same district, but families leaving the shooting school may enroll in those other schools, attenuating the estimate toward zero, as the authors point out.

Studies also define school shootings in various ways, which further complicates comparison. Some restrict to mass shootings only: Muñoz-Morales and Singh (2023) use shootings with more than three victims, and García-Montoya et al. (2022) use rampage shootings drawn from a Wikipedia list of multiple-victim events. Such samples capture the deadliest events but miss the low-level gun violence that most exposed students actually face, which Cabral et al. (2026) argue is the ordinary case. Broader counts that record every discharged firearm carry the opposite risk:

Levine and McKnight (2020) show that non-fatal incidents are reported more often in urban, high-media areas, so the sample correlates with race and income. In the next section, I describe the data in my setting.

3 Data

The most comprehensive dataset is the K-12 School Shooting Database (abbreviated K12SSD, Riedman 2024). In this dataset, an event counts as a school shooting if a gun was brandished or fired, or bullets hit school property, irrespective of time or day. These amount to approximately 2,600 school shootings, from 1966 to 2024. In my study period, 1990–2022, there are approximately 2,300 shooting incidents. This data is gathered continuously and extensively from government agencies, media or advocacy groups, websites, and blogs. Several other datasets, like the Washington Post School Shootings Database, are therefore a subset of this database. Each incident contains information about the perpetrator, the setting, and is tagged by its geolocation in latitude and longitude. I aggregate the data up to the school-year level. There are some schools that experience multiple shootings in one year. The relevant numeric variables, such as number of fatalities, are summed within year. I summarize the many datasets for school shootings in Appendix B.

I rely on two sources for school-level information. For administrative data, I use the Common Core of Data (CCD) by the National Center for Educational Statistics (NCES).[1] This covers all public elementary and secondary schools in the United States. The CCD contains school identifiers (such as NCESSCH codes, Local Education Agency code, latitude and longitude, etc.). It also contains information on the number of enrolled students in each grade and the number of Free or Reduced Lunch Recipients (a proxy measurement for number of poor students). Students from households with incomes at or below 130% of the federal poverty level are eligible for free lunch. For students from households earning between 130% to 185% of the federal poverty level, a reduced price lunch is available (U.S. Department of Agriculture, Food and Nutrition Service 2025). Students above these thresholds that are not in other lunch programs purchase a paid meal according to prices set by their school district. Data on the racial composition of schools and other socioeconomic indicators come from The Educational Opportunity Project at Stanford (Reardon et al. 2024).[2]

Data at the school-district level also come from the Stanford project. At this level, there are measures of segregation, such as exposure indices between poor and non-poor students. This data therefore allows us to indirectly measure selection into schools after a shooting. I create segregation indices at three geographic scales: between schools within a district, between geographic school districts within a county, and between counties within a state. The segregation analysis runs on the annual SEDA panel for 2009–2019, the years over which the segregation indices and the SEDA covariates that enter the estimator are available. Because the underlying enrollment counts are annual, event time e is measured in calendar years on a balanced yearly panel. The estimation window is set to ± 7 years, so the longer-horizon dynamic effects are identified off the earlier-treated cohorts that have enough post-shooting calendar years within 2009–2019.

1. As of June 2025, data can be accessed here: https://nces.ed.gov/ccd/.

2. The segregation and socioeconomic measures are built from the Common Core of Data which SEDA aggregates from the school level to districts, counties, and states (Reardon et al. 2024).

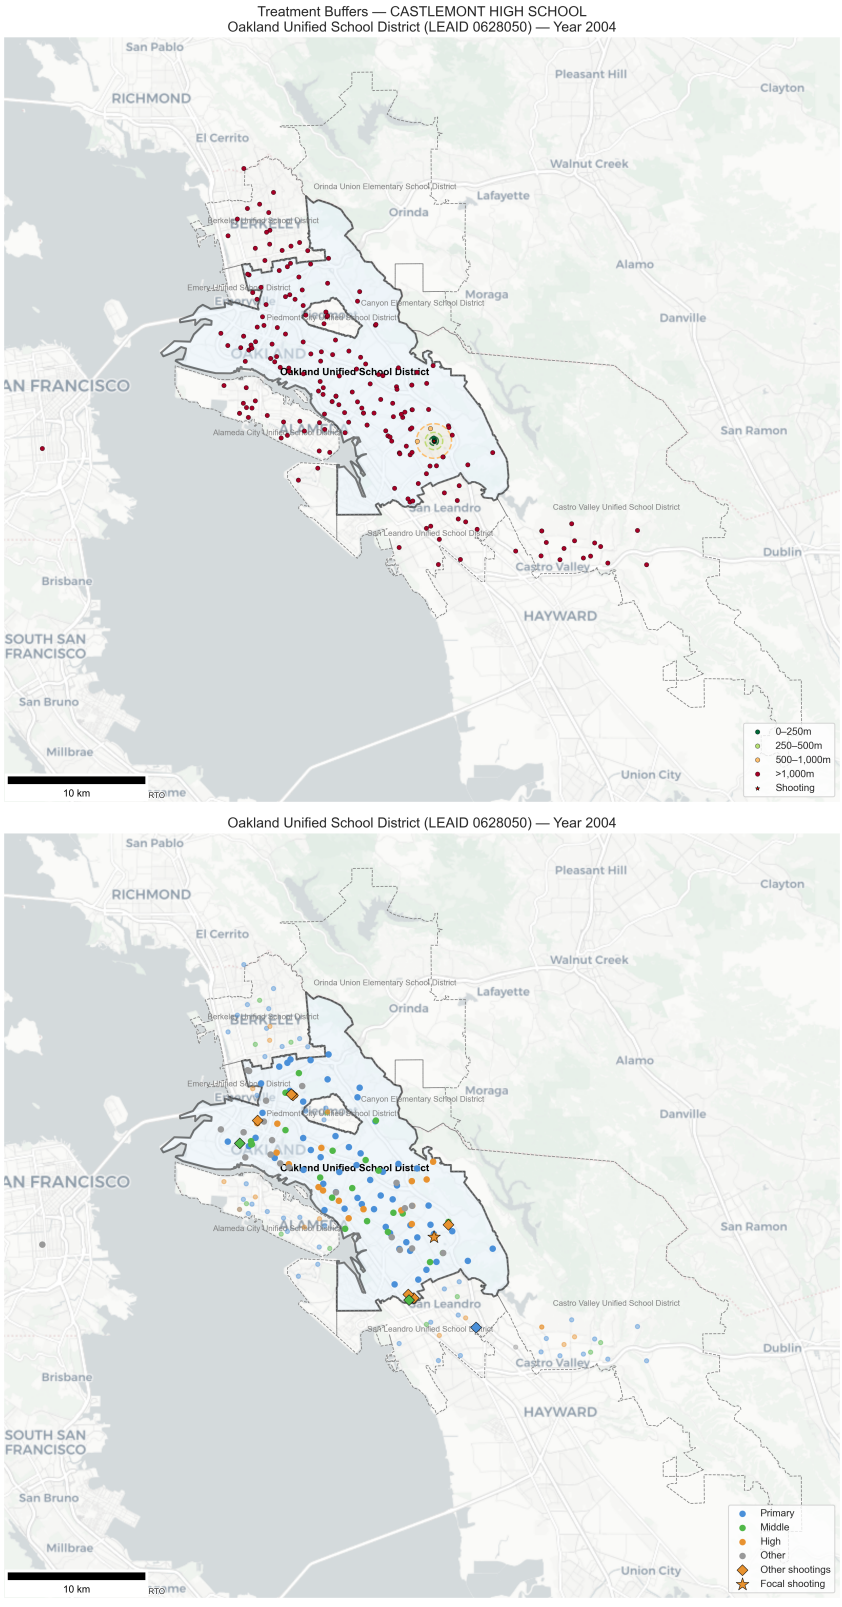
The main segregation outcome is the Normalized Exposure Index for Free Lunch to Non-Free Lunch recipients (Reardon and Firebaugh 2002). This two-group index measures the extent to which Free Lunch recipients are concentrated in schools with a different socioeconomic composition than Non-Free Lunch recipients, normalized to adjust for changes in the overall share of Free Lunch recipients in the unit. For subunits j nested within a wider area, let π_j be the share of Free Lunch eligible students in subunit j, Π the corresponding share in the area as a whole, m_j the enrollment of subunit j, and $M = \sum_j m_j$ the enrollment of the area. The index is

$$N \;=\; \frac{\sum_j \left(m_j/M\right)\left(\pi_j - \Pi\right)^2}{\Pi\left(1-\Pi\right)}, \tag{1}$$

the enrollment-weighted variance of subunit Free Lunch shares about the area mean, divided by its theoretical maximum $\Pi(1-\Pi)$. It is equivalently known as the variance ratio index, or η^2. A value of 0 indicates no segregation, meaning every subunit has the composition of the area as a whole, and a value of 1 indicates complete segregation. Since the normalization removes mechanical shifts driven by changes in the proportion of poor students, it isolates the distributional sorting component of segregation.

As shown in Figure 4, I perform a spatial join of the school and shooting data. The top panel shows my preferred treatment radius of 500m as well as two other radii. A robustness and spillover analysis is performed with other radii later in this paper, and Appendix E illustrates every treatment assignment used. The bottom panel shows the shooting school (star marker) with other schools color-coded based on their school-level in the same district and in neighboring districts. In this example, there are quite a few alternative high-schools not too far from the focal school.

Figure 4. School Districts and Shootings in the upper Bay Area



Notes: Zoomed in picture showing a school shooting in the focal point and 250m, 500m, and 1000m concentric rings. (Left). This school is also shown in its district and the surrounding districts. (Right)

Table 1 summarizes the sample at each scale of analysis. The school panel is built on the Common Core of Data and spans 1990–2022, with 2,504 treated schools entering the enrollment analysis and 2,450 entering the free/reduced-lunch-share analysis. The number of treated schools is greater than the number of shooting schools in the K12SSD precisely because of the spatial join that allows multiple schools within the treatment radius to be treated by the same school shooting event. The district and county panels are built on the SEDA data and so are confined to the SEDA observation window of 2009–2019, with 353 treated districts. The county between-district segregation analysis is observed across the longer panel but identifies effects only off treated cohorts whose first shooting falls inside the SEDA window, leaving 184 treated counties. Unit-year counts refer to treated observations with a non-missing outcome. The appendix Figure D.1 plots coverage by variable and year in detail.

Table 1. Sample Composition by Analysis Scale and Outcome Family

Scale	Outcome family	Window	Units	Treated units	Treated unit-years
School	Enrollment (CCD)	1990–2022	149,908	2,546	65,623
School	Free/reduced lunch share (CCD)	1998–2022	128,922	2,491	47,162
District	Poverty (SEDA)	2009–2019	12,798	353	3,881
District	Within-district segregation (SEDA)	2009–2019	12,797	353	3,881
County	Between-district segregation (SEDA cohorts)	2009–2019 cohorts	3,152	184	5,876

Notes: Distinct units, treated units, and treated unit-years entering each analysis, by scale and outcome family. The school panel is built on the Common Core of Data (CCD); the district and county panels on the Stanford SEDA data. "Treated" units are those ever experiencing a shooting; counts are over observations with a non-missing outcome. The county between-district segregation analysis restricts to treated cohorts whose first shooting falls in the SEDA window 2009–2019.

4 A Model of Safety, Mobility and Sorting

A shooting degrades the safety of a school. This is a loss in a priced local amenity that families weigh against the cost of leaving. In this section, I show that a minimal sorting model can produce predictions about whether enrollment falls, who leaves, and at what geographic scale re-sorting appears. The model builds on two key ingredients. The first is the single-crossing property of the canonical Epple–Sieg sorting model (Epple and Sieg 1999): the willingness to pay for a neighborhood amenity rises with household socioeconomic status (SES). The second is a moving cost that is *regressive* (lower for richer households) and *increasing in geographic scale* (switching schools is cheaper than moving across a county line). Treating school safety as the amenity and a shooting as a negative shock to it, the model collapses to a single threshold condition with clear comparative statics.

4.1 Setup

Consider a focal school in a location j. Households are indexed by a scalar SES type y, distributed according to $F(y)$, which we proxy empirically by Free/Reduced-Lunch (FRL) status. A household chooses a location j that bundles a *safety amenity* a_j, a cost of residing there p_j (rent or

the capitalized house price), and an idiosyncratic taste shock ε_{ij}. We take indirect utility to be quasi-linear in the numeraire:

$$V_{ij} \;=\; \alpha(y_i)\, a_j \;-\; p_j \;+\; \varepsilon_{ij}. \tag{2}$$

The only structural restriction is **single crossing**: the marginal willingness to pay for safety is strictly increasing in SES,

$$\alpha'(y) \;>\; 0. \tag{3}$$

This is the Epple–Sieg monotonicity of the marginal rate of substitution between price and amenity, specialized to safety as a single amenity. At the margin, richer families will pay more for safety than poorer families. This single-crossing assumption is a standard primitive of the structural sorting literature.

This restriction aligns with a consensus in the urban economics and economics of education literature that willingness to pay for school quality and neighborhood amenities is highly heterogeneous and strictly increasing in SES. Theoretically, this gradient arises because educational quality and safety are normal goods, so demand scales with endowed income (Epple and Romano 2003). Empirically, both hedonic pricing and discrete residential choice models consistently show that higher-income and more highly educated households have a significantly steeper marginal willingness to pay for academic quality and favorable neighborhood sociodemographics (Bayer et al. 2007; Black 1999; Burgess et al. 2015). Recent evidence exploiting centralized school assignment mechanisms reveals that this gradient is driven primarily by preference heterogeneity over peer composition. High-SES families are substantially more sensitive to the socioeconomic and academic profile of a school's student body than their lower-SES counterparts (Abdulkadiroğlu et al. 2020; Bjerre-Nielsen and Gandil 2024; Hastings et al. 2009). Because advantaged households have stronger preferences for peer quality and face fewer binding credit constraints, their bids set the marginal price for these amenities, capitalizing their value into local housing markets (Caetano 2019).

The shock. I introduce a shock to the focal school's safety of size $\sigma > 0$:

$$a_j \;\longrightarrow\; a_j - \sigma.$$

We can also index the severity of the shock, with $\sigma_{\text{fatal}} > \sigma_{\text{non-fatal}}$.

The friction. Why don't all families react to this negative shock, regardless of income or of the alternatives available to them? Because relocating is costly. We add a relocation friction to the framework. Reaching a *safe substitute*—a location whose safety has not itself been degraded by the same event—costs $c(y, d)$, where d indexes the geographic scale of the move. Two properties matter:

$$\text{(i) regressiveness:} \qquad \partial c/\partial y \;<\; 0, \tag{4}$$

$$\text{(ii) scale-increasing:} \qquad c(y, \text{switch school}) \;<\; c(y, \text{move within district}) \;<\; c(y, \text{cross county}). \tag{5}$$

Property (i) captures credit constraints, liquidity, and the lower relative burden of transaction and search costs for the affluent (the "regressive relocation frictions" of the dynamic-sorting literature). Property (ii) is geographic: the cost of relocating rises the farther one moves. A within-district

school transfer requires no residential move, whereas crossing a county line may require selling, searching, and relocating. Let d^* denote the cheapest scale at which a safe substitute actually exists.

4.2 The Move-Out Threshold

A household exits the focal location whenever the privately valued safety loss exceeds the cost of reaching the nearest safe substitute:

$$\alpha(y_i)\,\sigma \;>\; c(y_i, d^*)\,. \tag{6}$$

The left-hand side is increasing in y by (3); the right-hand side is non-increasing in y by (4). Their single crossing defines a cutoff type y^*,

$$\boxed{\alpha(y^*)\,\sigma \;=\; c(y^*, d^*)} \qquad \Longrightarrow \qquad \text{household exits} \iff y_i > y^*. \tag{7}$$

Families above the SES cutoff leave and families below it stay. Equation (7) is the formal counterpart of the three behavioral margins in the sorting literature: moving out, staying, and moving in. Affluent families, for whom (6) holds, move out, because the loss from staying in location j is greater than the cost of moving to the nearest safe alternative. Poorer families, for whom (6) fails, are stayers—and it fails for them from both sides: single crossing (3) makes their valuation of the lost safety small, while the regressive friction (4) makes their cost of reaching a substitute high. The third margin, moving in, works through the price channel. As affluent households exit, demand for the focal location falls and p_j drops. This is consistent with empirical findings that local property values decline following a school shooting (Muñoz-Morales and Singh 2023). A lower p_j in (2) partially offsets the safety loss for low-y households, drawing in poorer in-migrants and accelerating the compositional shift. I do not estimate the price response directly.

4.3 Comparative Statics

From this simple model, we can make a number of predictions. First, we expect that (P1) log enrollment falls. The share of families exiting the focal location after the shock is $1 - F(y^*) > 0$. Second, we expect that (P2) leavers are non-poor because exit is from the upper tail $\{y > y^*\}$, the families who leave are drawn from the top of the local SES distribution. The compositional effect then is that the FRL share among stayers rises. Third, we expect that (P3) severity of the shooting affects the magnitude of the effect. Differentiating (7) gives $\partial y^*/\partial\sigma < 0$. Thus, a larger shock lowers the cutoff and more families clear the move-out condition. A fatal school shooting should therefore produce a larger exit than a non-fatal one. Fourth, the (P4) greater the number of alternatives or the higher the ease of switching schools, the greater the decline in enrollment will be. The more same-level schools there are in a district, the lower the cost of switching at small distances, which lowers the threshold and increases the outflow from the focal school. The same logic applies where intradistrict open enrollment is available: the easier it is to switch schools, the larger and less selective the exit. The prediction is monotone, i.e. cheaper switching means a larger, less selective outflow. Fifth, (P5) re-sorting differs depending on geographic scale. What matters for *where* re-sorting appears is the scale d^* at which a genuinely safe substitute exists.

Panel A in Figure 5 shows how at the intersection of the privately valued safety loss equals the cost of relocating to a safe substitute defines the threshold y^*. The greater the shock, the steeper the $\alpha(y)\sigma$ line is, the greater the share of exiting families will be.

Panel B in Figure 5 shows how the threshold y^* is lower for between-school scale and higher for between district and between county. This is because, if within-district alternatives are contaminated by spillovers from the same incident, they are not safe substitutes, and the move-out condition (6) fails at $d =$ school and at $d =$ district. The nearest safe substitute, d^*, may then lay farther out.

4.4 The Predicted Effects on Segregation

Let segregation be summarized by the dispersion of poverty (FRL) shares across neighborhoods, the object empirically captured by the Normalized Exposure Index. Write π_j for the FRL share of the focal unit and Π for the FRL share of the wider area within which segregation is measured. The index is an enrollment-weighted variance of the π_j about Π (Equation (1)), so a marginal rise in the focal unit's own FRL share moves measured segregation in the direction of $\text{sign}(\pi_j - \Pi)$. Selective SES exit raises π_j, and it moves measured segregation in proportion to how *much* mass leaves and how *selected* that mass is:

$$\Delta\text{Seg}_{\text{county}} \;\propto\; \underbrace{\text{sign}\big(\,\pi_j - \Pi\,\big)}_{\text{position of the focal unit}} \;\times\; \underbrace{\big[\,1 - F(y^*)\,\big]}_{\text{mass that exits}} \;\times\; \underbrace{\big(\,\mathbb{E}[\,y \mid y > y^*\,] - \mathbb{E}[\,y\,]\,\big)}_{\text{SES selectivity of exiters}}. \qquad (8)$$

The last two factors set the magnitude. Both rise with severity σ and with the availability of safe substitutes (lower c), so the same forces that reduce enrollment also move segregation. The first factor sets the direction, and it is what makes the prediction conditional. Exit of the affluent from a unit that is already poorer than the wider area pushes that unit further from the mean and raises dispersion; the same exit from a unit that starts more affluent moves it *toward* the mean and lowers dispersion. The condition is satisfied in this setting: treated units sit on the poorer side of the distribution they belong to, with a pre-treatment mean FRL share of 0.53 among treated schools against 0.50 among controls (Table 3). This follows from the broad shooting definition adopted in Section 2. Indiscriminate rampage shootings do fall disproportionately on affluent districts, but they are the rarer of the two phenomena pooled here, and the modal event is everyday gun violence in poorer, urban schools. This force is restricted, however, to the scale at which safe substitutes are reachable. We can call this a segregation–inequality feedback.

4.5 Taking the Model to the Data

The model is not for structural estimations, but rather to guide the analysis. Still, I point out three caveats. First, the empirical design is reduced-form and identifies the comparative statics of (7) but not the general-equilibrium price response. Second, the binary FRL proxy is coarser than the model's continuous y and so cannot speak to mobility within the SES distribution. Third, the prediction of no between-school re-sorting rests on the auxiliary assumption that within-district alternatives are not safe substitutes. This is what I later find in a spillover analysis (Section 7.1). The same threshold can be derived, rather than assumed, from a standard constrained-optimization problem in which a household chooses consumption and location and safety is priced through a hedonic schedule; single crossing then emerges as a result rather than being imposed.

