Local Disamenities, Youth Human Capital, and Optimal Zoning: Evidence from Cannabis Dispensaries*

Hantao Wu †

Job Market Paper

This version: March 2026

Please click here for the most recent version

Abstract: How should zoning policy protect children from local disamenities? I answer this question in the context of minimum-distance buffers between cannabis retailers and schools in Los Angeles, combining hyperlocal reduced-form evidence with a spatial entry model. Using administrative data on 600,000 LAUSD students linked to geocoded retailer openings, I estimate event-study designs with student and zipcode-by-semester fixed effects. Retailer proximity increases drug-related suspensions and reduces GPA, with effects concentrated within 0.5 miles for school exposure and 0.3 miles for neighborhood exposure; both effects sharpen after California's 2017 recreational legalization, when retail expanded well beyond the prior medical market. Long-run estimates show that school exposure reduces high school graduation by 1.2 percentage points and college enrollment by 0.9 percentage points; neighborhood exposure reduces these outcomes by 0.4 and 0.5 percentage points, respectively. Embedding these causal estimates in a spatial location-choice model, I evaluate policies by the net of human-capital gains and foregone tax revenue—trading off student welfare against fiscal costs to local government. The welfare-maximizing buffer is 0.5 miles, yielding $251.0M in human-capital gains and $112.7M in net welfare relative to the current 600-foot rule. A Pigouvian tax schedule calibrated to the estimated externality generates $210.5M in net welfare by preserving fiscal revenue from remaining retailers, at the cost of less aggressive student protection.

*I am very grateful to my advisors, Leah Boustan, Thomas Fujiwara, Ellora Derenoncourt, and Desmond Ang for their generous guidance and encouragement throughout this project. I also received valuable feedback from Alicia Adsera, Bocar Ba, Renjie Bao, Juan Camilo Castillo, Adam Kapor, Evan Soltas, and Yuci Zhou.

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

1 Introduction

Local governments routinely confront a zoning question with little empirical guidance: how close should commercial disamenities be allowed to operate near children? From cannabis retailers to bars, fast-food outlets, gambling venues, and methadone clinics, these establishments generate tax revenue and consumer surplus, but they also concentrate age-restricted products, foot traffic, and local externalities in children's immediate environment. A large literature shows that neighborhoods shape long-run human capital (Chetty, Hendren and Katz, 2016; Almond, Currie and Duque, 2018), and sector-specific evidence points to meaningful risks from nearby disamenities (Currie et al., 2010). Yet the dominant policy tool, the minimum-distance buffer, is typically set without the three inputs policymakers need: how harms decay with distance, how large they are for exposed students, and how firms re-sort when the buffer changes (Turner, Haughwout and van der Klaauw, 2014).

This paper provides an empirical framework that estimates all three objects and applies it to cannabis retail in Los Angeles. The setting is policy-relevant: recreational cannabis is now legal in 24 states with annual tax revenues exceeding $4.4 billion, and minimum-distance rules around schools remain among the most contested elements of local cannabis regulation.[1] The proximity of cannabis retailers to schools is a prominent source of public concern.[2] I combine reduced-form causal estimates of hyperlocal exposure effects with a spatial entry model that captures equilibrium relocation under alternative regulations. Empirically, I assemble a comprehensive student-semester panel by geocoding retailer addresses and school locations, linking LAUSD administrative student records to city licensing files, LAPD incident data, and high-frequency commercial data, and computing precise distances from each student's home and school to the nearest cannabis retailer in each semester.

The empirical strategy exploits hyperlocal variation within zipcode-by-semester cells, comparing students whose homes or schools are very close to a newly opened cannabis retailer against students slightly farther away but in the same zipcode and time period. This design isolates the causal effect of retailer proximity by holding fixed both time-invariant individual characteristics and neighborhood-level trends common to students in the same zipcode

[1] Buffer distances for cannabis retailers range from 500 feet (San Francisco) to 1,000 feet (Denver), with no jurisdiction citing empirical evidence for its choice.

[2] In April 2017, parents forced the temporary closure of a Pacoima dispensary located one block from schools serving over 3,000 students (San Fernando Sun, 2017; NBC Los Angeles, 2017).

and semester. I separate two distinct exposure channels, proximity to a student's school (the school channel) and proximity to a student's home (the neighborhood channel), and estimate distance-decay functions using 0.1-mile bins to determine data-driven treatment thresholds. The data reveal sharp distance gradients: effects on drug-related suspensions are concentrated within 0.5 miles for school exposure and 0.3 miles for neighborhood exposure.

The reduced-form results establish a clear chain from retailer proximity to youth human capital losses. For school exposure, drug-related suspensions increase by 0.008 (32 percent relative to the control mean), while for neighborhood exposure they increase by 0.004 (17.4 percent). These increases reflect both first-time marijuana offenses (rising by 0.002 for school exposure and 0.001 for neighborhood exposure) and non-first-time marijuana offenses (rising by 0.003 for school exposure and 0.002 for neighborhood exposure), consistent with increased cannabis access through both initial uptake and repeated use. Turning to the primary academic outcome, event-study estimates for term GPA show flat pre-trends followed by significant post-opening declines. The opening within 0.5 miles of a student's school reduces GPA by 0.028 points (0.032 standard deviations), while the opening within 0.3 miles of a student's home reduces GPA by 0.012 points (0.013 standard deviations). These effects are robust to school and year-by-cohort fixed effects, controls for local crime rates, census tract-by-semester fixed effects, alternative walking-path distance measures, and alternative treatment distance thresholds. The effects are larger after recreational legalization in 2017: for school exposure, the post-legalization GPA effect exceeds the pre-legalization effect, while for neighborhood exposure the effect becomes economically and statistically significant only after legalization.

Heterogeneity analyses show that the term GPA effects are larger for students from low-income families and students with lower baseline academic ability—groups that are also disproportionately likely to live near cannabis retailers. I also document an asymmetric interaction between exposure channels: school exposure affects students regardless of whether they also live near a retailer, while neighborhood exposure matters primarily when the school is not also exposed, indicating that school proximity is the dominant channel and crowds out the residential channel for doubly-exposed students.

The channel asymmetry reflects differences in the social environments through which each channel operates. The effects are amplified in school-cohorts with higher baseline suspension rates, consistent with peer environments playing a reinforcing role. To further explain

why the neighborhood channel is weaker when school exposure is absent, I show that neighborhood effects are concentrated among students who live in zipcodes where a larger share of their school-cohort peers also reside—consistent with the neighborhood channel operating through peer social networks rather than through an independent residential pathway. Together, this evidence suggests that peer interactions, rooted in the school setting, are the primary mechanism through which cannabis retailer proximity transmits to academic harm.