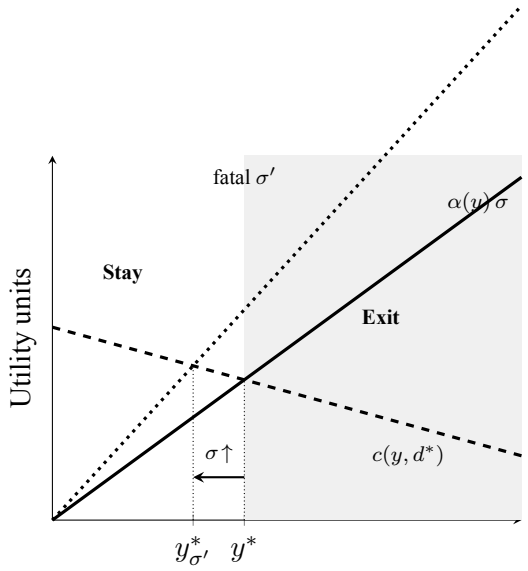
I use this single-threshold model to predict effects on school-level enrollment and SES composition, and on SES segregation at different geographic scales. I perform heterogeneity analysis by the number of alternative schools a household faces and by the open-enrollment rules that make switching easier or harder (Section 7.3). I also estimate the effects on enrollment for fatal and non-fatal shootings separately, to assess heterogeneity by the size of the negative amenity shock (Section 7.2).

4.5.1 Summary of predictions

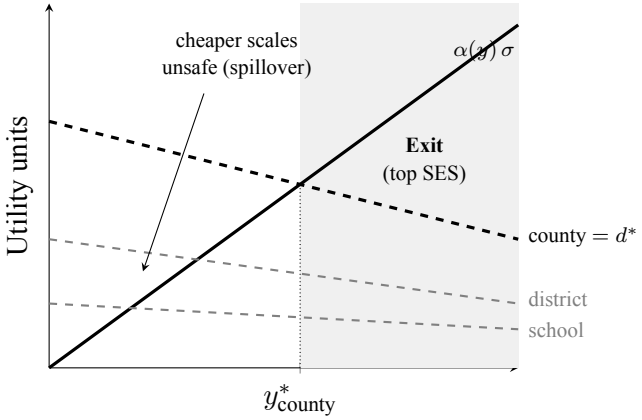
Table 2. Comparative statics of the move-out threshold

Comparative static	Mechanism
Exit share $1 - F(y^*) > 0$	Safety loss exceeds moving cost for high-y
Exit is the upper SES tail	Single crossing $\alpha'(y) > 0$
$\partial y^* / \partial \sigma < 0$	Bigger shock, lower cutoff
Lower c at small $d \Rightarrow$ lower y^*	More/cheaper alternatives
Re-sorting only at scale d^*	Neighbors unsafe (spillovers); safe substitute across county line
High c at county scale	Slow relocation
$p_j \downarrow$ as affluent exit	Compensating differential; predicts low-SES in-migration

Figure 5. The Move-Out Threshold and the Nested-Scale Cost Function



(a) The threshold y^* and severity



(b) Nested scales: d^* at the county line

Notes: Panel (a): the benefit of moving, $\alpha(y)\sigma$, rises with SES (single crossing) while the cost of reaching a safe substitute, $c(y, d^*)$, falls with SES; their intersection is the cutoff y^*, above which families exit. A more severe shock σ' (fatal) steepens the benefit line and shifts the cutoff left to $y^*_{\sigma'}$, enlarging the exit set. Panel (b): the moving cost increases in geographic scale (school $<$ district $<$ county). Because within-district substitutes are spillover-exposed and hence unsafe, the cheapest *safe* substitute is at the county scale, $d^* =$ county; the operative cutoff y^*_{county} is high, so only the highest-SES families exit, and re-sorting registers as between-county segregation.

5 Empirical Approach

I employ a staggered Difference-in-Differences (DiD) approach that exploits exogenous variation in the timing of shootings. This approach presupposes that school shooting incidents are quasi-random events to estimate their causal effects on mobility as captured by enrollment and sorting at different geographic levels.

My identification strategy relies on the key assumption that conditional on observed covariates, enrollment and other mobility outcomes would have had the same trends in schools that experience a shooting as in schools that have not yet experienced a shooting, had there been no shooting.

On average, shooting schools are larger, have a higher share of poor and Hispanic students, and sit in counties with more homicides. These differences in levels are absorbed by unit fixed effects and the covariates entering the estimator. This is not a direct threat to identification since identification rests on parallel trends, not on similarity in levels.

My primary outcome of interest is annual log enrollment at the school-level. This is because if we do not find any mobility responses to a shooting, then any potential link between school shootings and residential sorting by SES cannot be explained by mobility. Simply put, if households do not react to a shooting by switching schools or relocating, then there can be no re-sorting of students caused by the shooting. The secondary outcomes are the share of FRL students at the school-level and the Normalized Exposure Index at the between-school, between-district and between-county levels.

If the treatment changes over time, i.e. the characteristics of schools that experience a shooting over time changes, then we would be worried that the not-yet-treated schools are an imperfect control group. For that, we would need an alternative control group.

The main challenge is to find a suitable control group for the shooting schools. One potential solution is schools in the same neighborhood serving as control schools since they share the same neighborhood characteristics and other potential confounders. However, it is very likely that these other schools in the same neighborhood are directly or indirectly affected by the shooting. A neighboring school might be directly affected, by experiencing the same threat and engaging in lock-downs and other measures to protect their students. This would be a direct effect where neighboring schools are essentially themselves treated schools. Alternatively, neighboring schools could be a refuge for students that leave shooting schools. The neighboring school could temporarily be taking in students from the affected school. Therefore, although the neighboring school is not itself treated, there could be spillover effects making it an unsuitable control school. I perform a spillover analysis to answer this (Section 7.1).

Identifying variation

Two claims must be kept separate. Where a shooting happens falls short of random, but observables such as urban location predict it, so unit fixed effects and covariates absorb the location component. When a shooting happens, conditional on location, no one predicts, and that timing

creates the identifying variation. Contagion and imitation do not threaten the exogeneity of shooting timing assumption. Both contagion and imitation operate at the state or national level, not the individual school, and the annual panel absorbs the daily and weekly contagion documented for mass shootings (Fox et al. 2021; Kissner 2016). Counties that experience mass shootings look statistically indistinguishable from those that do not across political, demographic, economic, and crime characteristics (Yousaf 2021), and observable school characteristics do not predict shootings (Rossin-Slater et al. 2020). Section 8 returns to timing with a randomization-inference test and with evidence on how often shootings repeat at the same school.

5.1 Staggered Difference-in-Differences Exploiting Exogenous Timing

The traditional two-way fixed effects (TWFE) estimator can be biased in settings with staggered treatment and heterogeneous treatment effects across cohorts or time. School shootings are such a setting. In the case of heterogenous treatment effects, the TWFE would give us biased estimates as it uses already-treated schools as controls (leading to the so-called forbidden comparisons, as discussed in De Chaisemartin and d'Haultfoeuille 2020).

I employ the staggered Difference-in-Differences estimator of Callaway and Sant'Anna (2021), which addresses these issues by computing group-time average treatment effects. Units are grouped by their treatment timing (cohort g), and effects are estimated separately for each cohort at each post-treatment period t. The covariates used in the regression adjustment are the school's level (primary, middle, high, or other) and its locale (city, suburb, town or rural). The literature has identified these two characteristics as most likely to be associated with the incidence of a school shooting, and they may also be associated with differential enrollment trends.

I aggregate the group-time ATTs into event-study parameters. This aggregation produces event-study estimates showing how effects evolve with time since treatment, averaging across cohorts that experienced shootings in different years. I report the dynamic Average Treatment Effect on the Treated (ATT) in an event-study plot. For analyses that are run on separate samples, I also report the simple overall ATT by sample in a plot. The coefficients in the event-study plots for log-transformed outcomes are interpreted as average percentage changes in the outcome a certain number of periods after exposure to a shooting. For percentage outcomes, the coefficients are interpreted as percentage point changes.

Estimations are done at the school, district and county levels. In all estimations, the control group is the not-yet-treated and never-treated units, and the estimation method is outcome regression. The doubly robust and inverse-probability-weighting estimation methods would both require estimating a propensity score, and a school shooting is too rare an event to model as a first stage. A fitted propensity score would be noise, and weighting by its inverse would affect every estimate. I allow for one anticipation period ($a = 1$), so units treated in the first available year of their panel are dropped for lack of a clean pre-period. At the school level, where data begin in 1990, this drops units treated at or before 1991. Group-time effects use a universal base period such that each $ATT(g, t)$ is measured against the same within-cohort reference period $g - 1 - a$, which with $a = 1$ is $g - 2$. This differs from a varying base period, under which pre-treatment effects are instead measured against the immediately preceding period. The event-study coefficient at $e = -2$ is therefore zero by construction, while $e = -1$ is freely estimated. I calculate multiplier-bootstrap

standard errors with 1,000 iterations and report pointwise confidence bands. Standard errors are clustered at the unit of analysis in each regression: school, district or county.

6 Main Results

6.1 School Enrollment Falls

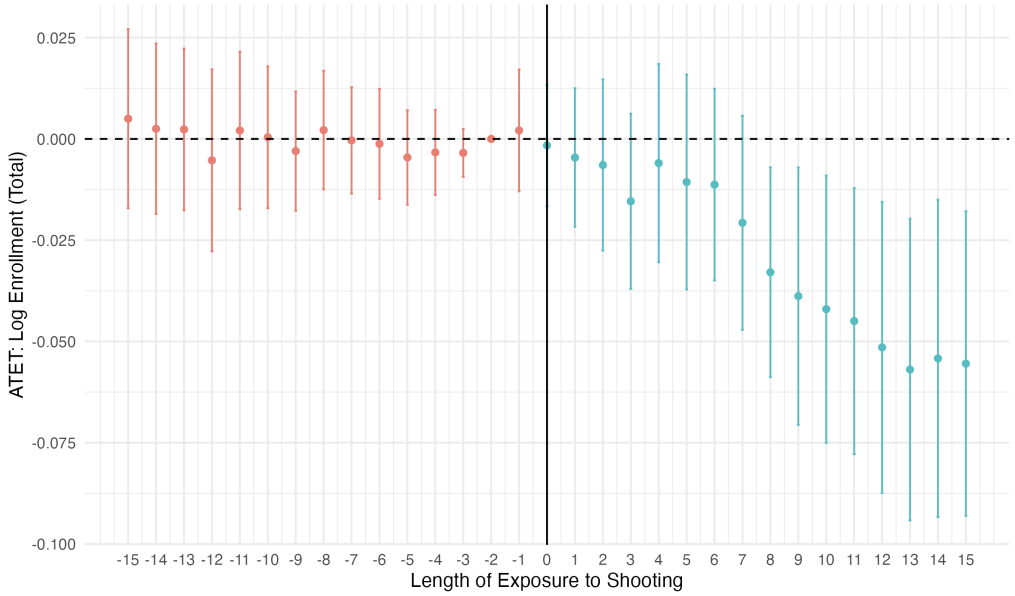
The overall (simple) ATT on log enrollment is -0.042, with the dynamic aggregation by length of exposure to the shooting shown in Figure 6. Given that the average shooting school had approximately 773 students (log enrollment of 6.65) pre-shooting, this decrease of about 4% corresponds to 32 students.

Putting these estimates in perspective, enrollment following the biggest mass shootings in the U.S. declined by 5%, 12%, 1% and 7% in the year following a shooting for the Columbine, Sandy Hook, Parkland, and Santa Fe shootings. Considering that we are looking at a collection of mass shootings as well as non-fatal gun brandishing, an average effect of around -4% is within a reasonable range of estimates.

These results are in line with the enrollment declines found in Yang and Gopalan (2023) (5%, with schools in non-shooting districts as comparison). Similar to them, I find that the long-run effects are more negative than the average effect. Studying high-schools rather than K-12, Abouk and Adams (2013) found a 0.4%–1.3% decline in public high schools state-wide while Beland and Kim (2016) found 5.8% decline in earliest grade enrollment in high schools at the school-level.

It is worth noting the discrepancy between the simple ATT and the overall dynamic ATT. First, the enrollment decline grows with event time. However, since school shootings have increased over time, most treated schools became treated in recent years. If a school experienced a shooting in 2018 and is observed through 2022, it contributes effects out to event-time $e = 4$. Since the simple ATT weighs each cohort by its size *and* its number of post-periods, these recent cohorts count for little and the average is heavily influenced by early cohorots followed over long horizons. The overall dynamic ATT instead weights each horizon equally, so that small short-run effects count just as much as the large long-run ones, pulling the average toward zero. This is why the simple ATT (-0.042) exceeds the dynamic ATT (-0.034) in magnitude.

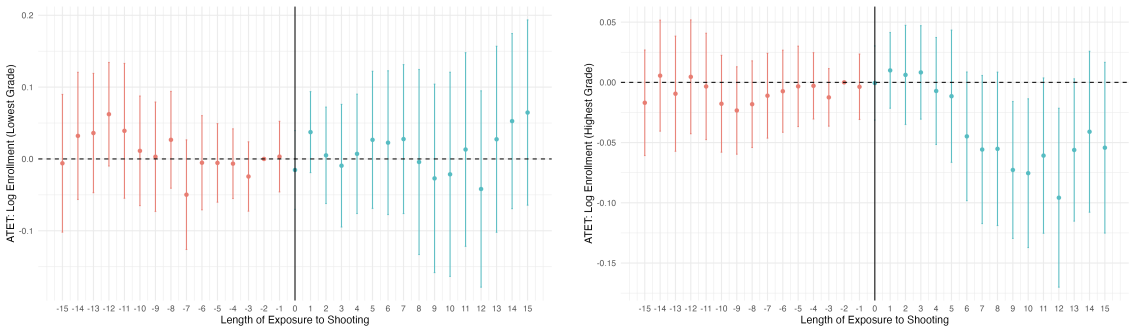
Figure 6. Log Enrollment



Notes: Dynamic ATT on log total enrollment by event time, relative to the year before the first shooting. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The estimate uses 1,307 treated and 51,568 control schools (1,744,875 school-year observations) on a balanced panel. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. A joint Wald test of the pre-treatment coefficients ($e \in [-15, -1]$) does not reject flat pre-trends ($p = 0.917$).

Lowest vs. Highest Grade Enrollment. Another dimension of changes in total enrollment is whether they are driven by fewer incoming students or by fewer graduating students. Naturally, both forces may be at play here. Splitting school enrollment by the lowest versus highest grade offered in Figure 7 shows an interesting pattern. The estimates are noisy, but there is a slight decrease in highest-grade enrollment six years after the initial exposure and no change in lowest-grade enrollment.

Figure 7. Log Enrollment by Grade Level



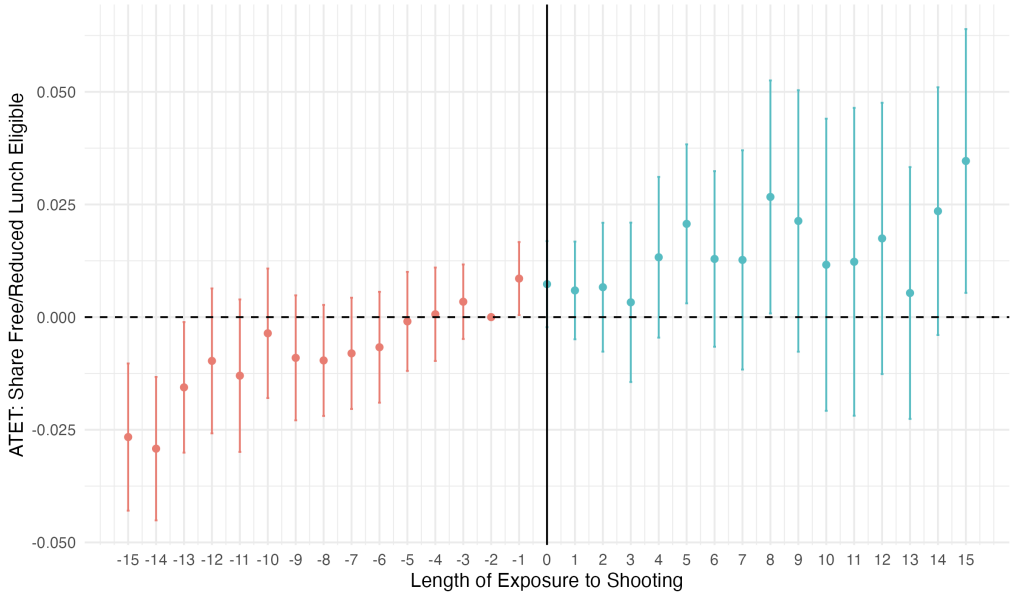
(a) Lowest Grade (b) Highest Grade

Notes: Dynamic ATT on log enrollment in the lowest (left) and highest (right) grade the school offers, by event time relative to the year before the first shooting. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The lowest-grade estimate uses 872 treated and 25,896 control schools (883,344 school-year observations); the highest-grade estimate uses 1,289 treated and 50,210 control schools (1,699,467 observations). Both panels are balanced, which is why the lowest-grade sample is roughly half the size: a school that does not report enrollment by grade in every year drops out entirely. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Joint Wald tests of the pre-treatment coefficients ($e \in [-15, -1]$) do not reject flat pre-trends in either panel ($p = 0.687$ lowest grade, $p = 0.884$ highest grade).

6.2 The Leavers Are Non-Poor

How does the school's student composition change in terms of socioeconomic backgrounds? Figure 8 shows a positive effect on the share of Free or Reduced Lunch Eligible students following a shooting, though imprecisely estimated. From 2014 the Community Eligibility Provision lets high-poverty schools report every student as eligible, which threatens this outcome directly. Section 8 shows that dropping the schools that adopt it does not shrink the estimate.

Figure 8. Free Lunch Eligible



Notes: Dynamic ATT on the share of students eligible for free or reduced-price lunch, by event time relative to the year before the first shooting. Coefficients are percentage-point changes. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The estimate uses 600 treated and 32,576 control schools (829,400 school-year observations) on a balanced panel; the sample is smaller than for enrollment because lunch eligibility is reported later and less completely. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. A joint Wald test of the pre-treatment coefficients ($e \in [-15, -1]$) *rejects* flat pre-trends ($p = 0.001$). This outcome therefore fails the pre-trend test that enrollment passes, and the post-period path should be read as suggestive rather than as a clean causal estimate.

In summary, Table 3 collects the four school-level estimates that were discussed thus far regarding total enrollment, enrollment in the lowest and the highest grade a school offers, and the share of students eligible for free or reduced-price lunch. Enrollment falls, the fall is concentrated in the highest grade, and the share of poor students rises.

6.3 Socioeconomic Sorting across Geographic Scales

To investigate whether the non-random enrollment translates into broader socioeconomic sorting, I apply the same staggered Difference-in-Differences framework to segregation outcomes. The analysis is conducted at three geographic scales: between schools within a district, between geographic school districts within a county, and between counties within a state. The main outcome at

Table 3. School-Level Effects of a Shooting on Enrollment and Composition

Outcome	Treated × Post	SE	N	Treated	Control	Treated Mean	Control Mean
Log Enrollment (Total)	-0.042***	(0.012)	1,744,875	1,307	51,568	6.650	6.106
Log Enrollment (Lowest Grade)	0.005	(0.039)	883,344	872	25,896	5.448	4.733
Log Enrollment (Highest Grade)	-0.050**	(0.020)	1,699,467	1,289	50,210	4.994	4.423
Share Free/Reduced Lunch Eligible	0.014*	(0.008)	829,400	600	32,576	0.534	0.501

Notes: Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Each row is a separate estimate of the overall (simple) ATT. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state × locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The panel is balanced within each estimate, so the treated and control counts are the schools actually estimated on. N counts school-year observations. Treated and control means are pre-treatment. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. The joint Wald test of the pre-treatment event-time coefficients ($e \in [-15, -1]$) does not reject flat pre-trends for total enrollment ($p = 0.917$), lowest-grade enrollment ($p = 0.687$), or highest-grade enrollment ($p = 0.884$). For the free/reduced-price lunch share it *does* reject ($p = 0.001$), so that row should be read as descriptive rather than causal. $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

each scale is the Normalized Exposure Index for Free Lunch to Non-Free Lunch recipients, which adjusts for changes in the overall share of Free Lunch recipients and thereby isolates distributional sorting from underlying increase in the share of Free Lunch recipients.

Do families sort by socioeconomic status following a shooting, and at which geographic scale? Figure 9 plots the overall ATT of the Normalized Exposure Index across all three geographic scales. The results are unambiguous: there is no detectable effect on between-school or between-district segregation once compositional change is accounted for. However, between-county (within-state) segregation increases significantly, with an overall ATT of approximately 0.022.

To gauge the magnitude of this effect, Table 4 reports the pre-period level of the Normalized Exposure Index at each scale. Among treated states, the between-county index averages 0.057 in the pre-period, so the estimated increase of 0.022 represents a rise of roughly 40% over baseline. The effect is also large relative to the cross-sectional spread: the pre-period interquartile range of between-county segregation runs from 0.024 to 0.093, so a 0.022 shift moves the average treated state by about a third of that range. By contrast, the between-school baseline is far higher (a treated mean of 0.110), yet the estimated effect there is a precise null.

Figure 10 shows the event-study version. The between-county effect is not immediate but emerges gradually from around three years post-shooting and continues to grow over the observed horizon, consistent with a slow-moving residential relocation process. Pre-period estimates are flat across all scales, supporting the parallel trends assumption. The between-school and between-district estimates are precise nulls throughout, ruling out within-district school-switching as the primary sorting mechanism. Taken together, the pattern is consistent with residential relocation across county lines: non-poor families leave the county of the affected school, increasing between-county socioeconomic segregation while leaving within-district sorting unchanged.

We can compare this with Sezer (2022), who finds that directly exposed survivors are less likely to leave the district than non-exposed individuals, yet are *less* likely to reach a top-income

district or county by age 30. These results are compatible with, and plausibly complementary to, mine, because the two papers concern different populations. Sezer (2022) identifies the directly exposed cohort, whereas my county-level effect reflects all families (most of whom were never exposed). Unchanged individual exit rates can therefore coexist with aggregate between-county sorting, provided those who leave differ in composition from those who stay. A potential solution to this paradox is the following. Suppose the directly exposed students are disproportionately lower-SES, a population that is less likely to move, which is consistent with my result. Sezer's results would then describe a population that, when it does relocate, tends to move *laterally or downward* rather than into more affluent areas, leaving its members in lower-income districts. The affluent, not-directly-exposed families face no such constraint: better able to absorb the costs of moving, they are the ones most likely to exit the county for higher-income areas following a shooting. The two flows act in the same direction; lower-SES families remaining or sorting down, higher-SES families sorting up and out. Each if these effects push toward greater socioeconomic divergence *between* counties, which is precisely the scale at which I detect the effect. I do not observe these moves directly and therefore infer the mechanism from the aggregate compositional pattern and from the nested scale at which segregation rises.

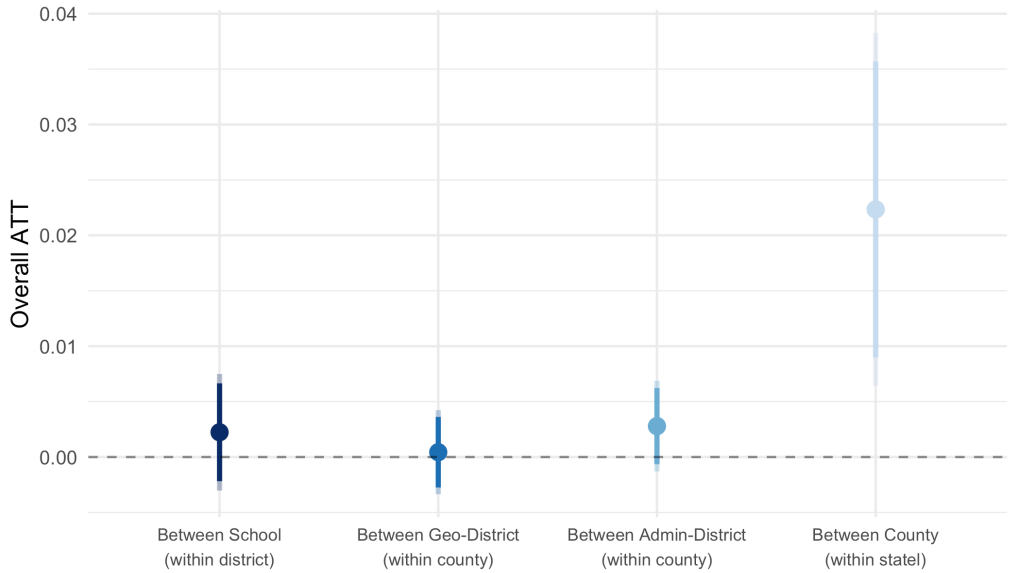
Table 4. Pre-Period Baseline of the Normalized Exposure Index by Geographic Scale

Scale	Treated mean	Treated SD	Treated IQR	Treated N	Control mean	Control SD	Control IQR	Control N
Between schools (within district)	0.110	0.099	[0.028, 0.170]	326	0.019	0.042	[0.000, 0.017]	12,281
Between geo-districts (within county)	0.067	0.084	[0.000, 0.105]	574	0.030	0.043	[0.000, 0.041]	2,563
Between counties (within state)	0.057	0.045	[0.024, 0.093]	37				

Notes: Pre-period level of the Normalized Exposure Index (Free Lunch vs. Not-Free-Lunch recipients) for treated versus control units at each geographic scale. Each unit contributes one number: for treated units the mean over its years before the first shooting, for never-treated controls the mean over all its years. Mean, SD and IQR are taken across units. IQR is reported as the [25th, 75th] percentile of those unit-level means. No never-treated control states exist at the between-county scale, so that row reports treated states only.

Table 5 reports the underlying overall ATTs. Each panel is a separate estimation panel with its own unit of analysis, so the three panels are not nested estimates of one regression but three regressions run at three scales. Reading down the panels reproduces the pattern in Figure 9: the normalized index is flat within district and between districts, and moves only between counties.

Figure 9. Normalized Exposure Index: Overall ATT by Geographic Scale



Notes: Overall (simple) ATT for the Normalized Exposure Index (Free Lunch vs. Non-Free Lunch recipients) at four geographic scales: between schools within district, between geographic districts within county, between administrative districts within county, and between counties within state. The index adjusts for the share of Free Lunch recipients, so it isolates distributional sorting from compositional change. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Each scale is a separate estimation panel: no covariates enter the regression adjustment at any scale, and standard errors are clustered on the panel's unit of analysis (district, county, or state). Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Sample sizes are 315 treated and 9,035 control districts between schools; 559 treated and 2,563 control counties at both between-district scales; and 32 treated states with *no never-treated control states* between counties, where identification is off not-yet-treated states only. The control pool is restricted to units sharing a (state $\times$ locale) cell with a treated unit in the district panel, but no such restriction is applied in the county and state panels. The joint Wald pre-trend test reported for the school-level figures is not computed at these scales; see Table 5 and the discussion of the between-county comparison group there.

Table 5. Composition and Segregation Effects by Geographic Scale

Panel A. District panel: composition and segregation within district

	Treated × Post	SE	N	Treated	Control	Treated Mean	Control Mean
Share Free Lunch	-0.006	(0.003)	103,439	315	9129	0.500	0.397
Share Reduced Lunch	0.002**	(0.001)	103,439	315	9129	0.076	0.080
Share Free/Reduced Lunch	-0.003	(0.003)	103,439	315	9129	0.576	0.477
SES Composite	0.003	(0.013)	102,378	315	8999	-0.139	0.248
Poverty Rate	-0.001	(0.001)	102,378	315	8999	0.160	0.127
Segregation: Information Theory (FL/NFL)	-0.002	(0.003)	103,362	315	9129	0.089	0.016
Segregation: Relative Diversity (FL/NFL)	-0.002	(0.003)	103,362	315	9129	0.103	0.018
Between-School: Exposure (FL/NFL)	0.010***	(0.003)	102,347	315	9035	0.451	0.592
Between-School: Isolation (FL/NFL)	-0.010***	(0.003)	102,347	315	9035	0.549	0.407
Normalized Exposure Between Schools (FL/NFL)	0.002	(0.003)	102,346	315	9035	0.108	0.021

Panel B. County panel: segregation between districts within county

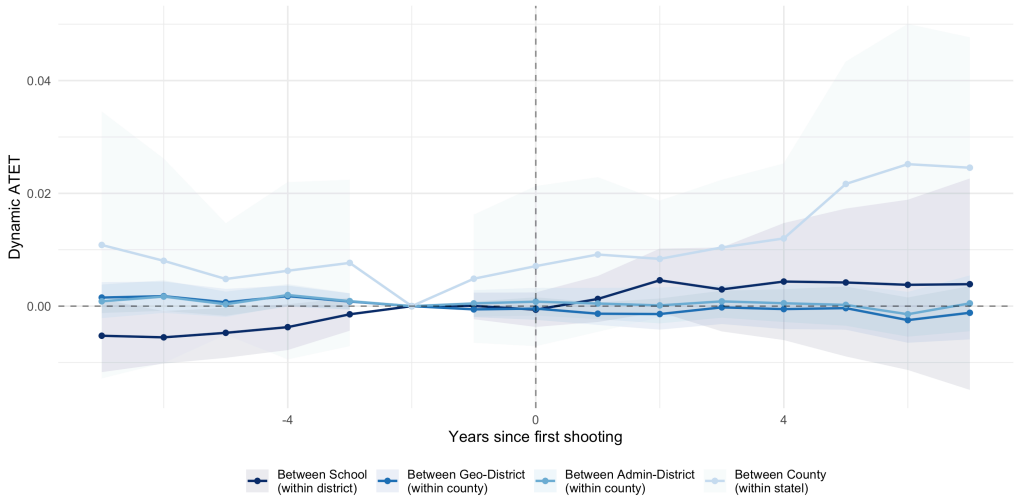
	Treated × Post	SE	N	Treated	Control	Treated Mean	Control Mean
Between Admin-Districts: Exposure (FL/NFL)	-0.006*	(0.003)	99,076	559	2558	0.583	0.586
Between Admin-Districts: Isolation (FL/NFL)	0.006*	(0.003)	99,076	559	2558	0.417	0.414
Between Admin-Districts: Normalized Exposure (FL/NFL)	0.003	(0.002)	99,190	559	2563	0.065	0.031
Between Geo-Districts: Exposure (FL/NFL)	-0.006*	(0.003)	99,047	558	2561	0.585	0.587
Between Geo-Districts: Isolation (FL/NFL)	0.006*	(0.003)	99,047	558	2561	0.415	0.413
Between Geo-Districts: Normalized Exposure (FL/NFL)	0.000	(0.002)	99,195	559	2563	0.063	0.029

Panel C. State panel: segregation between counties within state

	Treated × Post	SE	N	Treated	Control	Treated Mean	Control Mean
Between Counties: Exposure (FL/NFL)	0.014	(0.011)	1,024	32	0	0.696	—
Between Counties: Isolation (FL/NFL)	-0.014	(0.011)	1,024	32	0	0.304	—
Between Counties: Normalized Exposure (FL/NFL)	0.022***	(0.008)	1,024	32	0	-0.045	—

Notes: Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Each row is a separate estimate of the overall (simple) ATT, and each panel is a separate estimation panel: the unit of analysis is the school district in Panel A, the county in Panel B, and the state in Panel C. No covariates enter the regression adjustment at any of these three scales, and standard errors are clustered on the unit of analysis of the panel in question. Event time is restricted to $e \in [-7, 7]$. N counts unit-year observations; treated and control counts are units. Treated and control means are pre-treatment. In Panel A the control pool is restricted to districts sharing a (state × locale) cell with at least one treated district; *no such restriction is applied in Panels B and C*, whose panels are built at the county and state level. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Unlike the school-level estimates, the joint Wald test of pre-treatment event-time coefficients is not computed at these scales, so the flat pre-periods in Figure 10 rest on the individual coefficients rather than on a joint test. Every state is treated at some point and the last treated cohort enters in 2012, so Panel C has *no never-treated comparison units* (control count 0, control mean shown as —) and identifies off not-yet-treated states only; those states are the ones that went longest without a shooting anywhere in the state. Panel C should be read with that limitation in mind. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Figure 10. Normalized Exposure Index: Dynamic ATT by Geographic Scale



Notes: Dynamic ATT for the Normalized Exposure Index (Free Lunch vs. Non-Free Lunch recipients) at four geographic scales, by event time relative to the year before the first shooting. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Each scale is a separate estimation panel: no covariates enter the regression adjustment at any scale, and standard errors are clustered on the panel's unit of analysis (district, county, or state). Event time is restricted to $e \in [-7, 7]$ and the panel is balanced. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Shaded areas are therefore pointwise, so they should not be used to judge the whole path jointly. Sample sizes are 315 treated and 9,035 control districts between schools; 559 treated and 2,563 control counties at both between-district scales; and 32 treated states with *no never-treated control states* between counties, where identification is off not-yet-treated states only. The control pool is restricted to units sharing a (state $\times$ locale) cell with a treated unit in the district panel, but no such restriction is applied in the county and state panels. The joint Wald pre-trend test reported for the school-level figures is not computed at these scales, so the flat pre-periods shown here rest on the individual coefficients rather than on a joint test.

7 Mechanisms and Heterogeneity

The main results establish the headline pattern: enrollment falls, the students who remain are poorer, and socioeconomic sorting emerges between counties. This section asks *why* re-sorting appears at the county scale rather than nearer, and how the response varies with the severity of the shock and the ease of switching schools. I begin with the spatial spillovers that determine where a genuinely safe substitute lies.

7.1 Spatial Spillovers and the Scale of a Safe Substitute

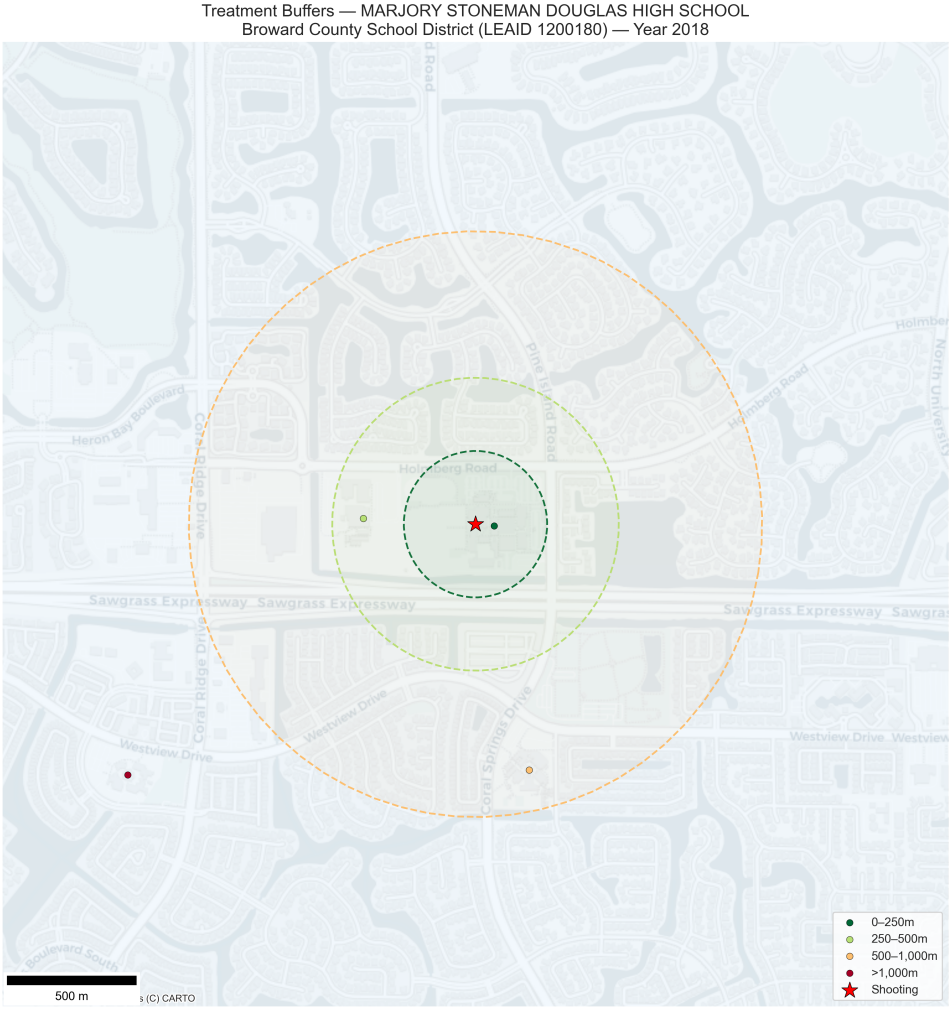
In the preceeding analyses, I used the proximity of schools to a school shooting to assign treatment. Schools within 500 meters of a school shooting are assumed to be treated. The alternative to a spatial join was to merge the school shootings dataset and the school directory based on string identifiers such as school name, city, and county. However, since some shootings occur in close proximity to multiple schools (e.g. on a yard for an elementary school and a high school), assigning these shootings to one school only would be problematic. We would risk identifying a school as a control school when it in fact also has experienced a shooting. This also applies to a school building that has multiple school entities (e.g. a school for the gifted and a regular school sharing the same building). My treatment assignment overcomes these issues by relying solely on schools' and shootings' geographic coordinates.

For example, the gruesome mass shooting at the Marjory Stoneman Douglas High School (serving grades 9–12) in Parkland, Florida in 2018 occurred in proximity to the Westglades Middle School (serving grades 6–8). The perpetrator of this shooting was himself enrolled in this nearby middle school in the years prior to carrying out this attack. In Figure 11, a map with increasing radius of treatment assignment shows that students of this middle school would not have been considered treated had the spatial join only captured school units within 250meters of a school shooting. Within 500 meters, both the focal high school and the middle school across the yard are considered affected. A three minute drive away lays the Country Hills Elementary School, captured by the 1km radius in the map. Testimony from a school police, Dave Dittman, jumping into action to help protect the students at Marjory Stoneman Douglas High School, shows how resources and individuals in the nearby community are affected (Smith 2023).

The two natural follow-up questions are: How sensitive are the results to the exact treatment assignment of 500 meters? And can we learn anything about spillover effects? To answer this, I rerun the main analysis using six radii of increasing size. As seen in Figure 12, for radii between 250m and 5km, the results are robust, hovering around -4%. The two other prominent patterns are that the effects of shootings on log enrollment goes toward zero (and eventually turns positive) as the radius increases, and the estimates become more precise. For the latter pattern, as our treatment radius increases, so too does the number of treated schools and thereby increasing the precision. For the former pattern, these results tell a story where closeby schools are affected negatively, while a radius of 50km may reflect that schools far away from the focal shooting school may become *more* attractive and gain in student enrollment. This can be interpretted as the nearest safe alternative d^*.

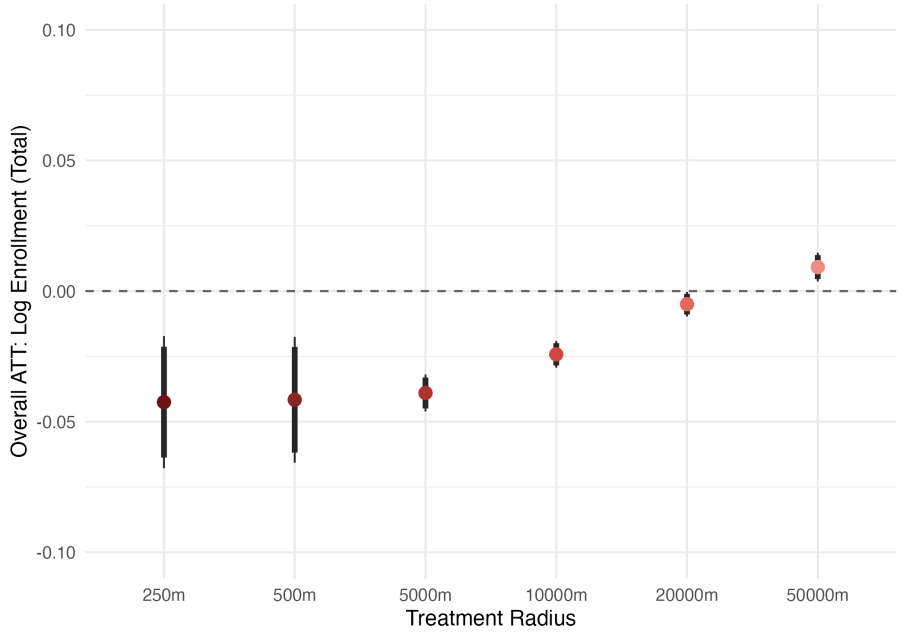
Because every disc contains the focal school, the sweep cannot show where the effect stops. Appendix G re-runs it on disjoint rings that exclude the focal school, against one control pool common to every ring (Figure E.1, panel D). The innermost ring reproduces the headline effect at -0.040, and the 5–10 km ring puts the effect at zero. The 2–5 km ring points to a loss of about 3% at schools that were never near a shooting themselves, though it narrowly fails its pre-trend test. Beyond 10 km the ring estimates turn positive, at 0.017 for 10–20 km and 0.022 for 20–50 km, in line with the sign change in the sweep. Their pre-period coefficients drift only early in the fifteen-year window and are flat over the last five years before the shooting.

Figure 11. Expanding Treatment Radius Example Map

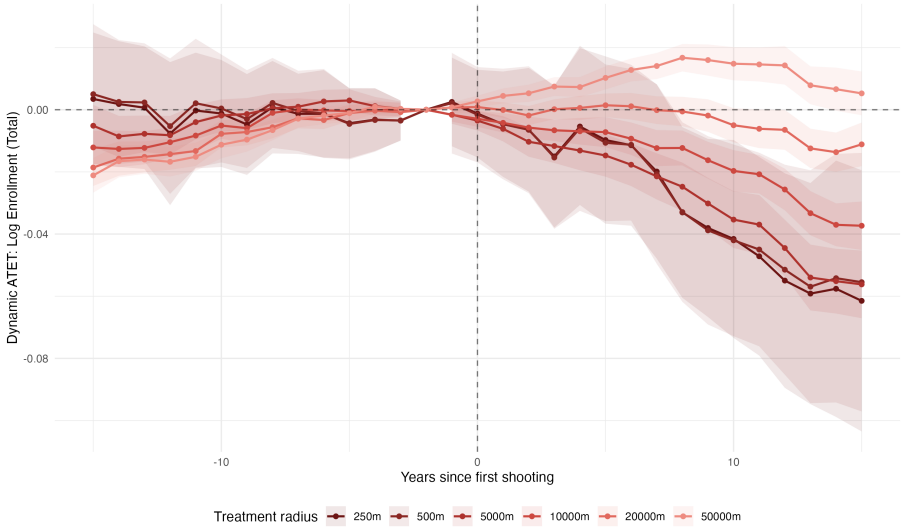


Notes: Map showing Westglades Middle School included in treatment of 500 meters but not in 250 meters. Rings correspond to increasing treatment radius from 250 meters to 1+ kilometers.