To translate the reduced-form evidence into policy guidance, I then embed these estimates in a parsimonious spatial entry model to evaluate two policy instruments: a minimum-distance buffer and a distance-varying Pigouvian tax. Since buffers and taxes alter where retailers locate in equilibrium, I model entry at the census block level following Seim (2006), estimating how retailers select locations based on local demand, operating costs, and competitive pressure from nearby incumbents. For each counterfactual policy, I simulate the new spatial distribution of retailers and map locations to changes in student exposure through the estimated distance-decay functions for graduation and college enrollment. Monetizing attainment effects using standard calibrations from the human-capital literature (Card, 1999; Oreopoulos and Salvanes, 2011; Chetty, Friedman and Rockoff, 2014), I compute net welfare as the difference between human-capital gains and foregone tax revenue. The welfare criterion assigns positive weight to student protection and fiscal revenue, and zero weight to cannabis consumers' surplus, due to distinction between involuntary pollutees and voluntary market participants. The optimal school buffer is 0.5 miles, yielding net welfare of $112.7 million relative to the current 600-foot rule; a distance-varying Pigouvian tax achieves smaller human-capital gains but generates fiscal revenue from externality-correcting add-on rates, yielding net welfare of $210.5 million. The two policies present a fundamental tradeoff: buffers maximize student protection by aggressively restricting entry near schools, while Pigouvian taxation maximizes net social welfare by balancing externality reduction against fiscal consequences.

This paper contributes to three literatures. First, it bridges the local externality and spatial entry literatures to provide a unified framework for evaluating place-based regulations of commercial disamenities. The local externality literature documents causal effects of proximity on health, housing values, and educational outcomes (Currie et al., 2015, 2010; Davis, 2011; Greenstone and Gallagher, 2008), but takes firm locations as given and therefore cannot answer how to design zoning policy; spatial entry models (Seim, 2006; Caoui, Hollenbeck and Osborne, 2022) characterize location choices under competition but abstract from the ex-

ternalities those choices impose on surrounding populations. I connect the two by estimating continuous distance-decay functions for youth human-capital harm—building on methods that exploit distance gradients around fixed points (Linden and Rockoff, 2008; Currie and Walker, 2011; Currie, Neidell and Schmieder, 2009)—and embedding them in a spatial entry model that predicts equilibrium relocation under counterfactual regulations. Grounded in Sandmo (1975)'s additivity result, a conceptual framework motivates a restricted welfare criterion and characterizes the student-protection–fiscal-revenue frontier traced by each instrument class, showing that price-based instruments achieve larger net welfare gains than categorical prohibitions at the cost of smaller human-capital gains—revealing a fundamental tradeoff between fiscal efficiency and student protection. Because the framework requires only a credibly estimated distance-decay of harm and a model of firm location choice, it applies directly to alcohol outlets (Livingston, 2008), fast-food restaurants (Currie et al., 2010), gambling venues, vaping shops (Henriksen et al., 2008), payday lenders (Melzer, 2011), and other disamenities for which policymakers must choose between quantity- and price-based instruments.

Second, I advance the neighborhoods and youth human capital literature by identifying effects at 0.1-mile resolution, the scale at which zoning rules actually bind. Prior evidence that places matter for children (Chetty, Hendren and Katz, 2016; Chetty and Hendren, 2018; Kling, Liebman and Katz, 2007) operates at tract, zipcode, or metro-area scales; even environmental-health studies linking pollution (Currie and Walker, 2011; Currie et al., 2014), lead (Aizer and Currie, 2019), and violence (Sharkey, 2010; Billings and Schnepel, 2018) to child outcomes typically identify effects at census-tract resolution or coarser. I contribute along three dimensions: estimating effects at 0.1-mile increments; separately identifying school and neighborhood exposure channels and showing they interact asymmetrically (school exposure affects outcomes regardless of neighborhood exposure, while neighborhood exposure primarily matters when school exposure is absent); and documenting that peer environments amplify individual exposure (effects are substantially larger in school-cohorts with higher baseline drug-suspension rates)—connecting the neighborhoods literature to the peer-effects literature in education (Sacerdote, 2001; Hoxby, 2000).

Third, I extend the cannabis-legalization literature by documenting that retail proximity imposes meaningful costs on school-age children at the hyperlocal scale where those costs operate. An earlier literature establishes that adolescent marijuana use carries lasting human capital consequences: marijuana use in high school reduces years of schooling completed

(Chatterji, 2006; Pacula, Ross and Ringel, 2003), lowers the probability of high school graduation (Yamada, Kendrix and Yamada, 1996), and substantially reduces subsequent employment (DeSimone, 2002). Adolescent marijuana use also reduces school attendance, a proximate channel through which academic performance declines (Roebuck, French and Dennis, 2004). Yet recent work on cannabis laws examines effects on youth substance use (Dills, Goffard and Miron, 2017; **?**; Powell, Pacula and Jacobson, 2018), crime (Sabia et al., 2021), adult health (Mathur and Ruhm, 2023; Jayawardhana, Hou and Maclean, 2025; Markowitz and Leinenbach, 2025), and labor supply (**??**), typically finding small or null aggregate effects at the state or county level. Two features limit its relevance for zoning policy: county- and state-level designs are too coarse to detect externalities that operate within blocks, and the near-exclusive focus on adults or aggregate youth surveys overlooks localized harm to minors who cannot legally purchase cannabis yet are exposed to retail presence simply because of where they live or attend school. I fill both gaps by estimating effects at hyperlocal resolution and tracing a causal chain—from first-time marijuana offenses through GPA declines to reduced graduation and college enrollment—at the individual student level, showing that retailer proximity causes substantial harm to nearby youth even where aggregate studies find limited effects and underscoring that evaluating legalization requires studying the right populations at the right geographic scale.

The remainder of the paper is organized as follows. Section 2 describes the institutional background of cannabis regulation in Los Angeles and the LAUSD. Section 3 presents the data sources. Section 4 details the empirical strategy and identification. Section 5 presents the reduced-form results. Section 6 develops the spatial entry model and counterfactual analysis. Section 7 concludes.

2 Background

This section describes the institutional setting: the Los Angeles cannabis market, the regulatory framework governing retail location, and the Los Angeles Unified School District.

2.1 The Los Angeles Cannabis Market

Legal cannabis markets have expanded rapidly since recreational legalization swept U.S. states in the mid-2010s. Nationwide sales rose from $13.6 billion in 2019 to $26.8 billion in 2022, generating nearly $3 billion in state tax revenue in 2022 alone. California is the largest segment of this market. The state legalized medical cannabis in 1996 under Proposition 215; Proposition 64 in November 2016 extended legalization to recreational use for adults aged 21 and older, with retail licenses available beginning January 2018. Crucially, Proposition 64 preserved local control: municipalities retained authority to permit, restrict, or ban cannabis businesses within their borders.