Figure 12. Expanding Treatment Radius



(a) Overall ATT



(b) Dynamic ATT

Notes: Overall ATT (top) and dynamic ATT (bottom) on log enrollment as the radius defining treatment widens from 250 metres to 50 km. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The (state $\times$ locale) control-pool restriction is re-applied separately at each radius, so the control pool is rebuilt rather than held fixed as the radius grows. Treated and control schools are, by radius: 1,308 and 51,412 at 250 m; 1,307 and 51,568 at 500 m; 11,978 and 41,797 at 5 km; 19,165 and 33,329 at 10 km; 24,862 and 24,978 at 20 km; and 32,823 and 12,104 at 50 km. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Joint Wald tests of the pre-treatment coefficients ($e \in [-15, -1]$) do not reject flat pre-trends at 250 m ($p = 0.898$) or 500 m ($p = 0.917$), but *reject at every wider radius* ($p = 0.011$ at 5 km and $p < 0.001$ at 10, 20 and 50 km).

7.2 Heterogeneity by Fatality and Underlying Violence Level

Shooting Fatality. The results so far pool all shootings together. A natural question is whether the effects are driven by fatal shootings. Figure 13 overlays the enrollment estimates from the full sample with those restricted to fatal shootings. Focusing on shootings that result in at least one fatality, the decline is more pronounced, particularly at longer horizons where it reaches approximately a 10% decline in enrollment, compared to about 4% decline on average for all shootings. Due to there being too few county-year observations with fatal shootings, I am underpowered in the analysis of the effects of school shootings on segregation when limiting to the fatal shootings only.

Historical county violence. The enrollment response may also vary by the level of pre-existing violence in the county. In a peaceful neighborhood, a school shooting may be seen as an extreme outlier event, not prompting any action on the parents part. In an otherwise violent, one might also expect a school shooting not to be as salient given the high levels of violence, and therefore not prompting any action on the part of the parents. That is unless a school shooting is seen as the event that triggers a response to an otherwise violent environment. Figure 14 splits the sample by quartiles of pre-shooting county violence (as measured by the Supplementary Homicide Reports victim count). For treated schools, I use the average pre-treatment SHR county victim count. Never-treated schools get their mean value of SHR victim count based on their whole period in the data. Once there are is a unique pre-treatment (or overall) mean value of SHR county victim count, I split the sample into five quantiles. There is strong indication that schools in the most violent counties (Q4) see the largest enrollment decline following a shooting, while schools in less violent counties (Q2) show near-zero effects. The dynamic plot shows that Q4 counties drive the overall negative effect, with enrollment declining steadily in the years following a shooting.

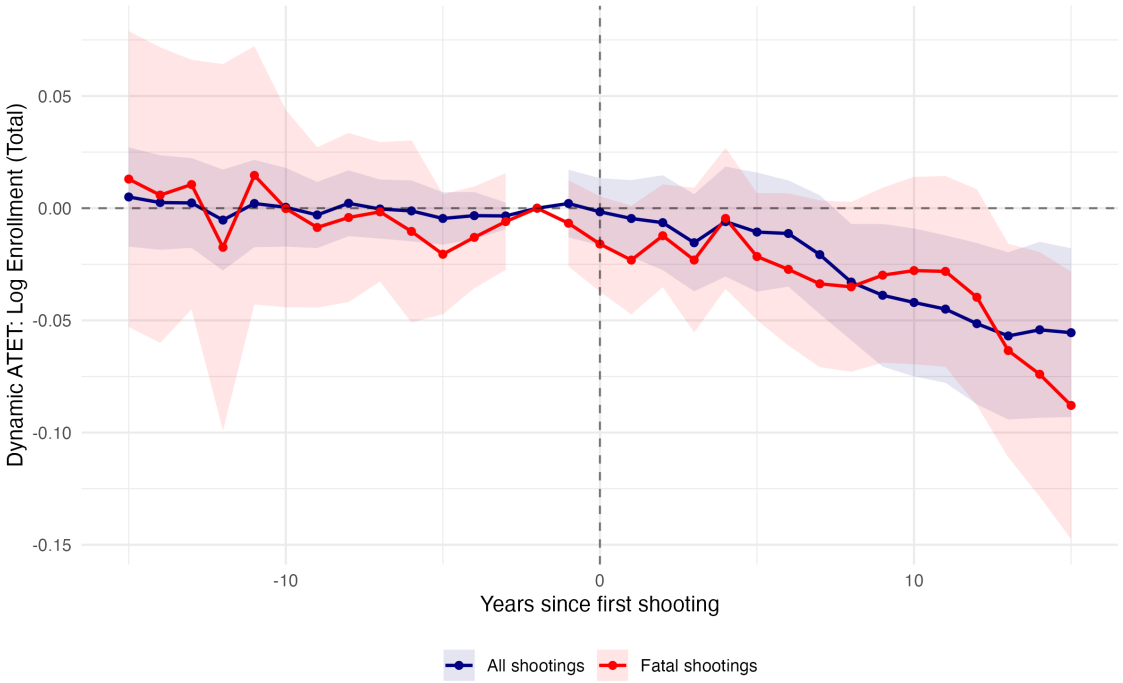
7.3 Heterogeneity by Ease of Switching Schools

Open Enrollment laws. There may be institutional reasons for why we see a strong mobility response in some schools and not in others. In the U.S., the ease of switching schools is a state-matter. Figure 15 shows which states have the most stringent open enrollment laws. We can broadly categorize states into those that allow students to switch school within a school-district, and those that allow students to switch between school-districts. Note that the exact regulations are set by school-districts. For example, while a state might permit intra-district change that does not mean that the district must allow intra-district changes, only that it may.

Figure 16 estimates the dynamic average treatment effects on log enrollment, splitting the sample between states that permit within-district school switching and those that don't permit that. In states where intradistrict open enrollment is permitted, we see a smaller decline in enrollment from schools following a shooting, *counter* to the idea that ease of switching facilitates mobility responses. Figure 17 splits the sample by states that allow between-district school switching and those that don't. States where school switching between districts is not permitted have more negative effects on enrollment. However, inspecting the number of shootings per state over time (Figure C.1) shows that—with the exception of Florida—9 of the 10 most affected states have strict interdistrict open enrollment rules. Figure 18 puts these overall ATT in relation to the full sample ATT.

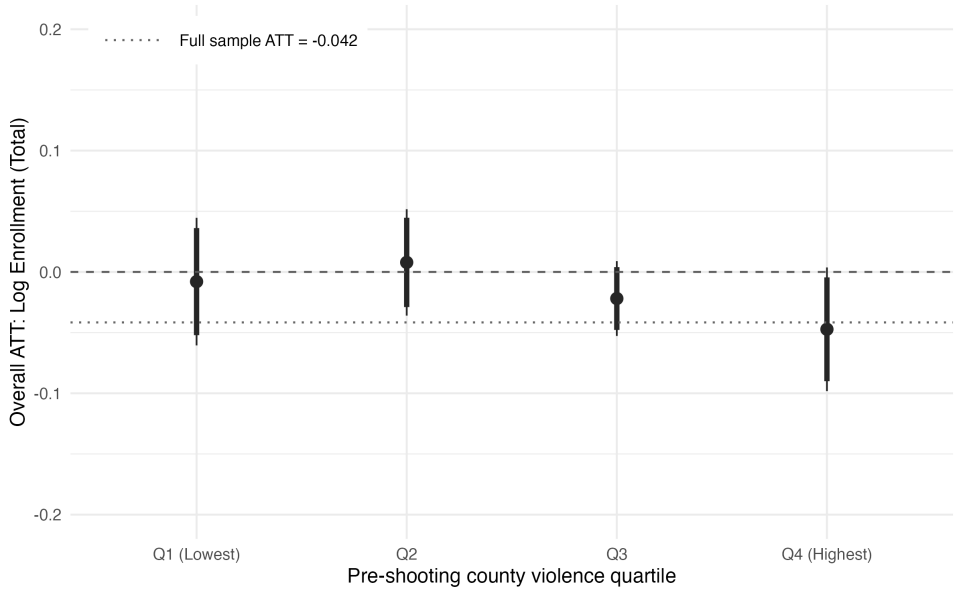
Number of Schools in District. A central thesis is that we cannot expect students to switch schools if it is difficult to do so. One dimension of ease of switching schools is the number of

Figure 13. Log Enrollment: All Shootings vs Fatal

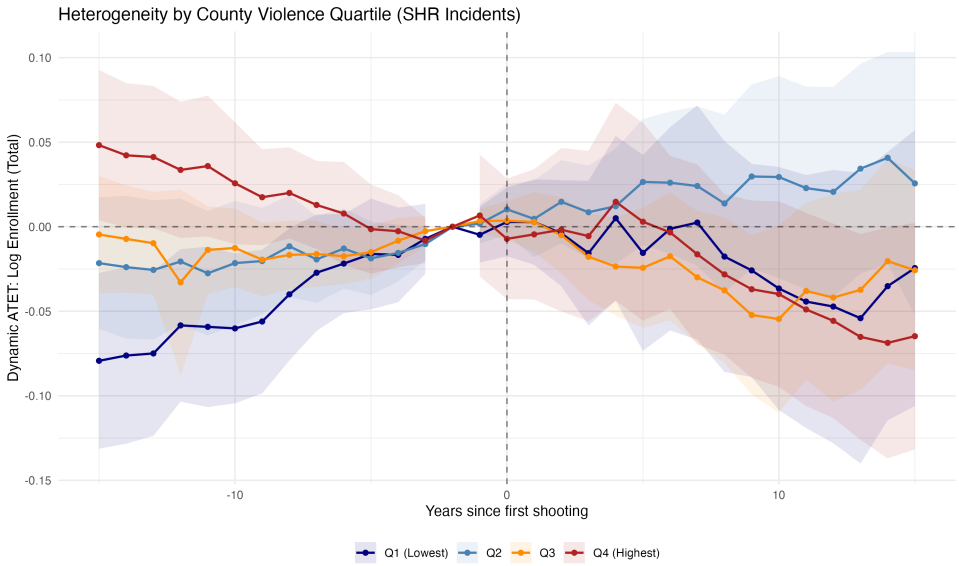


Notes: Dynamic ATT on log enrollment, overlaying all shootings with fatal shootings only, by event time relative to the year before the first shooting. The fatal series keeps only fatal-shooting schools as treated and retains the full control pool. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Standard errors are clustered on the school and The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The all-shootings series uses 1,307 treated and 51,568 control schools; the fatal series uses 336 treated and the same 51,568 controls. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Joint Wald tests of the pre-treatment coefficients ($e \in [-15, -1]$) do not reject flat pre-trends in either series ($p = 0.917$ all shootings, $p = 0.799$ fatal). Note that the two series differ in more than the treatment definition: the all-shootings series conditions on school level and locale, whereas the fatal series and every other heterogeneity split in this section use no covariates, because the level-and-locale specification is not estimable on the thinner sub-samples. The strictly comparable baseline for the fatal series is therefore a no-covariate re-estimate of the all-shootings specification, not the main-table estimate plotted here.

Figure 14. Effects on Log Enrollment by Historical County Violence



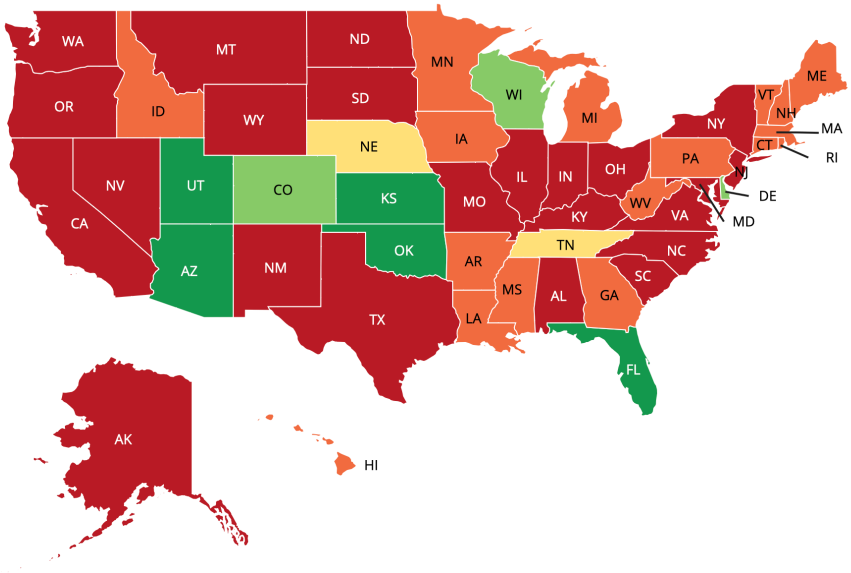
(a) Overall ATT



(b) Dynamic ATT

Notes: Overall ATT (top) and dynamic ATT (bottom) on log enrollment, split by quartile of pre-shooting county violence (Supplementary Homicide Reports victim count), by event time relative to the year before the first shooting. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. No covariates enter the regression adjustment in this split; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. Treated and control schools by quartile are 130 and 14,607 (Q1, lowest violence), 242 and 13,144 (Q2), 390 and 12,537 (Q3), and 537 and 10,836 (Q4, highest). Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Joint Wald tests of the pre-treatment coefficients ($e \in [-15, -1]$) do not reject flat pre-trends in any quartile ($p = 0.218, 0.522, 0.795$ and 0.804 for Q1 to Q4).

Figure 15. Open Enrollment Best Practices by state



Notes: 5-point system for Mandatory Cross-District OE, Mandatory Within-District OE, Transparent Reporting by State Education Agency, Transparent and Public Posting of Vacancies, and (Tuition) Free Access to All Public Schools. Dark green indicates 5/5, dark red indicates 0/5. Source: https://reason.org/policy-brief/public-schools-without-boundaries-a-50-state-ranking-of-k-12-open-enrollment/#california-diversions

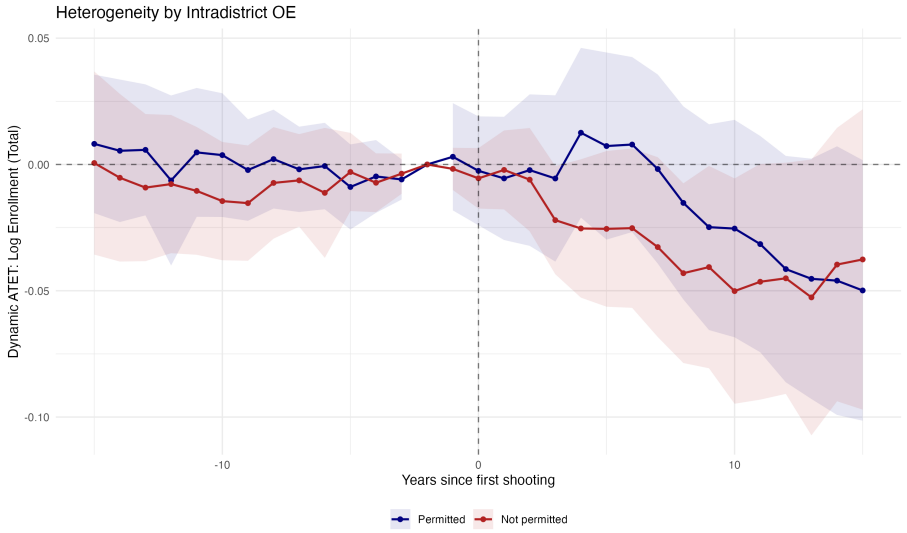
outside options, i.e. other schools of the same school-level within the same district. For example, if a shooting happens in an elementary school and it is the only elementary school in the district, mechanically, we cannot expect any student to switch schools *within* the district. As the number of schools of the same school-level increases, we would expect a stronger negative effect on log enrollment.

To test this, I bin the number of schools and run the same staggered difference-in-differences estimation separately per bin and school-level. The distribution of number of schools per district is heavily right-skewed, with very few big districts and many small ones. The biggest geographic school districts are New York City Public Schools (1,600 unique schools), Los Angeles Unified (1,100 unique schools) and the City of Chicago School District (1,000 unique schools). To put this in perspective, the median district has 3 schools and the mean district has 6 schools. To account for the skewness of the distribution of count of schools per district, I calculate quintiles on the $\log_2$ scale. This gives the bins roughly equal numbers of school-observations in the $\log_2$ space. For example, a district with 1–4 schools has log_2 between 1 and 2, while a district with 51 schools has a $\log_2(51) = 5.7$.

Figure 19 shows the overall ATT of shootings on log enrollment run separately by bin and school-level. [3] Two results stand out. First, the effects are always more negative for primary schools (grades K–3) than middle or high schools. Secondly, in above-average sized districts, some effects on high schools are positive and statistically significant. This is counter to what we

3. The CCD raw data does not include the variable school level before 2000, but school level is constructed based on information about the the lowest grade offered, available from 1986. A school level is determined solely by the lowest grade offered.

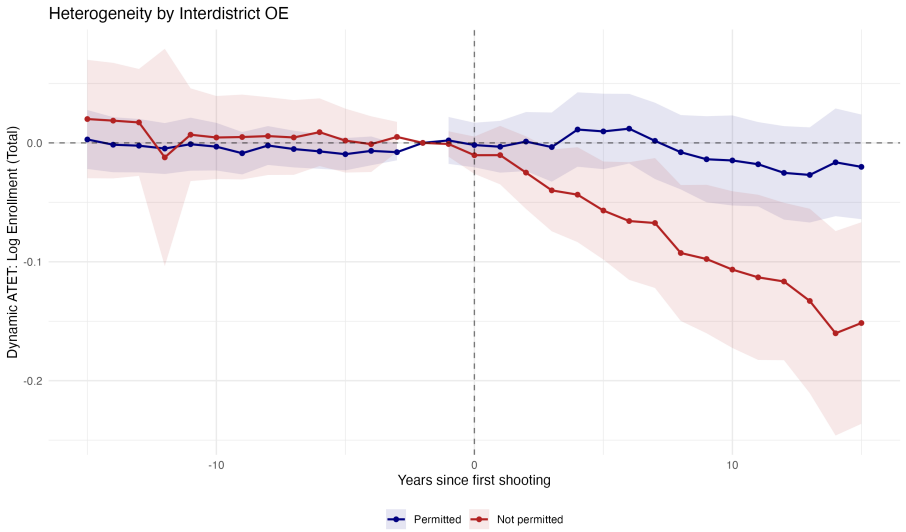
Figure 16. Effects on Log Enrollment by School Choice *Within* District



Notes: Dynamic ATT on log enrollment by event time relative to the year before the first shooting, splitting schools by whether their state permits *intradistrict* open enrollment. 21 of 50 states permit it. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Estimation is at the school level, not the district level. No covariates enter the regression adjustment in this split; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The permitted series uses 907 treated and 34,751 control schools; the not-permitted series uses 400 treated and 16,817 control schools. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Joint Wald tests of the pre-treatment coefficients ($e \in [-15, -1]$) do not reject flat pre-trends in either series ($p = 0.826$ permitted, $p = 0.736$ not permitted).

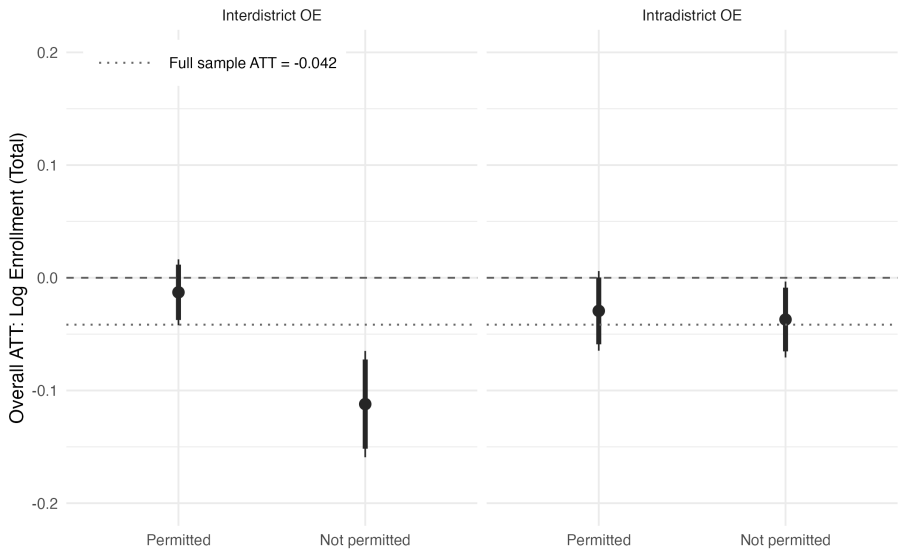
would expect from the theoretical model. As the number of same-grade alternatives increase, one would expect more negative effects on enrollment. One potential explanation is that districts with many schools are in urban areas or big cities that households choose not to leave for other reasons. This garners further investigation.