Los Angeles opted to permit recreational retail, establishing a licensing process that produced staggered entry across the city from 2018 onward. Over the 2010–2024 sample period, 1,371 licensed store entries occurred within the city. Some locations that had operated as medical dispensaries obtained recreational licenses and continued in place; others opened de novo. This staggered entry, combined with the city-approval requirement and location restrictions applicable to all licensed retailers, provides the event-study variation that underlies my empirical strategy. The large number of openings across an extended time period and a wide geographic spread ensures that many students experience a nearby retailer opening during their LAUSD enrollment, generating sufficient statistical power to detect effects at narrow distance bands.

Several features of the retailer data merit clarification. First, "licensed" in the pre-2018 medical-dispensary era refers to establishments registered with the City of Los Angeles Office of Finance for cannabis-specific tax compliance; after Proposition 64, it refers to retailers with both state approval from the Department of Cannabis Control and local authorization from the City's Department of Cannabis Regulation. Both periods draw from the same administrative source (city licensing records), ensuring comparability across the sample. Second, unlicensed

outlets operating outside the formal regulatory system are not observed in my data.[3] Third, each physical retail site has a single recorded opening date corresponding to its first entry into the licensed market; a site that transitioned from medical to recreational status at the same location retains its original opening date.

2.2 Zoning Buffer Regulation

Minimum-distance buffers are among the most widely used instruments in U.S. land-use regulation. Municipalities routinely prohibit certain businesses from locating within a specified radius of sensitive uses—most commonly schools, parks, daycare centers, and residential zones—with distances reflecting a range of underlying concerns: youth access to age-restricted substances, air-quality and safety hazards, and the perceived incompatibility of certain commercial uses with child-serving institutions.

A few examples illustrate the range. California's Alcoholic Beverage Control Act restricts off-sale alcohol retailers within 600 feet of schools. San Francisco's Tobacco-Free Youth Ordinance prohibits tobacco and vaping retailers within 1,000 feet. California's AB 1137 (2022) established a 3,200-foot setback for new oil and gas wells near schools. Across these cases, buffer distances are typically set through legislative bargaining or administrative convention rather than systematic cost-benefit analysis calibrated to local conditions.

Cannabis retail has been incorporated into this tradition as legalization expanded. Different jurisdictions adopted different thresholds: Detroit requires a 500-foot school buffer, Los Angeles adopted the 600-foot standard under California state law, and Denver mandates 1,000 feet. These cross-city variations reflect local politics and analogies to alcohol and tobacco regulation rather than evidence on how cannabis-related harms decay with distance from schools.

California operates a dual licensing system: retailers must obtain both a state license from the Department of Cannabis Control and local authorization—in Los Angeles, from the Department of Cannabis Regulation. The local process is the binding constraint; the state cannot issue a license without proof of local compliance. The Los Angeles Municipal Code imposes minimum-distance buffers from schools, daycare centers, parks, libraries, and other licensed cannabis businesses, collectively defining the feasible set of locations available to

[3]Unlicensed dispensaries were a small share of the Los Angeles market; coordinated city enforcement substantially reduced their prevalence (Associated Press, 2020; Office of the City Controller, City of Los Angeles, 2021).

prospective retailers and shaping the spatial distribution of licensed retail across the city.

2.3 The 600-Foot School Buffer

Under California state law and the Los Angeles Municipal Code, licensed cannabis retailers are prohibited from locating within 600 feet of K–12 schools, daycare centers, and youth centers. The distance is measured as a straight line between the nearest property boundaries of the cannabis establishment and the protected site—property line to property line rather than entrance to entrance—so the effective separation between a storefront and a school entrance can be substantially shorter than 600 feet when parcels are large or corner lots are involved.

The 600-foot threshold reflects regulatory convention rather than evidence-based calibration: like comparable buffers for alcohol and tobacco, it was not grounded in causal evidence on how proximity affects youth outcomes. The central question motivating this paper is whether the threshold is appropriately placed—a buffer that is too narrow exposes nearby students to harmful externalities, while one that is too wide unnecessarily restricts market entry, reduces consumer access, and forgoes tax revenue.

Available evidence suggests that even these statutory protections are inconsistently enforced in practice. Shi and Pacula (2021) conducted a statewide audit of 700 licensed recreational cannabis dispensaries in California and found pervasive noncompliance with youth-protection rules: only 12 percent of dispensaries checked identification outside the retail area, nearly 68 percent failed to post required age-limit signage, and 35 percent displayed products that may appeal to minors in areas accessible before ID verification. These findings, widely covered by CBS News Los Angeles (CBS News Los Angeles, 2021), suggest that the regulatory framework provides substantially less protection to youth than the statutory text implies.

Community concern about dispensary proximity to schools has also produced organized opposition. In April 2017, parents of students at the Vaughn charter school cluster in Pacoima protested a medical marijuana dispensary that opened within one block of five schools serving over 3,000 students, forcing a temporary closure and drawing local media attention to concerns about cannabis odor and the visibility of drug use near campuses (San Fernando Sun, 2017; NBC Los Angeles, 2017). In May 2022, LAUSD Superintendent Alberto M. Carvalho publicly raised concerns about dispensaries operating near school campuses, launching an initiative to work with local authorities to close or relocate retailers in close proximity to

schools (Beverly Press and Park La Brea News, 2022). Taken together, the combination of a conventionally calibrated buffer, weak enforcement, and sustained community mobilization underscores the potential for cannabis retailers to impose meaningful externalities on nearby youth—externalities that motivate the empirical analysis in this paper.

2.4 Los Angeles Unified School District

LAUSD is the second-largest school district in the United States, enrolling approximately 600,000 students across more than 1,000 schools spanning kindergarten through twelfth grade. This study focuses on high school students, whose age makes them the most plausible margin of effect for cannabis retail access. The student body is predominantly Hispanic (approximately 74 percent) and low-income (approximately 80 percent eligible for free or reduced-price lunch), with significant Black, White, and Asian minorities. Schools span a wide range of neighborhoods, from affluent westside communities to high-poverty areas in South and East Los Angeles.

The district's administrative data are well-suited to this study. Home addresses and school locations are geocoded, allowing me to compute precise distances between students and nearby retailers at both their place of residence and their enrolled school—the two exposure channels central to the analysis. The district records time-stamped disciplinary infractions with detailed codes that distinguish drug-specific offenses—including marijuana possession, other controlled substance violations, and drug sales—from other behavioral infractions, enabling examination of the "first stage" through which retailer proximity shapes youth drug behavior. Academic records include term-level GPA, attendance, and mental health service flags, enabling me to trace behavioral responses through to academic outcomes within the same student. Long-run data on high school graduation and college enrollment allow me to quantify the cumulative human capital cost of exposure.

3 Data

I construct a student-semester panel of LAUSD high school students that links precise cannabis retailer locations and opening dates to individual student outcomes. The analysis combines four principal data sources: administrative records on cannabis retailers, student-level records from LAUSD, geocoded 911 call data from the LAPD, and commercial and mobility data from SafeGraph and Advan. This section describes each source, the key variables, and the construction of the analysis sample.

3.1 Cannabis Retailer Records

I obtain administrative records on licensed cannabis retailers from the City of Los Angeles Office of Finance. These records cover the universe of licensed retail cannabis establishments operating in the city from 2009 through 2023, spanning both the medical dispensary era (pre-Proposition 64) and the post-legalization recreational market. Because licensing jurisdiction is limited to city boundaries—whereas LAUSD enrollment and LAPD records extend across a broader footprint of Los Angeles County—I restrict the analysis sample to students whose home addresses and enrolled schools both fall within city limits, ensuring complete retailer coverage for all observations.

I geocode each retailer's street address to latitude-longitude coordinates and construct a panel of operating status for each location. At the census block–semester level, I compute two variables: the count of new store entries (the number of stores that first open in that block-semester) and the stock of active stores (the number of stores currently operating). The count of new entries is used in the spatial entry model of Section 6, while the stock of active stores captures competitive pressure from incumbents. The opening date at each site indicates when licensed retail activity first began at that location.

The data cover all licensed outlets, which ensures that I observe the complete set of retailers to which any student might be exposed. This full-market coverage is important for accurately computing distances and for the spatial entry model developed in Section 6. Figure A1 presents annual snapshots of the spatial distribution of licensed cannabis retailers in Los Angeles for 2016 (prior to Proposition 64), 2019 (two years after recreational legalization), and 2023 (the final year of the sample). Because retailers enter and exit over time, each snapshot captures only those outlets operating at the end of the indicated year. Retailers cluster

along commercial corridors but exhibit meaningful within-zipcode spatial variation, a feature I exploit in the identification strategy.

3.2 LAUSD Student Administrative Records

The second data source is student-level administrative records from the Los Angeles Unified School District, covering all enrolled high school students over the 2009–2023 academic years.

Each student's home address is recorded annually and geocoded to latitude-longitude coordinates. I also observe the location of the student's enrolled school. These two geographic anchors allow me to compute distances to the nearest cannabis retailer from both school and residence, defining the two exposure channels, school exposure (school channel) and neighborhood exposure (home channel), that are central to the analysis.

The district records every disciplinary infraction with a date stamp and an offense code. I focus on drug-specific codes, which encompass first-offense possession of marijuana (not more than 1 ounce), non-first-time marijuana offenses (repeat marijuana possession and related marijuana violations), and other controlled substance categories such as counterfeit or prescription drug violations, sale or offer of controlled substances, and offering substitutes represented as controlled substances. I examine effects on total drug-related suspensions and then decompose the response into first-time and non-first-time marijuana offenses to assess whether the effects operate primarily through increased cannabis access.

Term-level GPA (on a 4.0 scale) serves as the primary measure of academic performance. I also observe daily attendance records, which I aggregate to semester-level attendance rates. Annual asthma diagnosis records are available for all students, and schools record mental health flags (indicators that a student has been referred for or received emotional disturbance services), which I use as a proxy for psychological well-being.

For students who reach the end of high school, I observe on-time graduation status and, through linked administrative records, first-fall college enrollment and enrollment at a four-year institution. I also observe leave and exit reasons, including transfers, dropouts, juvenile detention, and adult incarceration. The demographic variables include student age, gender, race and ethnicity (Hispanic, Black, White, Asian, other), parental education (college or above vs. below college), poverty status (eligibility for free or reduced-price lunch), and primary

language spoken at home.

3.3 LAPD 911 Call Records

I supplement the student data with geocoded 911 call records from the Los Angeles Police Department, covering 2009–2023. Each record includes the geographic coordinates of the incident, a timestamp, and a categorization of the call reason (violent crime, property crime, nuisance, drug-related offense, and others).

I aggregate 911 calls to distance rings around student homes and schools at the semester level, creating measures of hyperlocal crime exposure that I use to assess whether retailer openings coincide with changes in neighborhood safety. These crime measures also serve as controls in robustness specifications that assess whether the GPA effects of retailer proximity operate through increased crime or through other channels such as direct cannabis access, peer effects, or environmental degradation.

3.4 Commercial Environment and Mobility Data

To characterize the economic environment around retailer locations, and to estimate the spatial entry model, I draw on two proprietary data sources covering the period from 2019 to 2023. SafeGraph provides aggregated, anonymized credit and debit card transaction data at the point-of-interest level, recorded daily. For each point of interest, I observe total spend, median spend per customer, median spend per transaction, the number of customers and transactions, and online versus offline spending breakdowns. I aggregate these daily–point-level records to the census block–semester level, computing weighted means of per-customer and per-transaction spending along with total block-level expenditure within each semester. These variables proxy for revenue intensity and help characterize the commercial context of retailer sites. Because cannabis customers frequently transact in cash, card-based spending data alone may understate the true volume of retail activity. I therefore also draw on Advan foot traffic records, which provide visitor counts and visit counts at the census block level, aggregated to the same semester frequency. Foot traffic captures the total volume of customers visiting establishments in a block regardless of payment method. In the entry model, I include both spending and foot traffic measures jointly as demand shifters, since they capture distinct dimensions of a location's commercial appeal, spending intensity and physical traffic volume, and together provide a more complete picture of demand conditions than either source alone.

I augment these commercial data with block-level controls for population, income, rent, commercial land-use, and parking. Population and income controls are drawn from the ACS 5-year files (2019–2023) at the census block-group level and assigned to blocks within each block group: total population, per-capita income, and median gross rent. For zoning and parking, I construct countywide spatial proxies from OpenStreetMap features and map them to each block centroid: a commercial land-use indicator and local parking density (the count of parking amenities within 0.25 miles). I then map annual controls to the semester panel by calendar year and merge them into the block-semester entry dataset. In the entry model, these controls absorb differences in local market size, purchasing power, operating costs, commercial feasibility, and vehicle accessibility, so that the competition coefficients are identified net of these confounding neighborhood characteristics.