Figure 17. Effects on Log Enrollment by School Choice *Between* District



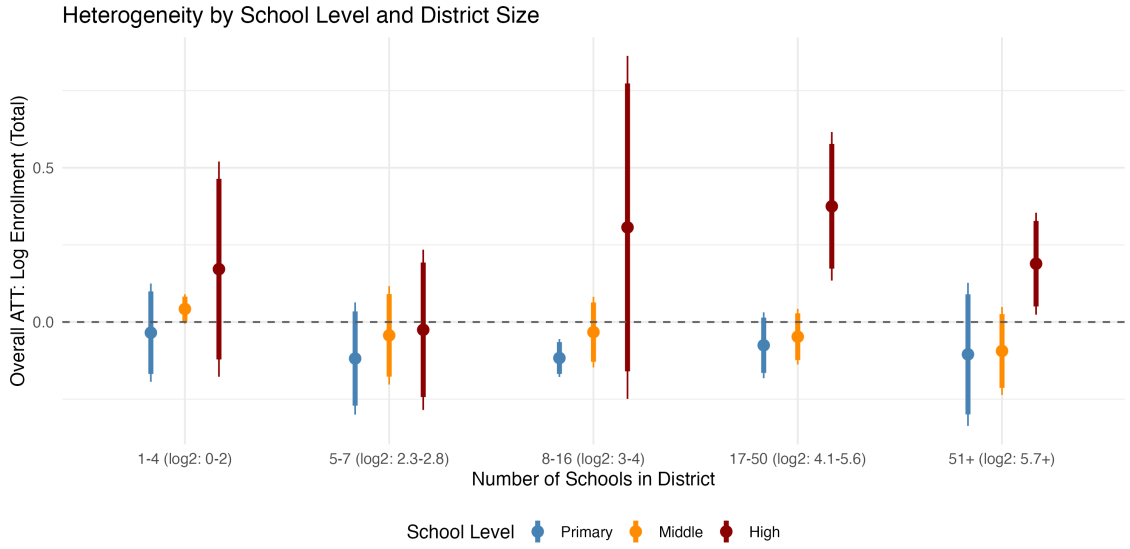
Notes: Dynamic ATT on log enrollment by event time relative to the year before the first shooting, splitting schools by whether their state permits *interdistrict* open enrollment. 43 of 50 states permit it. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Estimation is at the school level, not the district level. No covariates enter the regression adjustment in this split; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The permitted series uses 1,045 treated and 44,135 control schools; the not-permitted series uses 261 treated and 7,433 control schools, so the not-permitted arm rests on seven states and is thinly identified. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Joint Wald tests of the pre-treatment coefficients ($e \in [-15, -1]$) do not reject flat pre-trends in either series ($p = 0.524$ permitted, $p = 0.985$ not permitted).

Figure 18. Effects on Log Enrollment by School Choice Regimes



Notes: Overall (simple) ATT on log enrollment for the four open-enrollment regimes: intradistrict permitted and not permitted, interdistrict permitted and not permitted. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Estimation is at the school level. No covariates enter the regression adjustment in these splits; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. Treated and control schools are 907 and 34,751 (intradistrict permitted), 400 and 16,817 (intradistrict not permitted), 1,045 and 44,135 (interdistrict permitted), and 261 and 7,433 (interdistrict not permitted). Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Joint Wald tests of the pre-treatment coefficients ($e \in [-15, -1]$) do not reject flat pre-trends in any of the four regimes ($p = 0.826$, 0.736, 0.524 and 0.985 in the order listed above).

Figure 19. Effects on Log Enrollment by Number of Same-Level Schools in District



Notes: Overall (simple) ATT on log enrollment by the number of same-level schools in the district, estimated separately for primary, middle and high schools across five district-size bins (15 estimates). The x-axis gives the range of the number of schools in each bin, followed by its $\log_2$ equivalent in parentheses. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. No covariates enter the regression adjustment in this split, and standard errors are clustered on the school. Unlike the other estimates in this paper, these are run on unbalanced repeated cross-sections rather than a balanced panel, because balancing empties the thinner level-by-size cells. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. Across the 15 cells the treated count ranges from 75 to 642 schools and the control count from 6,126 to 22,245. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Joint Wald tests of the pre-treatment coefficients ($e \in [-15, -1]$) do not reject flat pre-trends in 14 of the 15 cells (p from 0.133 to 0.995); the exception is the largest district-size bin for middle schools, which rejects ($p < 0.001$) and should be discounted.

8 Robustness

In this section, I address five potential sources of bias: contamination of the control pool by schools near a shooting, selective attrition through school closures, a break in how free-lunch eligibility is reported, inference that rests on the timing assumption, and shootings that repeat at the same school. The figures for each check are in Appendix H.

Donut Difference-in-Differences. The radius analysis raises a concern about the control pool: schools just outside the 500m treatment ring, but still close to a shooting, may themselves be indirectly affected—through media exposure, district-wide responses, or by absorbing students who leave the treated school. Such contamination could bias the estimate either way. Controls that share the shock lose students too, which would understate the effect, while controls that absorb leavers gain students, which would overstate it. To address this, I re-estimate the main log-enrollment specification on a "donut" sample: I keep the standard 500m treatment definition but restrict the control pool to schools that are never within 50km of any shooting, leaving a cleaner counterfactual of plausibly unaffected schools (Figure E.1, panel C). This retains all 1,307 treated schools but reduces the control pool to the 11,033 of 51,568 schools that are far from any shooting. The overall ATT is -0.033 (SE 0.013), against -0.042 (SE 0.012) in the main specification (Figure H.1). The decline is somewhat smaller, but the confidence intervals overlap and pre-trends remain flat ($p = 0.942$). The result is therefore not driven by contamination of nearby controls.

School closures. Do schools survive a shooting? To understand this, we need to know what happens to a school's operational status in the years following a shooting. This is important to know for evaluating what direction bias we may have due to survivorship. In a standard Difference-in-Differences setting, we get a meaningful result in the intensive margin if our estimation sample is balanced. If we have non-random attrition, such that a shooting school is more likely to close and drop out of the dataset, and the remaining schools are positively selected (e.g. due to their good handling of the aftermath of a shooting), then a balanced estimation sample would lead us to underestimate the true negative effects on log enrollment. Figure H.2 plots raw data showing that the share of ever-treated schools that close increases drastically after the shooting year. The average closure rate is close to zero pre-shooting, and jumps to around 1.5% five years post-shooting. That's equivalent to 1 in every 66 school that has any kind of bullets fired with intent closing down.

Community Eligibility Provision. The free-lunch share has a measurement problem of its own. From the 2014–15 school year, the Community Eligibility Provision lets a high-poverty school serve free meals to all its students without individual applications, so a school that adopts it reports a free-lunch share near one whatever its students' incomes. The rise in the free-lunch share after a shooting could then reflect adoption rather than who leaves. I flag the 5.7% of schools whose reported share jumps to 99% or more from 2014 (Appendix H.3). Dropping them raises the estimate from 0.014 (SE 0.008) to 0.018 (SE 0.008), and restricting the panel to the years before 2014 gives 0.011 (SE 0.009; Figure H.3). The point estimate does not shrink. Both variants soften the pre-trend failure noted in Section 6.2 without curing it: the joint test gives $p = 0.052$ without the adopters and $p = 0.043$ before 2014, against $p = 0.001$ in the main specification.

Randomized Timing Assignments. The Staggered Difference-in-Differences with the not-yet-treated and never-treated as a control group relies on the assumption of exogeneous timing of the shootings. If there were no true effect of shooting on log enrollment, then a reshuffling of

treatment years among the treated units would produce similar ATT estimates, and our ATT would fall well within the range of ATTs produced by this reshuffling.

To test this, I reshuffle the first-shooting years among the 2,550 ever-treated schools 500 times and rerun the main analysis on each permuted assignment. The actual overall ATT is -0.042. Nine of the 500 permuted estimates are at least as large in absolute value, a two-sided permutation p-value of $(1 + 9)/(1 + 500) = 0.020$ (Figure H.4). The permutation distribution is not centred on zero. Its mean is -0.024, because the reshuffle moves treatment years but leaves outcomes untouched, so permuted post-periods that overlap real post-shooting years inherit part of the true decline. Measured from that centre rather than from zero, the p-value is 0.040.

Repeat shootings. Identification requires the timing of a school's first shooting to be unpredictable. A related worry is that shootings recur at the same school, so that a first shooting signals more to come. Repeats are rare: 263 of the 2,813 school-years with a shooting follow an earlier shooting at the same school. Schools that ever have a shooting are more likely to have one in a given year than schools in general, which is the location component that the unit fixed effects absorb (Figure H.5, panel b). Within that group, a past shooting does not predict the next. In almost every year, a school that has already had a shooting is less likely to have another than the average ever-shooting school, and so less likely than one that has not yet had its first (panel a).

9 Conclusion

Neighborhoods that were previously on similar trajectories can experience widening socioeconomic gaps after violent events. It is not a random sample of families that leave the neighborhood in response to a violent shock. Selection along a socioeconomic gradient drives spatial inequality. Establishing this causally has been difficult however, because violence and segregation share many confounders and the causation runs in both directions. I exploit the exogenous timing of school shootings as a shock to safety as a local neighborhood amenity to identify the causal effect of violence on socioeconomic segregation, and I measure that effect at three nested scales: between schools, between districts, and between counties.

School enrollment falls by 4%, a loss of 32 students at the average treated school. The decline is immediate and grows over time, reaching 10% in the long run for fatal shootings. The exits are consistent with non-Free or Reduced-Lunch eligible students leaving: the share of eligible students among those who remain rises, though that estimate is imprecise. I find no increase in socioeconomic segregation at the between-school or between-district level, but an increase of 40% above the treated baseline between counties, emerging around three years after the shooting and growing thereafter.

Where families relocate depends on where the nearest safe alternative is. A spillover analysis shows the estimate is stable from 250 metres to 5 kilometres and remains negative and precisely estimated at 10 kilometres. It approaches zero at 20 kilometres and turns positive at 50 kilometres, implying that schools far enough away absorb the students that nearby schools lose. Schools in the same district are hit by the same shock, so switching to one of them buys no safety. The radius at which nearby schools stop losing students and start gaining them is roughly the radius at which segregation begins to rise. Both coincide with the between-county scale.

The county scale also squares my results with the individual-level evidence. Cabral et al. (2026) find no increase in school switching or exit among exposed students, and Sezer (2022) finds that exposed individuals move less by age 30, not more. Neither finding conflicts with mine, because we measure different populations on different margins. Exit is also only one way a school empties. Families who would have enrolled can enroll elsewhere, and a design that tracks exposed students will not capture that. A school's composition shifts through who stops arriving as much as through who leaves. Non-entry fits the timing too, since it accumulates one cohort at a time, which is what the three-year delay and the steady growth thereafter look like.

A simple model of sorting predicts the drop in enrollment, the selective out-mobility, and a larger drop where families have more outside options. The heterogeneity analyses bear this out: outflows are larger where intradistrict open enrollment is permitted and where the district offers more schools at the same level. Schools in the most violent quartile of counties see the steepest declines, which suggests a shooting can act as a last straw rather than as a bolt from an otherwise safe sky. One result runs the other way. Enrollment falls more where interdistrict transfers are barred, which may reflect the skewness of the sample, since 9 of the 10 most affected states have strict interdistrict rules.

The 4% estimate is conservative. Shooting schools close at a rate of 1 in 66 within five years, and if the schools that shut down are those that would have lost the most students, my estimates

are a lower bound. Neighboring schools also lose students, so any comparison school close to a shooting understates the gap. Both push the measured response toward zero.

The scale carries a policy implication. Widening school choice within a district operates on a margin where nothing happens, because the alternatives inside the district are hit by the same shock; where it does bite, it speeds the compositional decline for the families who cannot follow. Disadvantaged neighborhoods stay disadvantaged partly because violence moves the families who can afford to move, and it moves them past the district line.

References

Abdulkadiroğlu, A., P. A. Pathak, J. Schellenberg, and C. R. Walters. 2020. "Do parents value school effectiveness?" *American Economic Review* 110 (5): 1502–1539.

Abouk, R., and S. Adams. 2013. "School shootings and private school enrollment." *Economics Letters* 118 (2): 297–299.

Ang, D. 2021. "The effects of police violence on inner-city students." *The Quarterly Journal of Economics* 136 (1): 115–168.

Baird, A. A., E. V. Roellke, and D. M. Zeifman. 2017. "Alone and adrift: The association between mass school shootings, school size, and student support." *The Social Science Journal* 54 (3): 261–270.

Bayer, P., M. Casey, W. B. McCartney, J. Orellana-Li, and C. Zhang. 2025. "Distinguishing causes of neighborhood racial change: A nearest-neighbor design." *American Economic Review* 115 (11): 3999–4039.

Bayer, P., F. Ferreira, and R. McMillan. 2007. "A unified framework for measuring preferences for schools and neighborhoods." *Journal of political economy* 115 (4): 588–638.

Bayer, P., and R. McMillan. 2012. "Tiebout sorting and neighborhood stratification." *Journal of Public Economics* 96 (11-12): 1129–1143.

Beland, L.-P., and D. Kim. 2016. "The effect of high school shootings on schools and student performance." *Educational Evaluation and Policy Analysis* 38 (1): 113–126.

Bjerre-Nielsen, A., and M. H. Gandil. 2024. "Attendance boundary policies and the limits to combating school segregation." *American Economic Journal: Economic Policy* 16 (1): 190–227.

Black, S. E. 1999. "Do better schools matter? Parental valuation of elementary education." *The quarterly journal of economics* 114 (2): 577–599.

Boustan, L. P. 2013. "Racial residential segregation in American cities."

Brodeur, A., and H. Yousaf. 2025. "On the economic consequences of mass shootings." *Review of Economics and Statistics* 107 (1): 109–124.

Broughman, S., B. Kincel, J. Willinger, and J. Peterson. 2021. *Characteristics of Private Schools in the United States: Results From the 2019–20 Private School Universe Survey (First Look)*. Technical report NCES 2021-061. Average annual tuition in 2019–20 was $9,720 at Catholic schools and $10,910 at other religious schools, versus $19,590 at nonsectarian schools—

roughly two to three times higher. Washington, DC: National Center for Education Statistics, U.S. Department of Education. https://nces.ed.gov/pubs2021/2021061.pdf.

Burdick-Will, J., L. Gebo, and A. D. Williams. 2024. "Anywhere but here: Neighborhood violence and local school preferences in Baltimore City." *Sociology of Education* 97 (1): 1–20.

Burdick-Will, J., K. M. Nerenberg, J. A. Grigg, and F. Connolly. 2021. "Student mobility and violent crime exposure at Baltimore City public elementary schools." *American Educational Research Journal* 58 (3): 602–634.

Burgess, S., E. Greaves, A. Vignoles, and D. Wilson. 2015. "What parents want: School preferences and school choice." *The Economic Journal* 125 (587): 1262–1289.

Cabral, M., B. Kim, M. Rossin-Slater, M. Schnell, and H. Schwandt. 2026. "Trauma at School: The Impacts of Shootings on Students' Human Capital and Economic Outcomes." *The Review of Economic Studies* 93 (1): 327–365. https://doi.org/10.1093/restud/rdaf027.

Caetano, G. 2019. "Neighborhood sorting and the value of public school quality." *Journal of Urban Economics* 114:103193.

Callaway, B., and P. H. Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." Themed Issue: Treatment Effect 1, *Journal of Econometrics* 225 (2): 200–230. https://www.sciencedirect.com/science/article/pii/S0304407620303948.

Caudillo, M. L., and F. Torche. 2014. "Exposure to local homicides and early educational achievement in Mexico." *Sociology of education* 87 (2): 89–105.

Chetty, R., and N. Hendren. 2018. "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects." *The quarterly journal of economics* 133 (3): 1107–1162.

Chetty, R., N. Hendren, and L. F. Katz. 2016. "The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment." *American Economic Review* 106 (4): 855–902.

De Chaisemartin, C., and X. d'Haultfoeuille. 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American economic review* 110 (9): 2964–2996.

Deb, P., and A. Gangaram. 2024. "The effects of school shootings on risky behavior, health, and human capital." *Journal of Population Economics* 37 (1): 1–28.

Epple, D., and H. Sieg. 1999. "Estimating Equilibrium Models of Local Jurisdictions." *Journal of Political Economy* 107 (4): 645–681.

Epple, D. N., and R. Romano. 2003. "Neighborhood schools, choice, and the distribution of educational benefits." In *The economics of school choice,* 227–286. University of Chicago Press.

Erwin, B., B. Kelley, and G. Silva-Padrón. 2022. *50-State Comparison: Open Enrollment Policies.* Online. Last accessed 2026–07–31. Education Commission of the States. https://www.ecs.org/50-state-comparison-open-enrollment-policies/.

Fox, J. A., N. E. Sanders, E. E. Fridel, G. Duwe, and M. Rocque. 2021. "The contagion of mass shootings: The interdependence of large-scale massacres and mass media coverage." *Statistics and Public Policy* 8 (1): 53–66.

Freilich, J. D., S. M. Chermak, N. M. Connell, B. R. Klein, and E. A. Greene-Colozzi. 2022. "Using open-source data to better understand and respond to American school shootings: introducing and exploring the American School Shooting Study (TASSS)." *Journal of school violence* 21 (2): 93–118.

Gammell, S. P., N. M. Connell, and M. G. Huskey. 2022. "A descriptive analysis of the characteristics of school shootings across five decades." *American journal of criminal justice* 47 (5): 818–835.

García-Montoya, L., A. Arjona, and M. Lacombe. 2022. "Violence and voting in the United States: How school shootings affect elections." *American Political Science Review* 116 (3): 807–826.

Gujral, K., A. M. Ellyson, A. Rowhani-Rahbar, and F. Rivara. 2023. "The community impact of school-shootings on stress-related emergency department visits." *Contemporary Economic Policy* 41 (3): 455–470.

Hailey, C. A. 2025. "Choosing Schools, Choosing Safety: How Multiple Dimensions of Safety Shape School Choices." *AERA Open* 11:23328584251329899.

Hastings, J., T. J. Kane, and D. O. Staiger. 2009. "Heterogeneous preferences and the efficacy of public school choice." *NBER working paper* 2145:1–46.

Hodges, J. C., D. T. Walker, C. F. Baum, and S. S. Hawkins. 2023. "Impact of School Shootings on Adolescent School Safety, 2009–2019." *American journal of public health* 113 (4): 438–441.

Kissner, J. 2016. "Are active shootings temporally contagious? An empirical assessment." *Journal of Police and Criminal Psychology* 31 (1): 48–58.

Levine, P. B., and R. McKnight. 2020. *Not all school shootings are the same and the differences matter.* Technical report. National Bureau of Economic Research.

———. 2021. *Exposure to a school shooting and subsequent well-being.* Technical report. National Bureau of Economic Research.

———. 2024. "The consequences of high-fatality school shootings for surviving students." *Journal of Policy Analysis and Management.*

Livingston, M. D., M. E. Rossheim, and K. S. Hall. 2019. "A descriptive analysis of school and school shooter characteristics and the severity of school shootings in the United States, 1999–2018." *Journal of Adolescent Health* 64 (6): 797–799.

Morenoff, J. D., and R. J. Sampson. 1997. "Violent crime and the spatial dynamics of neighborhood transition: Chicago, 1970–1990." *Social forces* 76 (1): 31–64.

Muñoz-Morales, J., and R. Singh. 2023. "Do school shootings erode property values?" *Regional Science and Urban Economics* 98:103852.

NBC Bay Area. 2022. *Parents, Local Leaders Demand Change After Shooting at Oakland Campus.* Online. Last accessed 2023–10–30. https://www.nbcbayarea.com/news/local/east-bay/parents-leaders-demand-change-oakland-school-shooting/3018366/.

Owens, A., S. F. Reardon, D. Kalogrides, H. Jang, and T. Tom. 2022. "Trends in racial/ethnic and economic school segregation, 1991–2020." *The Segregation Index Research Brief.*

Pah, A. R., J. Hagan, A.-L. Jennings, et al. 2017. "Economic insecurity and the rise in gun violence at US schools." *Nature Human Behaviour* 1 (2): 1–6.

Reardon, S. F., A. D. Ho, B. R. Shear, et al. 2024. *Stanford Education Data Archive (Version 5.0)*. Retrieved from https://purl.stanford.edu/cs829jn7849; Last accessed 2024-10-01. https://purl.stanford.edu/cs829jn7849.

Reardon, S. F., and G. Firebaugh. 2002. "Measures of Multigroup Segregation." *Sociological Methodology* 32:33–67. Accessed September 17, 2025. http://www.jstor.org/stable/3186153.

Reeping, P. M., L. Klarevas, S. Rajan, et al. 2022. "State firearm laws, gun ownership, and K-12 school shootings: Implications for school safety." *Journal of school violence* 21 (2): 132–146.

Riedman, D. 2024. "K-12 school shooting database."

Rossin-Slater, M., M. Schnell, H. Schwandt, S. Trejo, and L. Uniat. 2020. "Local exposure to school shootings and youth antidepressant use." *Proceedings of the National Academy of Sciences* 117 (38): 23484–23489.

Sampson, R. J. 2012. *Great American city: Chicago and the enduring neighborhood effect*. University of Chicago press.

Scribner, H. 2022. *Parents, Students Return to Uvalde Schools Months After Mass Shooting*. Axios, Online. Published 2022–09–06. Last accessed 2026–07–31. https://www.axios.com/2022/09/06/uvalde-school-district-shooting-return-to-school.

Sezer, A. H. 2022. "The Legacy of School Shootings: The Long-Term and Intergenerational Effects." *CentER Discussion Paper,* https://research.tilburguniversity.edu/en/publications/the-legacy-of-school-shootings-the-long-term-and-intergenerationa.

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

———. 2013. *Stuck in place: Urban neighborhoods and the end of progress toward racial equality*. University of Chicago Press.

———. 2018. "The long reach of violence: A broader perspective on data, theory, and evidence on the prevalence and consequences of exposure to violence." *Annual Review of Criminology* 1:85–102.

Smith, G. 2023. *NL Native on Responding to One of the Country's Worst School Shootings*. The Day (New London, CT), Online. Published 2023–09–30; updated 2023–10–01. Last accessed 2026–08–11. The live page renders its byline client-side and leaves the author field empty in its structured data; byline ("By Greg Smith, Day Staff Writer") verified against the Internet Archive snapshot of 2024–08–04: https://web.archive.org/web/20240804075940/https://www.theday.com/local-news/20230930/nl-native-on-responding-to-one-of-the-countrys-worst-school-shootings/. https://theday.com/news/508791/nl-native-on-responding-to-one-of-the-countrys-worst-school-shootings/.