Figure A2 displays the distribution of distances from high school students to the nearest cannabis retailer, both overall and separately for the pre- and post-legalization periods. The distribution shifts leftward after Proposition 64, reflecting the expansion of the retail market.

4 Empirical Strategy

The goal of the empirical analysis is to estimate the causal effect of cannabis retailer proximity on youth outcomes. The key challenge is that retailers do not open at random: they locate in commercially zoned areas with specific demographic and economic characteristics that may independently affect student outcomes. This section describes the identification strategy, the estimating equation, the data-driven approach to choosing distance thresholds, and the validity checks that support a causal interpretation.

4.1 Identification Strategy

I exploit hyperlocal variation within zipcode-by-semester cells. The core comparison is between students who live (or attend school) very close to a newly opened cannabis retailer and students who are slightly farther away but reside in the same zipcode and are observed in the same semester. The identifying assumption is that, conditional on student fixed effects and zipcode-by-semester fixed effects, the exact distance between a student's home (or school) and the nearest retailer is uncorrelated with student-specific trends in outcomes.

This design addresses several threats to identification. Student fixed effects absorb all time-invariant individual characteristics, including family background, baseline ability, and residential sorting, that might correlate with both proximity to retailers and academic performance. Zipcode-by-semester fixed effects absorb any neighborhood-level shocks that affect all students in a zipcode in a given semester, including changes in local economic conditions, policing intensity, or school resources that might coincide with retailer entry. What remains is variation in how close, within a zipcode, a particular student happens to be to the retailer that opens, variation driven by the precise siting of the retailer along a commercial corridor relative to the student's block.

Figure 1 illustrates the identification concept. Because cannabis retailers must locate in commercially zoned parcels within a given zipcode, the precise site they select shares a similar micro-neighborhood environment with any other feasible site in that zipcode. Students whose homes (or schools) fall within the data-driven treatment threshold $\bar{d}$ of the retailer form the treated group, while students farther away but still within two miles in the same zipcode serve as controls. Both groups are exposed to identical zipcode-level trends, but only the treated group experiences the hyperlocal disamenity of the retailer opening. The treatment threshold

$\bar{d}$ is not imposed *a priori* but instead recovered from distance-decay regressions that reveal how drug-related suspensions attenuate with distance, as described in Section 4.2 below.

[Insert Figure 1 Around Here]

I separately estimate effects for two exposure channels. The first is school exposure (the school channel), where treatment is defined by the distance from the student's enrolled school to the nearest retailer. This channel captures effects that operate through the school environment: disruption of the school perimeter, peer influences mediated by students who can easily access the retailer before or after school, and visibility of drug-related activity near campus. The second is neighborhood exposure (the home channel), where treatment is defined by the distance from the student's geocoded home address to the nearest retailer. This channel captures effects that operate through the residential environment: exposure to cannabis odor and smoke, increased foot traffic and loitering near the home, changes in the perceived safety of the walking route to school, and direct access to cannabis products. Distances are computed using geodesic (straight-line) measures in the baseline specification; as a robustness check, I verify that results hold when using actual walking distances calculated via OpenRouteService.

4.2 Spatial Decay Specification

A critical design choice is the threshold distance $\bar{d}$ that separates treated from control students. Rather than imposing an ad hoc cutoff, I estimate a spatial decay function that lets the data reveal how the behavioral response to a nearby retailer attenuates with distance.

Specifically, I estimate:

$$Y_{ist} = \sum_{d} \beta_d \cdot \mathbf{1}\{\text{dist}_{is} \in [d, d+0.1)\} \times \text{Post}_{st} + \alpha_i + \lambda_{\text{zip}\times t} + \varepsilon_{ist}, \tag{1}$$

where Y_{ist} is drug-related suspensions for student i near store s in semester t, the distance bins are 0.1-mile increments from 0 to 2 miles, and the omitted reference category is [1.9, 2.0) miles. Student fixed effects α_i and zipcode-by-semester fixed effects $\lambda_{\text{zip}\times t}$ are included as before. Post_{st} is an indicator equal to one in semesters after the retailer opens.

I use drug-related suspensions, the most direct behavioral marker of increased cannabis access, as the outcome in this specification. The threshold $\bar{d}$ is then chosen as the distance beyond which the estimated coefficients $\hat{\beta}_d$ become statistically and economically indistin-

guishable from zero. The school and home channels are estimated separately, allowing each to have its own threshold. Section 5.1 presents the resulting distance-decay profiles and the implied thresholds for each channel.

4.3 Estimating Equation

Given the channel-specific thresholds $\bar{d}$ obtained from the spatial decay estimates, I estimate event-study specifications of the form:

$$Y_{ist} = \sum_{\ell \neq -1} \beta_\ell \cdot \mathbf{1}\{\tau_{st} = \ell\} \times \text{Close}_{is} + \alpha_i + \lambda_{\text{zip}\times t} + \gamma\, \text{Close}_{is} + \delta\, \text{Post}_{st} + \varepsilon_{ist}, \tag{2}$$

where $\text{Close}_{is} = \mathbf{1}\{\text{dist}(i,s) < \bar{d}\}$ indicates that student i is within the treatment threshold of store s, and τ_{st} denotes the semester relative to the opening of store s, with $\ell = -1$ (the semester before opening) normalized to zero. The specification includes student fixed effects α_i and zipcode-by-semester fixed effects $\lambda_{\text{zip}\times t}$, and standard errors are clustered at the zipcode level throughout.

The coefficients of interest are the β_ℓ for $\ell \geq 0$, which trace out the dynamic response of student outcomes to a nearby retailer opening. Pre-treatment coefficients ($\ell < -1$) provide a test of the parallel trends assumption: if the identifying variation is valid, treated and control students should exhibit similar outcome trajectories before the retailer opens.

I also report difference-in-differences estimates from a pooled specification that replaces the event-time indicators with a single $\text{Post}_{st} \times \text{Close}_{is}$ interaction, yielding an average treatment effect across all post-opening periods.

A potential concern with event-study designs involving staggered treatment is that two-way fixed effects estimators may produce biased estimates when treatment effects are heterogeneous across cohorts (Goodman-Bacon, 2021; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021). To address this, my baseline specification uses a stacked difference-in-differences design. For each cannabis retailer opening event, I construct a separate stack containing the treated students (those within the threshold distance) and a clean control group: students who are slightly farther away (between the threshold and 2 miles) but not yet treated by any other retailer within 2 miles in the relevant event window, all within the same zipcode. I then estimate the event-study specification on the stacked dataset, including stack fixed effects

to allow for different baseline levels across events. This design is implemented throughout as the baseline empirical approach.