Smith, S. A., and R. Zimmer. 2022. "The Impacts of School District Consolidation on Rural Communities: Evidence from Arkansas Reform. EdWorkingPaper No. 22-530." *Annenberg Institute for School Reform at Brown University*.

Towers, S., A. Gomez-Lievano, M. Khan, A. Mubayi, and C. Castillo-Chavez. 2015. "Contagion in mass killings and school shootings." *PLoS one* 10 (7): e0117259.

U.S. Department of Agriculture, Food and Nutrition Service. 2025. "Child Nutrition Programs: Income Eligibility Guidelines." Notice. FR Doc. 2025-03821, *Federal Register* 90:11938. Accessed January 15, 2026. https://www.federalregister.gov/documents/2025/03/13/2025-03821/child-nutrition-programs-income-eligibility-guidelines.

Wang, K., A. Rathbun, and L. Musu. 2019. *School Choice in the United States: 2019.* Technical report NCES 2019-106. Washington, DC: National Center for Education Statistics. https://nces.ed.gov/pubs2019/2019106.pdf.

Wippell, J. G., D. L. Haynie, and D. Riedman. 2026. "Preventable Tragedies: A Longitudinal Analysis of State Firearm Laws and K-12 School Shootings in the United States (2000-2019)." *Social Science & Medicine,* 119049.

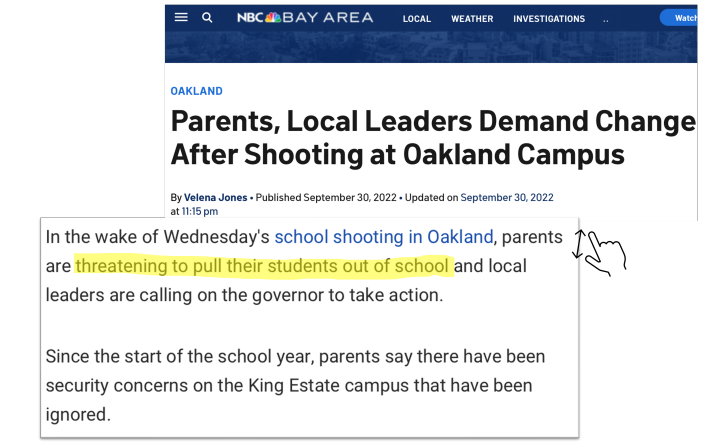
Yang, L., and M. Gopalan. 2023. "The Effects of Campus Shootings on School Finance and Student Composition." *Education Finance and Policy* 18 (2): 277–301.

Yousaf, H. 2021. "Sticking to one's guns: Mass shootings and the political economy of gun control in the United States." *Journal of the European Economic Association* 19 (5): 2765–2802.

Appendix

A News Reporting on Responses to Shootings

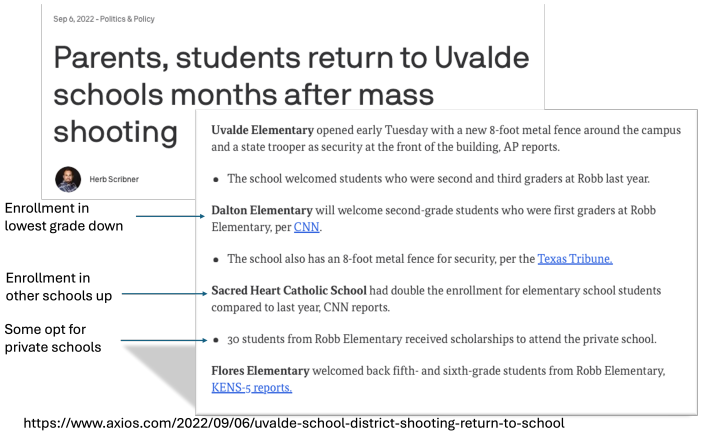
Figure A.1. News Reporting on Parental Responses and Nearby School Responses to Shooting



https://www.nbcbayarea.com/news/local/east-bay/parents-leaders-demand-change-oakland-school-shooting/3018366/



https://www.axios.com/2022/09/06/uvalde-school-district-shooting-return-to-school



https://www.axios.com/2022/09/06/uvalde-school-district-shooting-return-to-school

Notes: News reporting on parental responses to school shootings as well as enrollment spillovers to other schools.

B School Shootings databases

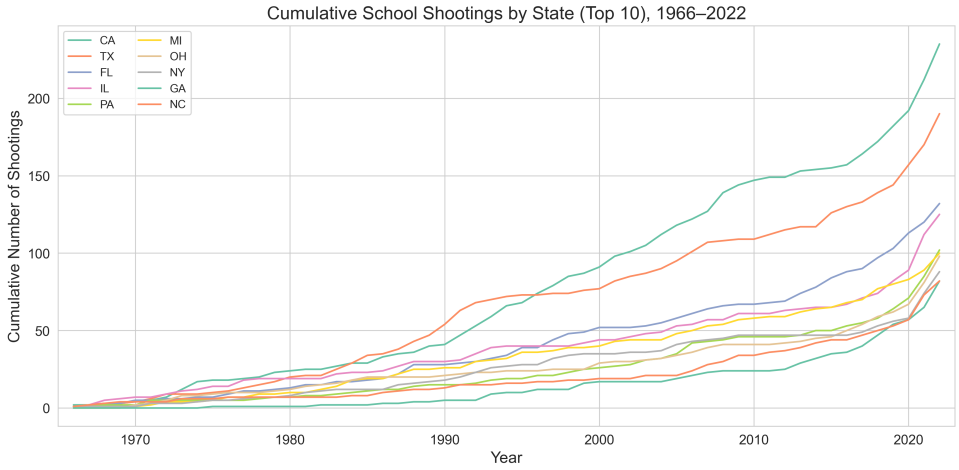
Table B.1. School Shootings Datasets Overview

Dataset	Period	# Events	Inclusion Criteria	Based on	Used by	Last accessed
Public K–12 School Shooting Database (K12SSD)	1970–2024	2,675	Gun brandished or gun fired or bullet hits school property irrespective of time or day.	Government agencies, media or advocacy groups, websites and blogs	(See CHDS)	February 9th 2024 (https://k12ssdb.org)
US School Gun Violence	1990–2013	381	firearm being discharged (inc. accidental), on a school campus, involves students or school employees, either as perpetrators, bystanders or victims.	Shultz dataset; Slate, Brady Campaign to Prevent Gun Violence, and Wikipedia.	Pah et al. 2017	October 23rd 2023 (https://amaral.northwestern.edu/resources/data-sets/gun-violence-us-schools)
Washington Post School Shootings Database	1999–2023	387	On campus immediately before/during/just after classes. Excludes: shootings at after-hours events, accidental discharges that don't injure anyone other than the perpetrator, suicides that occurred privately or posed no threat to other children	News articles, open-source databases, law enforcement reports, school websites, calls to schools & police	Cabral et al. 2026; Gujral et al. 2023; Livingston et al. 2019; Reeping et al. 2022; Rossin-Slater et al. 2020; Yang and Gopalan 2023	October 31st 2023 (https://github.com/washingtonpost/data-school-shootings)
CHDS K-12 School Shooting Database	1970–2022	2,070	Like K12SSD but not as up-to-date	(See K12SSD)	Cabral et al. 2026; Deb and Gangaram 2024; Gammell et al. 2022; Gujral et al. 2023; Hodges et al. 2023; Levine and McKnight 2020, 2021; Sezer 2022	January 23rd 2024 (www.chds.us/ssdb)
The American School Shooting Study (TASS)	1990–2016	652	Resulted in criminal justice response, firearm discharged, at K-12 school, on school grounds, & injured or killed at least one person	.	Freilich et al. 2022	N/A More information at (https://rockinst.org/wp-content/uploads/2022/08/Overview-American-School-Shooting-pdf)
School-Associated Violent Death	1992–2020	2,031	Homicide, suicide, or legal intervention death, fatal injury occuring on school campus, while victim to/from regular sessions or official school-sponsored event	.	Beland and Kim 2016	N/A (https://nces.ed.gov/programs/coe/indicator/a01/violent-deaths-and-shootings)
Stanford Mass Shootings of America	1966–2022	335	In an elementary, middle or high school, with 3+ victims excl. the perpetrator	Online reporting sources (at least 3 corrobarting)	Muñoz-Morales and Singh 2023	January 23rd 2024 (https://github.com/StanfordGeospatialCenter/MSA)

Notes: Summary of publically available datasets relating to school shootings.

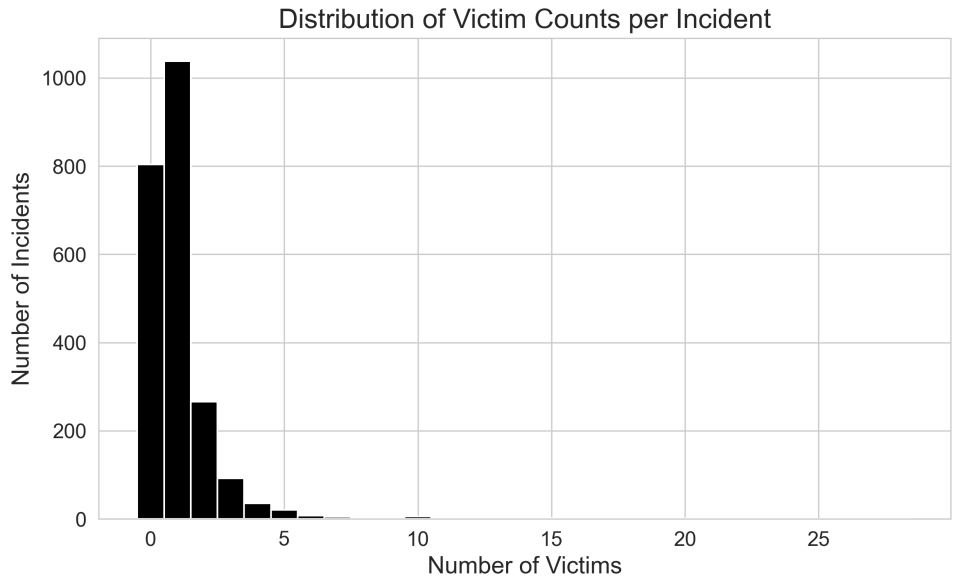
C Descriptive Statistics on Shootings

Figure C.1. Top 10 States with Shootings



Notes: Line plot of cumulative number of shootings per year and state for the top 10 states with most shootings in 2022.

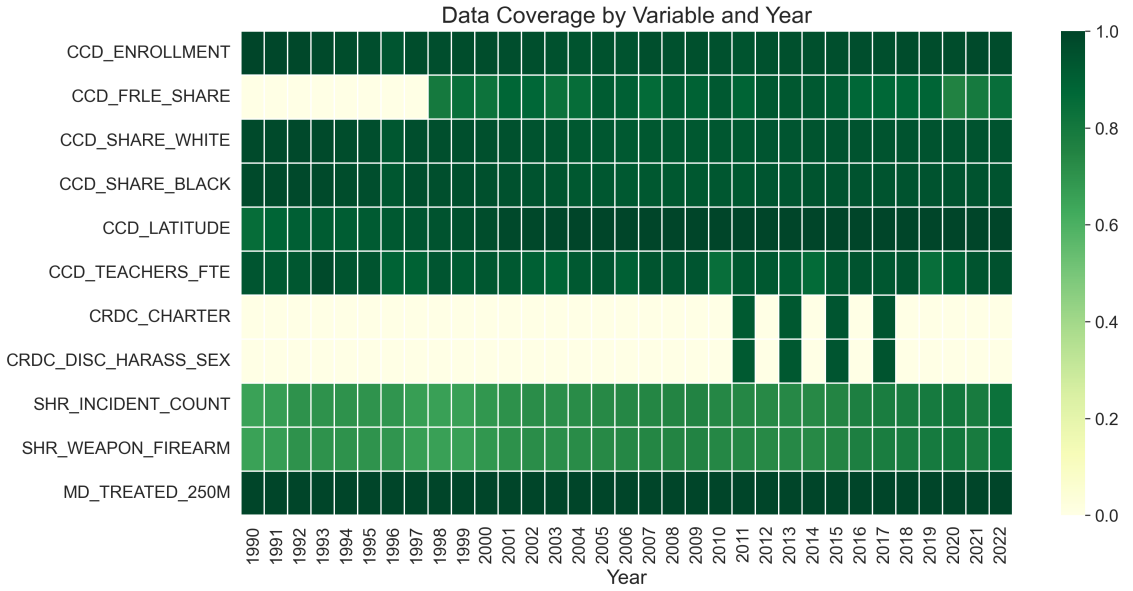
Figure C.2. Distribution of Number of Victims



Notes: Bar plot with count of number of victims.

D Data Coverage

Figure D.1. Missing data



Notes: Plot of data completeness over time. The variable name prefix reflects the data source. The prefix ”MD_” was added to own generated variables.

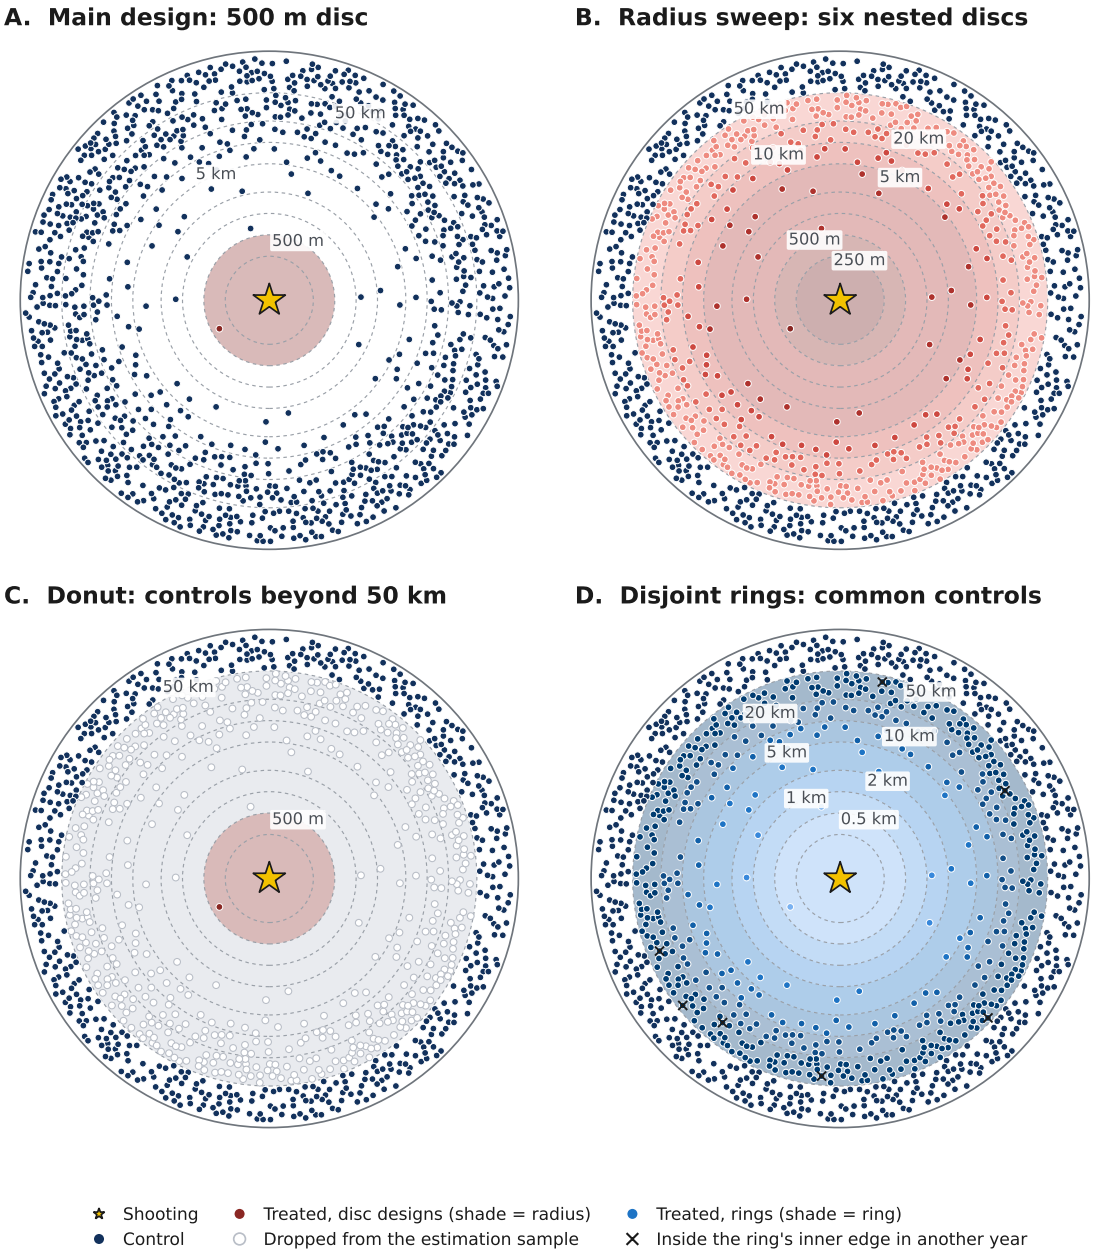
E Treatment Assignment Designs

The paper assigns treatment by distance in four ways. Figure E.1 draws each design on the same stylised neighbourhood of schools around one shooting, while Figure E.2 lays all of them out along a single distance axis.

The main design (panel A) treats a school if a shooting falls within 500 m of it in the same year. Every other school in a state $\times$ locale cell that holds a treated school is a control, so the control pool has no upper distance bound: a school 600 m from a shooting is a control. The radius sweep of Section 7.1 (panel B) re-runs the same rule at six radii from 250 m to 50 km. The discs are nested, so every disc contains the focal school. The state $\times$ locale restriction is re-applied at each radius, so the control pool shrinks as the disc grows. The donut of Section 8 (panel C) keeps the 500 m treatment but drops every control that was ever within 50 km of a shooting. The disjoint rings of Appendix G (panel D) treat a school whose nearest same-year shooting falls in a band $(a, b]$ and which was never within a metres of a shooting in any year. Every ring is measured against one common control pool, the schools never within 50 km of a shooting.

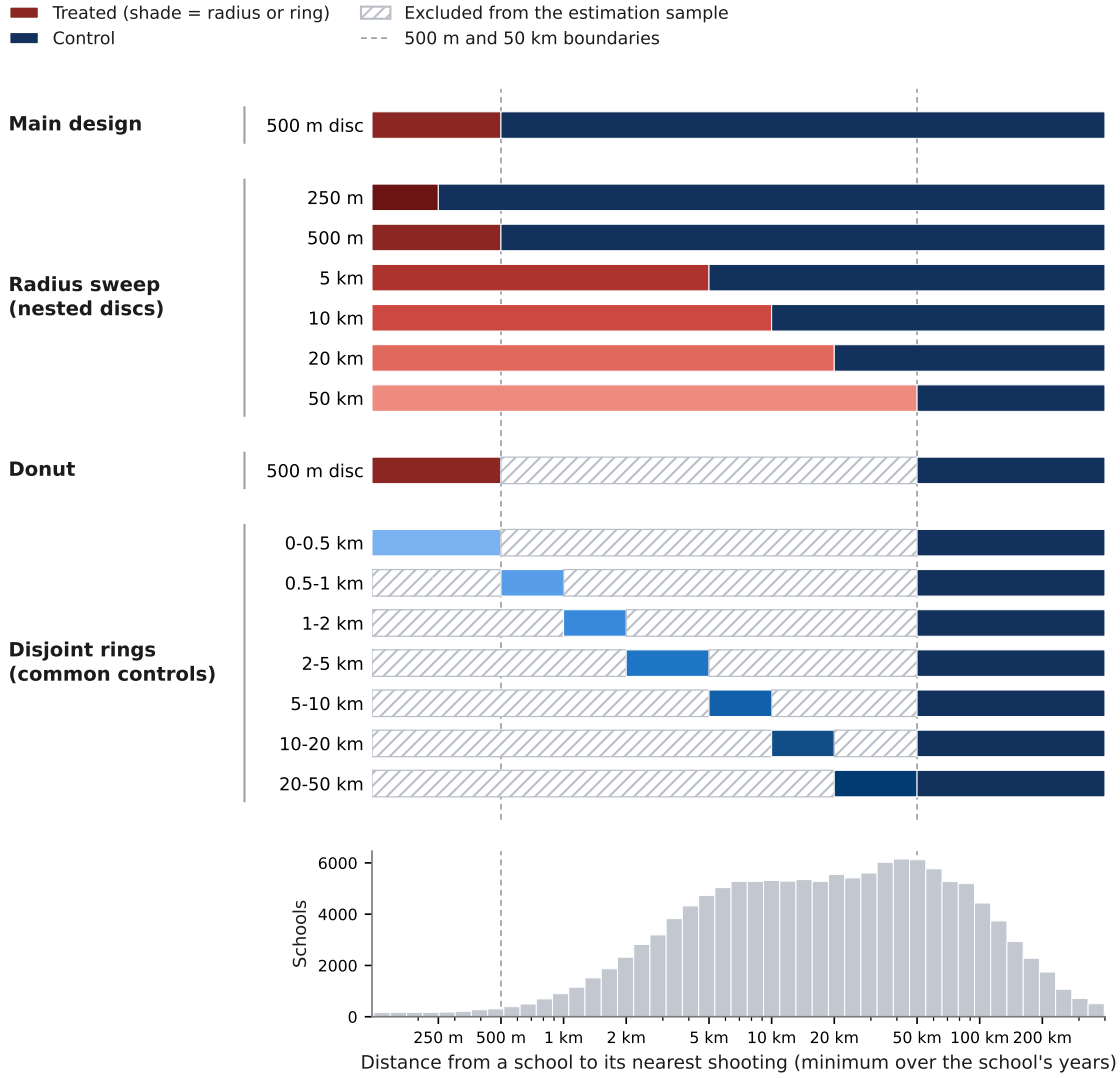
The disc designs and the rings measure distance differently. The discs use buffered spatial joins on each school's coordinates in that year, whereas ring membership uses each school's median coordinate over the panel, for the reason given in Appendix G. The histogram in Figure E.2 shows where schools sit relative to shootings. Few lie within 500 m of one and most lie between 2 and 200 km, which is why the wide discs treat a large share of all schools.

Figure E.1. Treatment Assignment Designs



Notes: Each panel applies one assignment rule to the same stylised neighbourhood: a shooting (star) and the schools around it (dots). Treated schools are shaded by radius in panel B and by ring in panel D; navy dots are controls, and hollow dots in panel C are controls dropped from the donut. The radial axis is logarithmic, from 60 m at the centre to 200 km at the rim, so that the 500 m and 50 km boundaries are both visible; areas are not to scale. Each distance band holds the median number of other schools found in that band around a treated school, and beyond 50 km about one school in four is drawn. Crosses in panel D mark schools whose nearest shooting falls in a ring but which were inside the ring's inner edge in another year, and which are therefore dropped from it. Estimation-sample counts for each design are given with its results.

Figure E.2. Treatment Assignment Designs along the Distance to the Nearest Shooting



Notes: Each row is one estimation sample, read along a school's distance to its nearest shooting: treated, control (navy), or excluded from the estimation sample (hatched). Control bars run to the right edge because no design bounds the control pool from above. Ring rows also exclude schools that were inside the ring's inner edge in any year. The histogram counts schools by the minimum, over the years in which they are observed, of their distance to the nearest shooting; schools with no usable coordinate in any year are omitted. The distance axis is logarithmic.

F Further notes on mobility

At the student-year level, Cabral et al. (2026) find that the probability of an individual student leaving the school is not statistically significantly different between students at shooting schools and control schools. I write this as follows:

$$\Pr\left(\text{Leave}_{i,t}|\text{ShootingSchool}_{s,t-1}=1\right) \approx \Pr\left(\text{Leave}_{i,t}|\text{ShootingSchool}_{s,t-1}=0\right) \tag{9}$$

First, I argue that since Texas has no state-wide mandatory open enrollment laws, the probability of any student switching schools is low to begin with (i.e. $\Pr\left(\text{Leave}_{it}\right)=0+\varepsilon$ for some arbitrarily small ε). The conditional probabilities of switching school given a school shooting would have to be very large to yield any significant difference.

Second, I compare this to school-year level estimates of enrollment. For example, Yang and Gopalan (2023), estimating a decline of around four to five percent, identifies something similar to the following:

$$\mathbb{E}(\Delta\text{Enrollment}_{s,t}|\text{ShootingSchool}_{s,t-1}=1) < \mathbb{E}(\Delta\text{Enrollment}_{s,t}|\text{ShootingSchool}_{s,t-1}=0) \tag{10}$$

Specifically, I argue that the approximate equality in Equation 9 does not force the two sides of Equation 10 to be equal: a null at the student level is consistent with a decline at the school-year level.

We can divide the change in enrollment into stayers, leavers and joiners. Let $S_{s,t}$ be the number of students enrolled in school s in both $t-1$ and t, $L_{s,t}$ the number enrolled at $t-1$ but not at t, and $J_{s,t}$ the number enrolled at t but not at $t-1$. Enrollment in the two periods is

$$\text{Enrollment}_{s,t-1} = S_{s,t} + L_{s,t}, \qquad \text{Enrollment}_{s,t} = S_{s,t} + J_{s,t}, \tag{11}$$

so the stayers are common to both periods and cancel in the difference:

$$\Delta\text{Enrollment}_{s,t} \;=\; J_{s,t} - L_{s,t}. \tag{12}$$

Only the two flows enter the change in enrollment; the incumbent stock does not.

If we redo the Cabral et al. (2026) analysis at the school-year level, we might well find that there is no difference in the probability of students switching schools, but still find an overall negative effect on enrollment. Differencing (12) across shooting and non-shooting schools,

$$\begin{aligned}
&\mathbb{E}\big(\Delta\text{Enrollment}_{s,t} \mid \text{ShootingSchool}_{s,t-1}=1\big) - \mathbb{E}\big(\Delta\text{Enrollment}_{s,t} \mid \text{ShootingSchool}_{s,t-1}=0\big) \\
&\quad= \underbrace{\Big[\mathbb{E}\big(J_{s,t} \mid \text{ShootingSchool}_{s,t-1}=1\big) - \mathbb{E}\big(J_{s,t} \mid \text{ShootingSchool}_{s,t-1}=0\big)\Big]}_{\text{joiner margin}} \\
&\quad- \underbrace{\Big[\mathbb{E}\big(L_{s,t} \mid \text{ShootingSchool}_{s,t-1}=1\big) - \mathbb{E}\big(L_{s,t} \mid \text{ShootingSchool}_{s,t-1}=0\big)\Big]}_{\text{leaver margin}}.
\end{aligned} \tag{13}$$

Equation (13) shows that a decline in total enrollment can arise through either flow: fewer students joining, or more students leaving. The estimates in this paper locate it in the leaver margin. Enrollment in the lowest grade a school offers is unchanged (0.005, standard error 0.039), while enrollment in the highest grade falls by about five percent (-0.050, standard error 0.020; Table 3).

A shortfall of joiners would show up at the entry grade, where intake occurs, and it does not. A shortfall concentrated at the top of the grade range points instead to incumbent students leaving before completing the grades their school offers.

This remains consistent with Equation 9, for the reason given above. The per-student probability of leaving is small to begin with, so a four percent decline at a school of the average treated size amounts to roughly thirty departures over a year. Spread across a student body, that is a difference in individual switching probabilities well inside the confidence intervals of a student-level design. The two levels of analysis are also not measuring the same population: Cabral et al. (2026) follow matched students within a single state, whereas the enrollment counts used here form a national panel of schools. A null on individual switching and a negative effect on school enrollment can therefore hold at the same time.

G Spillovers on Disjoint Distance Rings

The radius sweep in Section 7.1 expands *discs*. A school inside the 250 m definition is also inside the 500 m definition and inside the 50 km definition, so every estimate in that sweep still contains the focal school's own effect, and each radius is measured against a different control pool. Because a disc average is a weighted blend of everything inside it, it crosses zero at a larger radius than the underlying ring does. This appendix re-runs the same question on disjoint rings, which separate the effect at a distance from the effect it nests.

Ring $(a, b]$ treats a school whose nearest same-year shooting ever falls in $(a, b]$ and which is never within a metres of any shooting in any year, so each ring excludes both the inner rings and the focal school (Figure E.1, panel D). Two features differ from the disc sweep. Ring membership is measured from each school's median coordinate rather than its year-specific one, because roughly a quarter of schools move more than a kilometre across the panel as the CCD re-geocodes them, which nested discs absorb but a 500 m-wide ring does not. And the control pool is common to every ring—schools never within 50 km of any shooting in any observed year—because comparing effects across rings requires one shared counterfactual.

Table G.1 reports the rings, and it should be read from the pre-trend columns inward. Two rings are clean. The innermost reproduces the headline effect at -0.040 with a pre-trend p of 0.81, and the 5–10 km ring puts the effect at -0.002, indistinguishable from zero, with $p = 0.50$. Four fail outright, and Figure G.2 shows why. The two outer rings drift upward from roughly ten years before treatment and simply continue through it, so their positive coefficients are a pre-existing trend rather than a response to a shooting. The 0.5–1 km and 1–2 km rings drift downward over the same window, so their large negative coefficients are partly earned before the shooting happens.

The 2–5 km ring is the awkward case and is reported as such. It rejects the joint pre-trend test at $p = 0.028$, but its largest pre-treatment coefficient is 0.005—the second smallest in the table, and smaller than that of the innermost ring, which passes comfortably. With 7,518 treated schools the test has enough power to reject on movements too small to matter, which is why the table reports the largest pre-treatment coefficient beside the p-value: the two columns disagree here, and neither alone settles it. I therefore read the 2–5 km estimate of -0.030 as suggestive of a genuine spillover onto schools that were themselves never near a shooting, not as established.

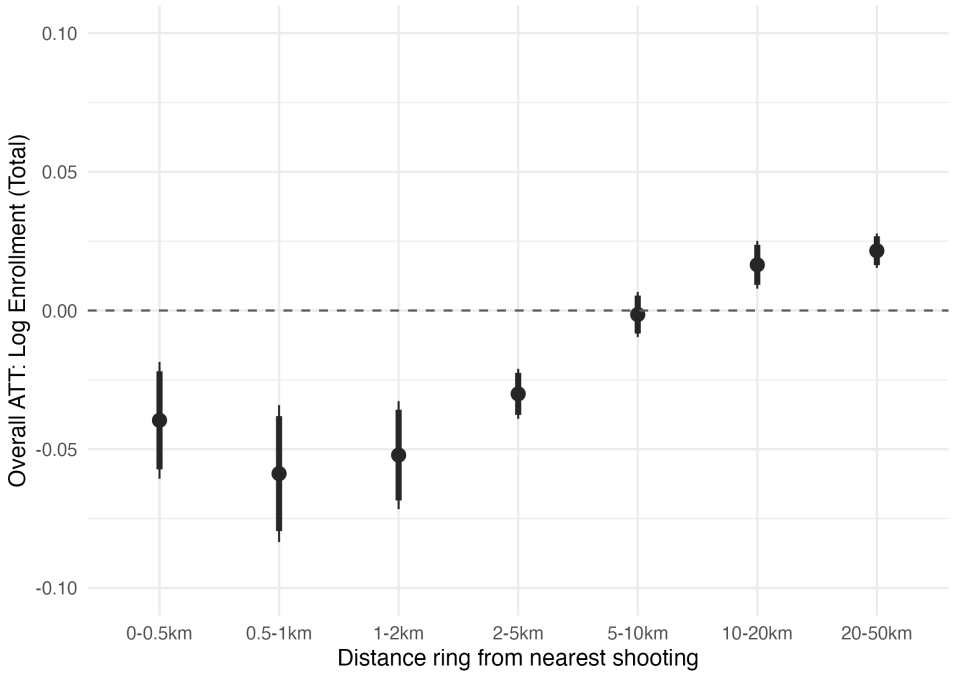
The composition columns explain the pattern. The control pool is fixed at roughly 9% City and 48% Rural, while the treated schools rotate from 54% City in the 1–2 km ring to 10% City in the 20–50 km ring. Urban schools were already on a different enrollment path from rural ones, and the rings that fail their pre-trends are the rings whose treated schools least resemble the control pool. The exercise therefore supports the main text on the near margin, where enrollment losses extend past the focal school, but it does not identify the distance at which those losses stop, because the rings that would pin it down are the contaminated ones.

Table G.1. Enrollment Effects on Disjoint Distance Rings

Ring	ATT	SE	Pre-trend p	Max |pre|	Treated	Control	Treated locale	Control locale
0-0.5km	-0.0396	(0.0107)	0.812	0.0073	1,261	11,333	City 43%, Rural 13%, Suburb 31%, Town 13%	City 9%, Rural 48%, Suburb 11%, Town 32%
0.5-1km	-0.0588	(0.0126)	0.000	0.0580	900	10,329	City 51%, Rural 7%, Suburb 31%, Town 11%	City 10%, Rural 47%, Suburb 11%, Town 32%
1-2km	-0.0521	(0.0099)	0.018	0.0232	2,485	11,501	City 54%, Rural 5%, Suburb 33%, Town 8%	City 9%, Rural 48%, Suburb 11%, Town 32%
2-5km	-0.0301	(0.0046)	0.028	0.0047	7,518	11,745	City 46%, Rural 6%, Suburb 42%, Town 6%	City 9%, Rural 49%, Suburb 11%, Town 31%
5-10km	-0.0015	(0.0042)	0.495	0.0036	8,140	11,754	City 34%, Rural 12%, Suburb 48%, Town 7%	City 9%, Rural 49%, Suburb 11%, Town 31%
10-20km	0.0165	(0.0044)	0.000	0.0283	7,887	11,836	City 17%, Rural 25%, Suburb 41%, Town 17%	City 9%, Rural 49%, Suburb 11%, Town 32%
20-50km	0.0216	(0.0032)	0.001	0.0158	11,894	12,093	City 10%, Rural 38%, Suburb 23%, Town 29%	City 9%, Rural 49%, Suburb 10%, Town 32%

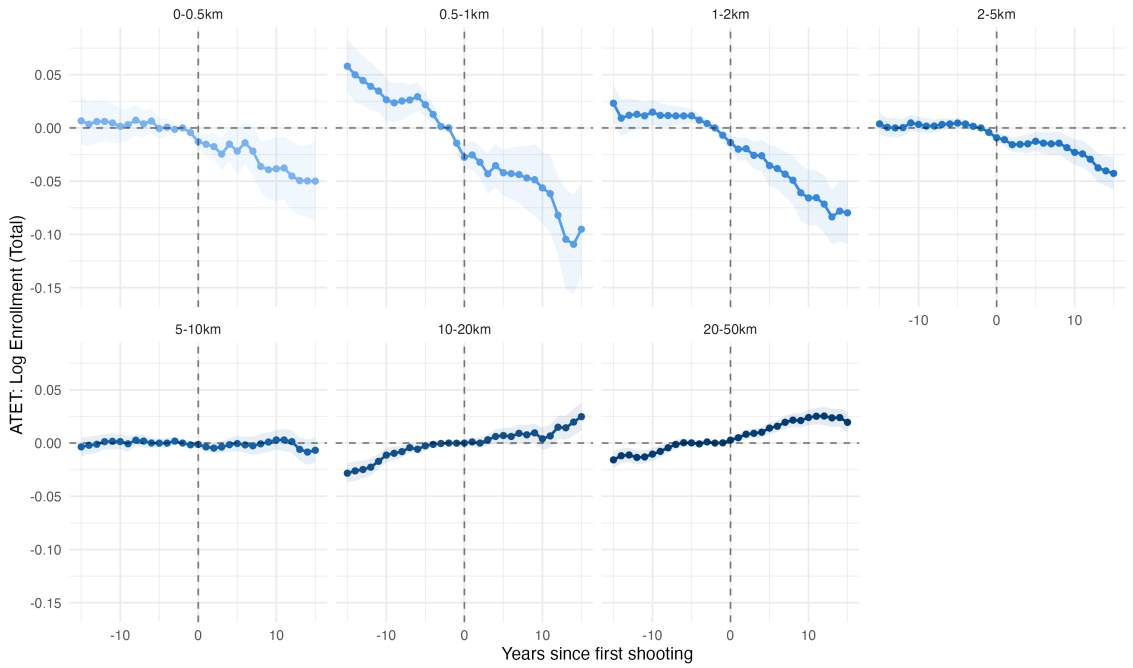
Notes: Each row is a separate estimation on one distance ring. Ring $(a, b]$ treats schools whose nearest same-year shooting falls in $(a, b]$ and which are never within a metres of any shooting, so rings are disjoint and exclude the focal school. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The control pool is common to every ring: schools never within 50 km of any shooting in any observed year. Pre-trend p is the joint Wald test on the aggregated event-study coefficients; Max |pre| is the largest pre-treatment coefficient in absolute value, reported alongside it because at these sample sizes the test rejects on movements too small to matter. The locale columns give the composition of the estimation sample, measured after panel balancing, and are the reason the failing rings fail.

Figure G.1. Overall ATT by Disjoint Distance Ring



Notes: Overall ATT on log enrollment for each disjoint ring, with 90% and 95% confidence intervals. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Point estimates beyond 10 km are positive but fail their pre-trend tests and should not be read as effects; see Figure G.2.

Figure G.2. Event Studies by Disjoint Distance Ring



Notes: One panel per disjoint ring, plotting the dynamic ATT on log enrollment against years since the shooting. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. The two outermost panels rise steadily from roughly ten years before treatment and continue through it, which is why their positive overall estimates are not read as effects.

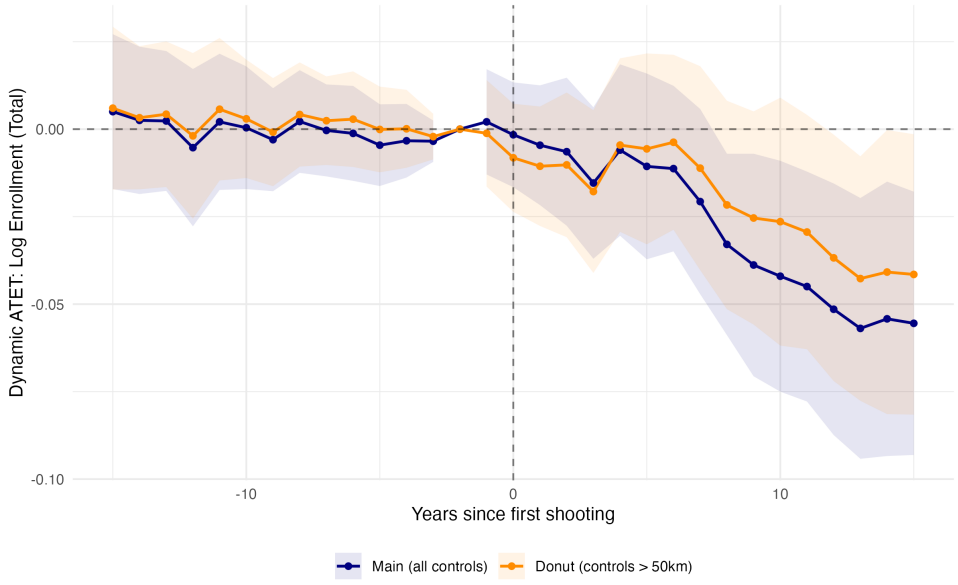
H Robustness Checks

This appendix collects the figures behind the robustness checks in Section 8.

H.1 Donut Difference-in-Differences

The donut keeps the 500 m treatment and drops every control that was ever within 50 km of a shooting (Figure E.1, panel C). Proximity is computed on the unrestricted panel, before the state $\times$ locale restriction.

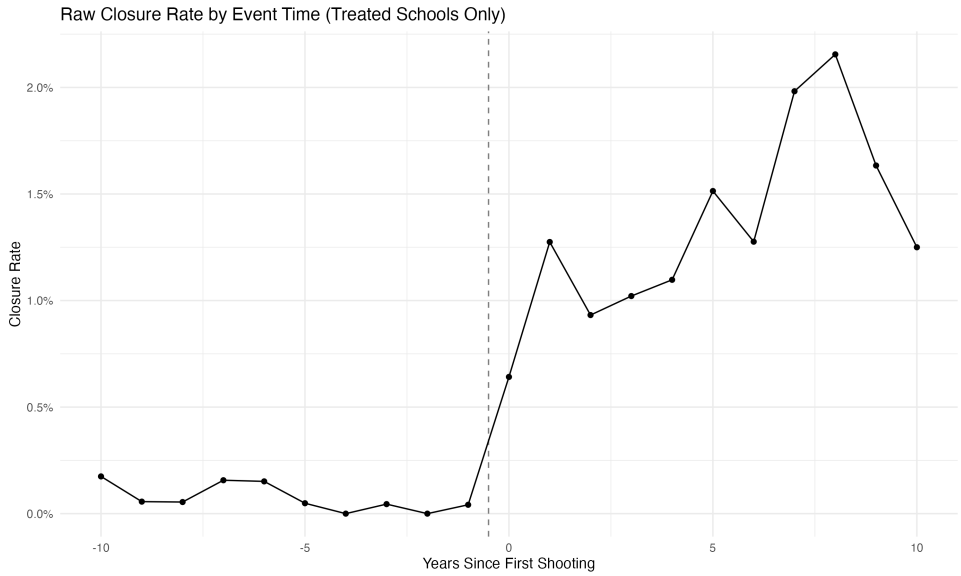
Figure H.1. Donut DiD: Excluding Controls Within 50km of a Shooting



Notes: Dynamic ATT on log enrollment by event time relative to the year before the first shooting, comparing the main specification with a donut specification that drops every control school ever within 50 km of a shooting. Treatment is still defined at the 500 m radius. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The donut specification keeps all 1,307 treated schools but only 11,033 of the 51,568 controls (407,220 school-year observations), since proximity is computed on the unrestricted panel and most controls sit within 50 km of some shooting. The overall ATT is -0.033 (SE 0.013) in the donut specification, against -0.042 (SE 0.012) in the main specification. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. A joint Wald test of the pre-treatment coefficients ($e \in [-15, -1]$) does not reject flat pre-trends in the donut specification ($p = 0.942$, against $p = 0.917$ in the main specification).