The event-study specification indexes treatment relative to each student's *first* nearby retailer opening, the semester in which the nearest store within the relevant threshold distance first appears. In a dense urban market where students may eventually be exposed to multiple retailers, this timing convention means that later openings occurring while a student remains in the post-treatment window are not separately coded as new events. Instead, the estimated coefficients capture the cumulative effect of all retailer exposure from the time of the first opening forward. This approach parallels the logic in Kleven, Landais and Søgaard (2019), whose event study of the child penalty indexes relative time to the birth of a mother's first child; later births are subsumed in the post-treatment trajectory rather than treated as independent events. Here, the first opening represents the initial treatment, and subsequent nearby openings contribute to the total exposure captured in the post-period coefficients. Consequently, the estimated effects should be interpreted as the overall impact of cannabis retailer proximity, inclusive of the possibility that treated students are exposed to additional retailers over time, rather than the isolated effect of a single marginal store. This interpretation aligns with the policy simulations in Section 6, which evaluate changes in the entire spatial configuration of stores and rely on the reduced-form estimates to predict the cumulative human-capital consequences of proximity.

4.4 Robustness Specifications

I probe the robustness of the main results along several dimensions. First, I add school fixed effects to ensure that comparisons are made among students attending the same school, controlling for school-level resources, culture, and policies. Second, I include year-by-cohort fixed effects to absorb cohort-specific trends in outcomes that might confound the estimates. Third, I include stack (event) fixed effects to allow different baseline levels across opening events. Fourth, I control for the violence rate at the Census Block Group level to assess whether the GPA effects operate through changes in local crime or through other channels. Fifth, I replace zipcode-by-semester fixed effects with census tract-by-semester fixed effects, absorbing variation at a finer geographic scale and ensuring that results are not driven by comparisons across distant neighborhoods within the same zipcode. Sixth, I drop students who first enrolled in the District in grades 10–12. These late entrants may have transferred into the school specifi-

cally because of neighborhood conditions—including retailer proximity—introducing selection that could confound the estimates; restricting to students with longer exposure windows addresses this concern. Seventh, I re-weight the stacked estimator so that each event contributes equally, ensuring that results are not driven by a small number of large events.

4.5 Threats to Identification

Three principal threats to the causal interpretation merit discussion.

The first concern is selection on unobservables. If cannabis retailers systematically locate near students who are on differential academic trajectories, even conditional on student and zipcode-by-semester fixed effects, the estimates would be biased. I address this concern in two ways. Table 1 shows that treated and control students are similar on observable characteristics (age, gender, race, parental education, poverty status, English proficiency) after conditioning on zipcode fixed effects. The unconditional differences in Black share, parental college education, poverty status, and English proficiency that arise for home exposure vanish once zipcode fixed effects are included, consistent with between-zipcode sorting rather than within-zipcode selection. In addition, the event-study estimates display flat pre-trends, providing direct evidence against differential pre-existing trajectories.

[Insert Table 1 Around Here]

The second concern is that the estimated gradients may capture generic commercial intensification rather than cannabis-specific exposure. As noted in Section 2 and Figure A1, cannabis retailers cluster along commercial corridors. Within the same zipcode-by-semester cell, students living or attending school immediately along a busy artery (within 0.1–0.3 miles) may differ from those in the residential interior (beyond 0.3 miles) in ways that correlate with unobserved hyperlocal changes in the commercial environment. If retailer openings coincide with broader corridor-level increases in retail foot traffic, infrastructure investment, or other commercial activity, the estimated effects on student outcomes might reflect general neighborhood commercialization rather than cannabis-specific exposure.

To test this directly, I re-estimate the same hyperlocal stacked event-study design at the block-semester level using measures of the broader commercial environment from SafeGraph spending and Advan foot-traffic data. Specifically, I collapse total commercial POI counts, visit counts, and total spending (excluding cannabis establishments) into block-semester aggregates

and estimate event studies using both the 0.5-mile (school channel) and 0.3-mile (neighborhood channel) treatment thresholds. If cannabis openings were part of a broader commercial intensification process, I would expect to see corresponding increases in non-cannabis retail activity, foot traffic, or spending in the same hyperlocal area following the opening.

Appendix Figure A3 presents the results. Across all measures and both treatment thresholds, the event-study coefficients remain centered around zero with no sharp post-opening break or upward trend. Pre-treatment coefficients are flat and close to zero, and post-treatment coefficients show no systematic change relative to the pre-period. This null pattern indicates that cannabis openings in my setting are not accompanied by a parallel surge in broad local commercial activity. The evidence reinforces the interpretation that the main student-outcome effects operate through cannabis-specific mechanisms—direct access to cannabis products, visibility of cannabis-related activity, and peer-mediated substance use—rather than through general neighborhood commercialization or corridor-level economic shocks affecting all nearby retail establishments.

The third concern is selective migration. Since the LAUSD dataset records only students' most recently reported addresses rather than tracking residential moves over time, families might relocate away from neighborhoods where cannabis retailers open. If students living very close to a newly opened retailer are more likely to move away, the estimates could be biased by differential attrition in the treated group.

I address this concern through three complementary analyses. First, to more directly assess residential mobility responses, I analyze individual-level migration patterns using L2 voter file data spanning 2014–2022. These administrative records provide annual snapshots of registered voters' residential addresses, allowing me to construct a migration indicator based on whether individuals change their home address between consecutive years. I implement an event-study design analogous to the main spatial analysis, restricting to registered voters within 2 miles of a retailer and using both the 0.5-mile and 0.3-mile treatment thresholds. Appendix Figure A4 presents the results. Across both thresholds, the estimated event-study coefficients are centered around zero with no post-opening increase in residential mobility, indicating a null response of individual migration to cannabis openings. While voter files capture only the citizen population and therefore may not fully represent the student sample, this analysis provides direct micro-level evidence that cannabis retailer openings do not trigger

selective residential mobility at the hyperlocal scale relevant to the identification strategy.

Second, I examine aggregate mobility patterns using American Community Survey data from 2016–2020. For each census block group, I calculate the tenure rate—defined as the share of residents living in the same house as one year prior—and compare block groups with versus without nearby cannabis retailer exposure. Tenure rates are balanced across treated and control block groups, with a mean difference of 0.09 percentage points ($p = 0.90$), suggesting no differential residential stability around openings.

Third, Figure A5 presents event-study estimates for student transfers out of school and dropout rates under both school and neighborhood exposure definitions. In all cases, there is no sharp post-opening increase in transfer-out or dropout, indicating limited composition change around openings. Because migration responses may differ by socioeconomic status, Appendix Figures A6 and A7 replicate the same exercise separately for non-poor and poor students and show similar patterns.

A related concern is that housing affordability pressures may push economically distressed families toward cheaper, retailer-proximate neighborhoods—generating selection into treatment through downward residential mobility rather than avoidance. Our hyperlocal design absorbs this threat: zipcode-by-semester fixed effects difference out any common trends among families sorting into the same neighborhood, and the balance results in Table 1 show that poverty status and other demographics are similar between treated and control students within zipcode boundaries.

5 Results

This section presents the reduced-form results. I begin with the distance-decay estimates that motivate the treatment thresholds, then document the effect of retailer proximity on drug-related disciplinary infractions, the most direct behavioral response to increased cannabis access. I report the main results on GPA, examine heterogeneity by student demographics and peer environments, and present evidence on additional outcomes including health, safety perceptions, crime, and long-run attainment.

5.1 Distance-Decay Estimates

Figure 2 presents the distance-decay estimates from Equation (1) for drug-related suspensions under both exposure channels. Panel A reports school exposure: effects are positive and significant through 0.5 miles, with a gradual decline beyond that point, motivating the treatment threshold of $\bar{d}^{\text{school}} = 0.5$ miles. Panel B reports neighborhood exposure: effects are largest within 0.1–0.2 miles, decline monotonically through 0.3 miles, and become statistically indistinguishable from zero beyond 0.3 miles, motivating $\bar{d}^{\text{home}} = 0.3$ miles. The wider school-exposure radius is consistent with peer-mediated effects and school-perimeter disruptions that propagate over a larger area than direct residential exposure.

[Insert Figure 2 Around Here]

5.2 First Stage: Drug-Related Suspensions

If cannabis retailer proximity increases youth substance use, the most direct evidence should appear in individual students' drug-related disciplinary infractions. Figure 3 presents event-study estimates for total drug-related suspensions under both exposure channels. For school exposure (Panel A), pre-trend coefficients are flat and close to zero, and suspensions increase sharply in the semester of the opening, with effects persisting and growing through four semesters post-opening. The pooled ATE is 0.008, significant at the 10 percent level. Relative to the control mean of 0.025, this represents a 32 percent increase in drug-related suspensions. For neighborhood exposure (Panel B), the pattern is similar but smaller in absolute magnitude. The pooled ATE is 0.004, significant at the 1 percent level, corresponding to a 17.4 percent increase relative to the control mean of 0.023. The larger school-exposure effect is consistent with the school channel capturing stronger peer-mediated effects and direct disruption

of the school perimeter. A caveat is that suspension records may reflect both changes in actual cannabis use and changes in detection or enforcement: students who purchase cannabis near their homes may consume it outside school, and schools near newly opened retailers may intensify surveillance. The channel-magnitude comparison should therefore be interpreted with caution. The analysis of GPA in Section 5.3.1 provides a check on this concern: because term GPA is recorded mechanically from course grades, it is not susceptible to enforcement variation, and similar channel patterns in that outcome would support the interpretation that the suspension results reflect genuine changes in substance use rather than detection bias alone.

[Insert Figure 3 Around Here]

Table 2 summarizes first-stage magnitudes from pooled post-opening regressions. For school exposure, retailer openings increase drug-related suspensions by 0.008 (s.e. 0.004), first-time marijuana use by 0.002 (s.e. 0.001), and non-first marijuana use by 0.003 (s.e. 0.002). Relative to control means, these correspond to increases of 32.0 percent, 28.6 percent, and 25.0 percent. For neighborhood exposure, the corresponding estimates are 0.004, 0.001, and 0.002, implying increases of 17.4 percent, 12.5 percent, and 20.0 percent relative to means. The larger magnitudes under school exposure are consistent with the school channel being the stronger first-stage margin.

[Insert Table 2 Around Here]

To understand the nature of the drug-related response, I decompose total drug suspensions into first-time and non-first marijuana offenses. Figure A8 shows effects on first-time marijuana offenses: for school exposure, the pooled ATE is 0.002 (significant at the 10 percent level), representing a 28.6 percent increase relative to the control mean of 0.007; for neighborhood exposure, the ATE is 0.001 (significant at the 1 percent level), a 12.5 percent increase. Figure A9 shows effects on non-first marijuana offenses (repeat violations and related marijuana incidents): for school exposure, the pooled ATE is 0.003 (significant at the 10 percent level), a 25 percent increase relative to the control mean of 0.012; for neighborhood exposure, the ATE is 0.002 (significant at the 1 percent level), a 20 percent increase. That both first-time and repeat offenses increase indicates that retailer proximity affects both initial cannabis uptake and continued use. The significant effects on both margins suggest that marijuana-related infractions drive the overall increase in total drug suspensions. The smaller magnitude for neighborhood exposure could partly reflect that students who use cannabis near their resi-

dence are less likely to be detected at school, leading to underreporting of neighborhood-based use in school disciplinary records.

5.3 Main Outcome: Grade Point Average

Having established that cannabis retailer proximity increases drug-related suspensions, I now turn to the primary measure of human capital accumulation: term GPA. The first-stage results indicate increased substance use, but the question remains whether this translates into measurable academic harm.

5.3.1 Event-Study Estimates

Figure 4 presents event-study estimates for GPA under both exposure channels. For school exposure (Panel A), pre-trend coefficients are flat and centered on zero, ruling out differential pre-existing trajectories between treated and control students. Following the retailer opening, GPA drops sharply on impact and remains persistently negative in subsequent semesters. The pooled ATE is -0.028, significant at the 1 percent level. Relative to the control mean GPA of 2.504, this represents a 1.12 percent decline and 0.032 standard deviations. For neighborhood exposure (Panel B), pre-trends are also flat, with post-opening declines that are smaller but still persistent. The pooled ATE is -0.012, significant at the 1 percent level. Relative to the control mean GPA of 2.651, this represents a 0.45 percent decline and 0.013 standard deviations. Appendix Figure A10 shows that the event-study pattern remains similar when using actual walking distances rather than geodesic distances, confirming that measurement of proximity is not driving the results.

The larger school-exposure effect mirrors the channel ordering documented for drug suspensions. The school environment concentrates students in close proximity for extended periods, which may amplify any given disamenity relative to the more diffuse residential setting; the mechanisms underlying this differential are examined in the heterogeneity analysis below. Crucially, the school-channel dominance documented for drug suspensions reappears here for GPA—an outcome recorded mechanically from course grades and therefore insensitive to changes in enforcement intensity or school surveillance—providing reassurance that the suspension results reflect genuine changes in substance use rather than shifts in detection or disciplinary practice.