H.2 School Closures

Figure H.2. Closure Rates among Shooting Schools

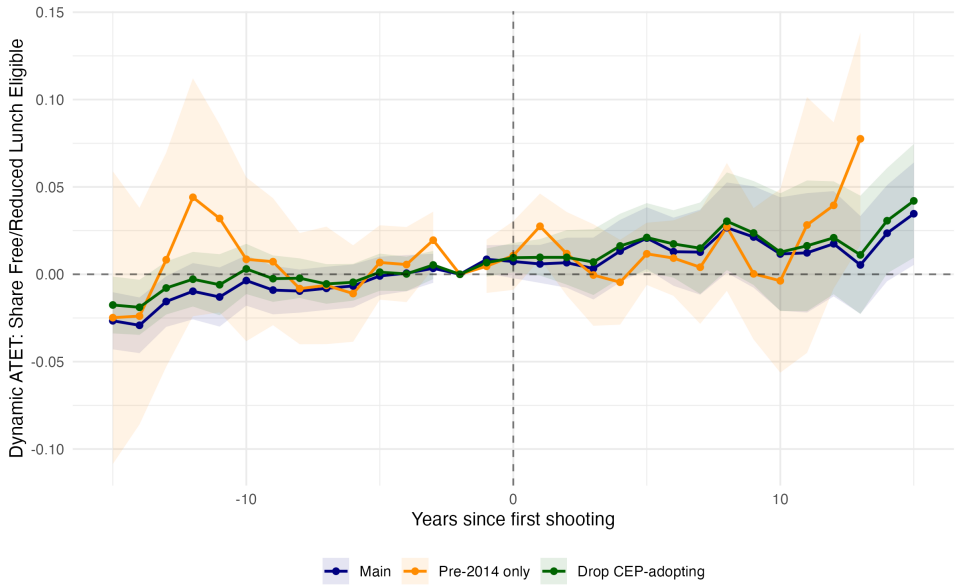


Notes: Raw closure rates by event time among ever-treated schools, where closure is an operational status of "closed" in the Common Core of Data.

H.3 Community Eligibility Provision

A school is flagged as adopting when its highest reported free-lunch share before 2014 is below 99% and its highest share from 2014 on is at least 99%. The rule flags 8,571 of 149,716 schools (5.7%). Because it reads adoption off the reported share alone, it misses adopters whose share had already reached 99% before 2014, and it flags any school whose share reaches 99% for another reason.

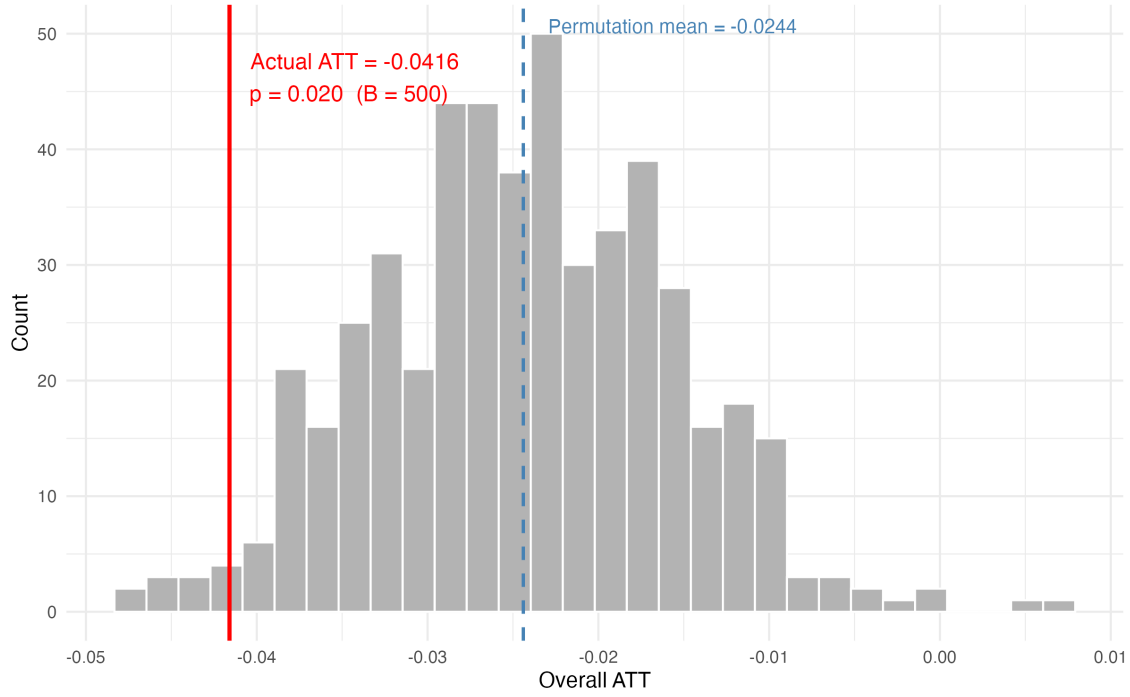
Figure H.3. Free Lunch Eligible, Excluding Community Eligibility Provision Adopters



Notes: Dynamic ATT on the share of students eligible for free or reduced-price lunch, by event time relative to the year before the first shooting, for the main specification, a panel restricted to the years before 2014, and a sample that drops the schools flagged as adopting the Community Eligibility Provision. Coefficients are percentage-point changes. Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. The main estimate uses 600 treated and 32,576 control schools, the pre-2014 estimate 784 and 41,255, and the estimate without adopters 522 and 29,928, each on a balanced panel. The overall ATTs are 0.014 (SE 0.008), 0.011 (SE 0.009) and 0.018 (SE 0.008) respectively. Standard errors come from the multiplier bootstrap with 1,000 replications. Bands are pointwise 95% confidence intervals. Joint Wald tests of the pre-treatment coefficients ($e \in [-15, -1]$) reject flat pre-trends in the main specification ($p = 0.001$) and in the pre-2014 panel ($p = 0.043$), and fall just short of rejecting without adopters ($p = 0.052$).

H.4 Randomization Inference

Figure H.4. Permutation Distribution of ATT

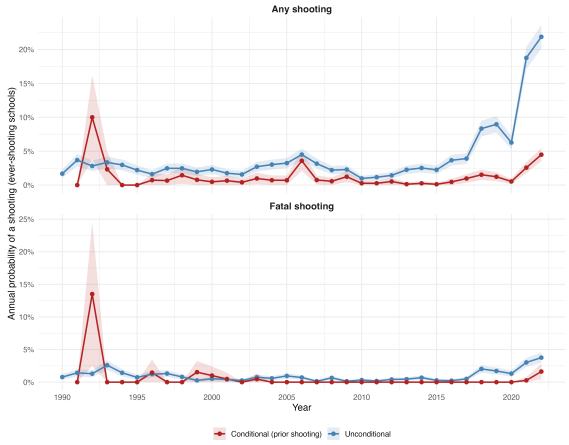


Notes: Permutation distribution of the overall (simple) ATT on log enrollment from 500 runs in which first-shooting years are reshuffled among the 2,550 ever-treated schools. The solid line is the actual ATT of -0.042 from the main analysis, and the dashed line is the mean of the permutation distribution, -0.024. Each run applies the identical specification to the reshuffled data: Estimates are group-time average treatment effects from the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), estimated by outcome regression (regression adjustment) and aggregated by event time. The comparison group is the not-yet-treated, which includes never-treated units. One anticipation period is allowed and group-time effects use a universal base period, so the event-study coefficient at $e = -2$ is zero by construction and is excluded from the joint pre-trend tests reported below. Covariates in the regression adjustment are school level and locale; standard errors are clustered on the school. The control pool is restricted to schools sharing a (state $\times$ locale) cell with at least one treated school. The restriction only removes control units that could never serve as a counterfactual but does not force comparisons within cell. Years are reshuffled before the sample filters, which are then re-applied to each permuted assignment, so the set of treated schools that survives balancing differs across runs. Inference here is the two-sided permutation p-value $(1 + k)/(1 + B)$, where $k = 9$ of the $B = 500$ runs are at least as extreme in absolute value as the actual ATT, giving $p = 0.020$; measured from the permutation mean instead of zero it is 0.040.

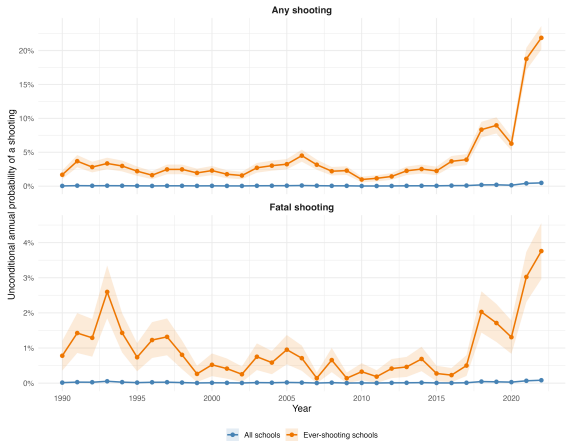
H.5 Repeat Shootings

The exercise uses every school-year in the panel: 3,160,215 school-years at 153,796 schools, of which 66,605 belong to the 2,550 schools that ever have a shooting. A shooting here is the main design's treatment, a shooting within 500 m of the school in that year. Of the 2,813 school-years with a shooting, 263 follow an earlier shooting at the same school; of the 642 with a fatal shooting, 22 follow an earlier fatal one. Across all schools, a shooting occurs in about 0.09% of school-years.

Figure H.5. Annual Probability of a Shooting



(a) Ever-shooting schools: given a prior shooting, and unconditional



(b) Unconditional: ever-shooting schools and all schools

Notes: Panel (a): among the 2,550 schools that ever have a shooting, the annual probability of a shooting given a shooting at the same school in an earlier year (red) and without that condition (blue). Panel (b): the unconditional annual probability among ever-shooting schools (orange) and among all schools (blue). In each panel the top row is any shooting and the bottom row fatal shootings, where a fatal repeat requires an earlier fatal shooting. Bands are pointwise 95% normal-approximation intervals for a proportion. Because the unconditional series in panel (a) pools the years before and after a school's first shooting, a conditional probability below it is also below the probability for schools that have not yet had a shooting.

I Household Mobility in the American Community Survey

I.1 Data and Design

The school data in the main text count students but do not show which families move. The American Community Survey (ACS) does. It asks every person aged one or older where they lived one year earlier. I use the IPUMS USA one-year samples for 2012–2021.

The finest geography at which both ends of a move are observed is the migration PUMA, a group of one or more Public Use Microdata Areas of at least 100,000 residents each. I use the 2010 definitions, which cover 982 migration PUMAs and apply to the origin of a move in exactly these ten years. For children aged 5–17 in households, I compute four annual rates in each migration PUMA: the share who moved house, the share who moved within the migration PUMA, the share who moved in from another U.S. migration PUMA, and the share who moved out. The first three are shares of the children living in the unit at the interview. The out-move rate is a share of the children who lived in the unit a year earlier, so it measures the response of the families who were there before. Rates use person weights.

I compute the same four rates for three groups, each against a comparison group. The first is children with a BA+ parent, meaning that the household head or the head's spouse or partner holds a bachelor's degree or more, against children aged 5–17 without one. The second is low-income children, in families below 200% of the poverty line, against children aged 5–17 at or above it. The third is children aged 1–4 against children aged 5–17. The outcome of interest is the gap: a group's rate minus its comparison group's rate in the same migration PUMA and year. Income and education are measured at the interview, after any move.

Each shooting is assigned to the migration PUMA containing its school, and a migration PUMA is treated from the calendar year of its first shooting. Units with any shooting between 1990 and 2010 are dropped, since they are neither clean controls nor first treated inside the panel. Alaska and Hawaii are dropped for lack of school data. Controls are restricted to states with at least one treated unit. The estimation sample has 661 migration PUMAs, 193 of them treated.

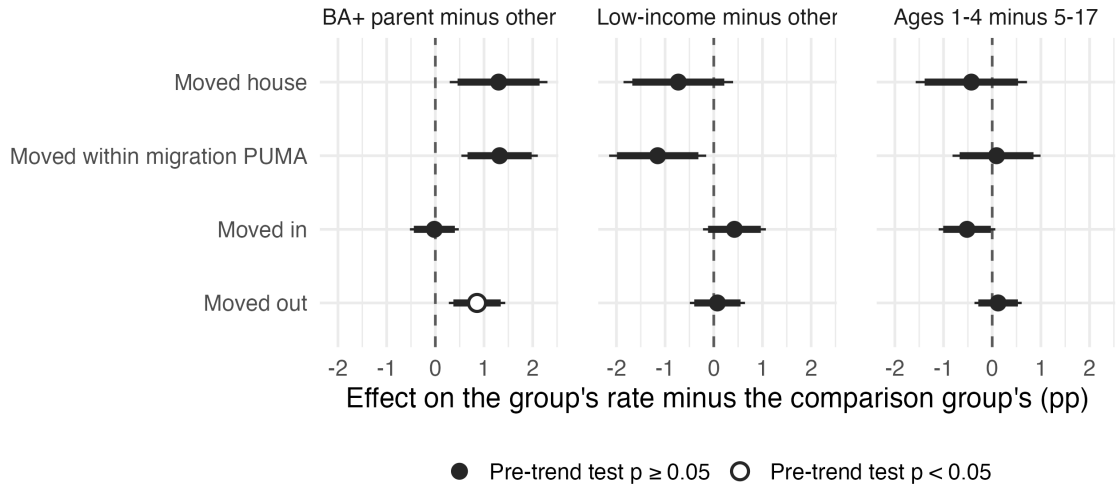
The main estimator is the imputation estimator of Borusyak, Jaravel and Spiess (2024). Unit and year effects are fitted on the untreated unit-years, the untreated outcome is imputed for every treated unit-year, and the effect is the average difference. Each treated unit-year is therefore compared with its unit's mean over all its pre-treatment years rather than with a single base year. This matters here, because a single ACS year in one migration PUMA is noisy. Standard errors come from a cluster bootstrap over migration PUMAs with 499 draws. The pre-trend test regresses the untreated outcomes on indicators for the five years before the shooting.

The base period is the year before the shooting, with no anticipation. The ACS asks about the past twelve months, so the file for the year before the shooting cannot contain a move made after it. Units first treated in 2011 or 2012 have no such year in the panel and drop out. The public ACS does not report the month of interview, so the file for the year of the shooting mixes moves made before and after it, and event time 0 is only partly treated.

I.2 Who Moves

Figure I.1 reports every gap on one axis. Children with a BA+ parent move more after a shooting. Low-income children and children aged 1–4 do not, and no gap for children aged 1–4 is significant at 5%.

Figure I.1. Who Moves after a Shooting: Families with a BA+ Parent

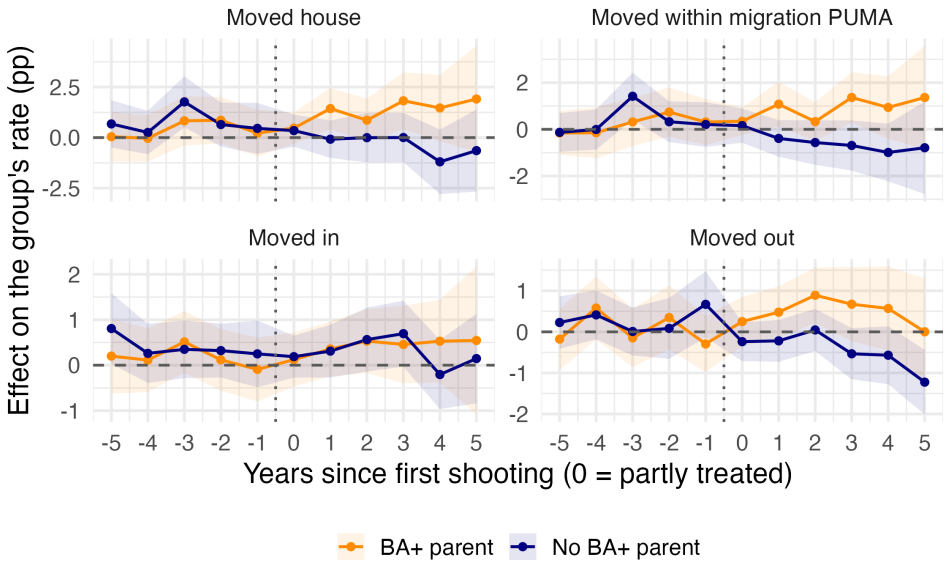


Notes: Overall ATT of a shooting on a group's mobility rate minus the comparison group's rate in the same migration PUMA and year, in percentage points. The comparison groups are children aged 5–17 without a BA+ parent, children aged 5–17 at or above 200% of the poverty line, and children aged 5–17. Estimates come from the imputation estimator: unit and year effects are fitted on the untreated unit-years, and the effect is the average difference between each treated unit-year's outcome and its imputed untreated outcome. Standard errors come from a cluster bootstrap over migration PUMAs with 499 draws. The unit is the 2010 migration PUMA, observed annually over 2012–2021: 661 migration PUMAs, 193 of them treated, on a balanced panel. Thick bars are 90% and thin bars 95% confidence intervals. A hollow point marks a gap whose pre-trend test rejects at the 5% level.

Figure I.2 separates the two sides of the BA+ gap. Children with a BA+ parent become 1.24 percentage points (pp) more likely to move house (SE 0.34), on a base of 9.4%. They become 0.85 pp more likely to move within the migration PUMA (SE 0.32), on a base of 5.6%. For children without a BA+ parent the same effects are -0.06 and -0.47. The gaps are $+1.30$ (SE 0.51) and $+1.32$ (SE 0.40), and neither fails its pre-trend test ($p = 0.77$ and $p = 0.36$). Most of the extra moves therefore stay inside the migration PUMA. The BA+ gap in moving out is $+0.86$ (SE 0.30), but it fails its pre-trend test ($p = 0.04$), so I do not read it as an effect.

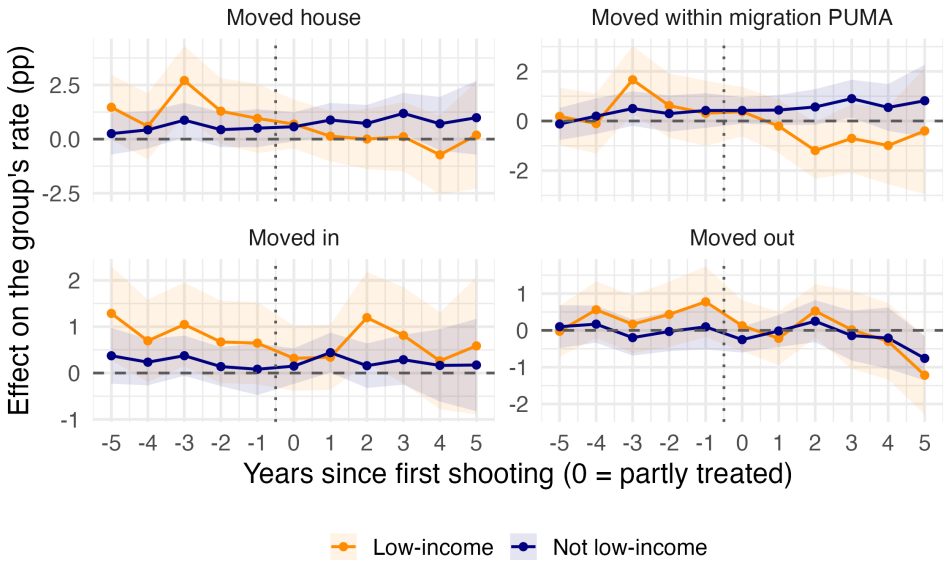
Figure I.3 shows the same split by income. Within the migration PUMA, children above the low-income line move more ($+0.64$ pp, SE 0.25, on a base of 6.3%). Low-income children do not (-0.51, SE 0.44, on a base of 11.6%). The gap is -1.15 (SE 0.51; pre-trend $p = 0.61$), so the local moves come from better-off families. Low-income children do move in more often ($+0.70$ pp, SE 0.27, on a base of 5.0%), but not more than other children (gap $+0.42$, SE 0.33).

Figure I.2. BA+ Families Move, Mostly within Their Migration PUMA



Notes: Effect of a shooting on the mobility rate of children aged 5–17 with and without a BA+ parent, in percentage points, by years since the migration PUMA's first shooting. Estimates come from the imputation estimator: unit and year effects are fitted on the untreated unit-years, and the effect is the average difference between each treated unit-year's outcome and its imputed untreated outcome. Standard errors come from a cluster bootstrap over migration PUMAs with 499 draws. Post-shooting coefficients are imputation estimates. Pre-shooting coefficients come from the pre-trend regression of untreated outcomes on leads, so every coefficient is measured against the unit's mean over its pre-treatment years and none is zero by construction. Event time 0 is only partly treated. The unit is the 2010 migration PUMA, observed annually over 2012–2021: 661 migration PUMAs, 193 of them treated, on a balanced panel. Bands are pointwise 95% confidence intervals. The figure stops at five years after the shooting; the overall ATTs also average event times six to eight, which rest on the earliest cohorts.

Figure I.3. The Local Moves Come from Better-Off Families



Notes: Effect of a shooting on the mobility rate of children aged 5–17 below and at or above 200% of the poverty line, in percentage points, by years since the migration PUMA's first shooting. Estimates come from the imputation estimator: unit and year effects are fitted on the untreated unit-years, and the effect is the average difference between each treated unit-year's outcome and its imputed untreated outcome. Standard errors come from a cluster bootstrap over migration PUMAs with 499 draws. Post-shooting coefficients are imputation estimates. Pre-shooting coefficients come from the pre-trend regression of untreated outcomes on leads, so every coefficient is measured against the unit's mean over its pre-treatment years and none is zero by construction. Event time 0 is only partly treated. The unit is the 2010 migration PUMA, observed annually over 2012–2021: 661 migration PUMAs, 193 of them treated, on a balanced panel. Bands are pointwise 95% confidence intervals. The figure stops at five years after the shooting; the overall ATTs also average event times six to eight, which rest on the earliest cohorts.