These ATEs reflect the average effect across all exposed students, combining direct ef-

fects on those who consume cannabis and spillover effects on their classmates and peers. I cannot directly estimate the effect on students who receive drug suspensions, as the sample of suspended students within narrow distance bins is too small to support separate within-suspension comparisons. Moreover, disciplinary records likely underreport actual cannabis use, so the true number of users exceeds observed suspensions. The estimated ATEs therefore capture the combined impact of direct use and peer spillovers, though the relative contribution of each component cannot be separately identified.

[Insert Figure 4 Around Here]

5.3.2 Sensitivity to Treatment Distance Definition

A natural concern is whether the main event-study results are driven by the particular treatment thresholds—0.5 miles for school exposure and 0.3 miles for neighborhood exposure—that are motivated by the distance-decay regression in Section 5.1. To assess this, I re-estimate the GPA event study using alternative treatment radii spanning a range above and below each baseline threshold.

[Insert Figure 5 Around Here]

For school exposure (Panel A), I consider thresholds of 0.3, 0.4, 0.5 (current), 0.6, and 0.7 miles. The underlying logic is straightforward: a threshold set below the effective harm radius of 0.5 miles reassigns students in the excluded outer band—those between the chosen cutoff and 0.5 miles—from treated to control. The resulting estimate may attenuate or remain similar depending on whether impacts in that marginal band are substantial or already close to zero. Conversely, a threshold set above 0.5 miles incorporates students beyond the effective impact radius into the treated group, diluting the estimate by averaging in observations with near-zero treatment effects. Consistent with this reasoning, expanding the radius to 0.6 or 0.7 miles weakens the post-opening decline relative to the baseline, while narrowing the radius yields changes in magnitude that reflect how quickly impacts decay across successive 0.1-mile bands rather than a determinate monotonic pattern.

For neighborhood exposure (Panel B), I consider thresholds of 0.1, 0.2, 0.3 (current), 0.4, and 0.5 miles. The same logic applies: widening the radius beyond 0.3 miles brings in students farther from the retailer, where effects are closer to zero, and therefore tends to attenuate the estimated post-opening decline through dilution. In contrast to the school-based results, tight-

ening the neighborhood radius to 0.1 or 0.2 miles yields smaller and less precisely estimated declines, consistent with discarding a nontrivial share of affected students in the 0.1–0.3 mile band.

Together, these patterns confirm that the distance-decay regression provides the appropriate treatment thresholds: the baseline results are neither artifacts of an overly broad treatment window nor sensitive to minor threshold adjustments. The finding that both narrower and wider definitions produce weaker estimates is precisely what one would expect if the true effect is concentrated at the distance ranges identified by the suspension distance-decay estimates.

5.3.3 Robustness

Table 3 presents the GPA results across eight specifications that progressively add controls and vary the estimation sample. Column (1) reports the baseline specification with student and zipcode-by-semester fixed effects. Column (2) adds school fixed effects. Column (3) adds year-by-cohort fixed effects. Column (4) adds stack (event) fixed effects for the stacked difference-in-differences design. Column (5) controls for the Census Block Group violence rate. Column (6) replaces zipcode-by-semester fixed effects with the more demanding census tract-by-semester fixed effects. Column (7) drops students who first entered the District in grades 10–12, focusing on those with longer exposure windows. Column (8) re-weights so that each stack contributes equally.

For school exposure (Panel A), the estimated effect on GPA ranges from -0.017 to -0.029 across all eight specifications, significant at the 1 percent level in every column. This corresponds to roughly 0.68 to 1.16 percent of the control mean GPA (2.504), or about 0.019 to 0.033 standard deviations. For neighborhood exposure (Panel B), the estimates range narrowly from -0.011 to -0.013, also significant at the 1 percent level throughout. This corresponds to about 0.41 to 0.49 percent of the control mean GPA (2.651), or about 0.012 to 0.014 standard deviations. The stability of the estimates across these demanding specifications, particularly the near-identical results when moving from zipcode-by-semester to tract-by-semester fixed effects, provides strong evidence that the effects reflect the causal impact of retailer proximity rather than confounding neighborhood trends.

[Insert Table 3 Around Here]

5.4 Heterogeneity

The average treatment effects documented above mask important heterogeneity in who bears the costs of cannabis retailer proximity and through what channels those costs are transmitted. Three patterns emerge from the heterogeneity analysis. First, the burdens fall disproportionately on already-disadvantaged students—those from low-income families and with lower baseline academic ability—widening existing achievement gaps. Second, the effects are substantially larger after recreational legalization, pointing to cannabis access as the operative mechanism rather than incidental neighborhood change. Third, the school-peer environment is the primary node through which proximity transmits to academic harm: peer amplification, asymmetric channel interaction, and peer mediation of the neighborhood effect all converge on the same conclusion. Together, these patterns do more than characterize who is harmed—they build a coherent picture of how retailer proximity operates.

5.4.1 By Student Demographics

Figure 6 presents heterogeneity in the GPA treatment effect by student demographics. For school exposure (Panel A), effects are broadly similar across racial groups, with no single racial group exhibiting markedly larger declines than others; the clearest gradient runs along socioeconomic lines, with students with below-college parental education (ATE $= -0.021$) and low-income students (ATE $= -0.020$) showing consistently negative and significant effects. For neighborhood exposure (Panel B), effects are statistically significant for Hispanic and Asian students, while estimates for Black and White students are small and not precisely estimated; the socioeconomic gradient again emerges clearly, with students with below-college parental education (-0.009), low-income students (-0.009), and students with low baseline academic ability (-0.019) all exhibiting significant declines. Taken together, the heterogeneity results point to a consistent income and ability gradient across both channels: cannabis retailer proximity disproportionately harms students from already-disadvantaged socioeconomic backgrounds, widening existing achievement gaps.

[Insert Figure 6 Around Here]

5.4.2 Pre- versus Post-Legalization Effects

Figure A11 separates the GPA event studies by regulatory regime: pre-2017 (medical dispensaries only) and post-2017 (recreational legalization). For school exposure, effects are neg-

ative in both periods but substantially larger after legalization. The pre-2017 ATE is -0.021 (significant at the 5 percent level), while the post-2017 ATE is -0.039 (significant at the 1 percent level), a substantially stronger effect after recreational legalization. For neighborhood exposure, the pre-2017 estimate is small and statistically insignificant (ATE $= -0.003$), while the post-2017 effect is clearly negative and highly significant (ATE $= -0.022$, significant at the 1 percent level).

This pattern points directly to cannabis access as the operative mechanism. Before 2017, purchasing required medical documentation—an effective if imperfectly enforced barrier to casual access. Recreational legalization removed that barrier entirely, lowering transaction costs and expanding the retail footprint well beyond the prior medical market. The result is a natural experiment in access intensity: the same hyperlocal identification strategy applied to the same neighborhoods yields substantially larger effects in the post-2017 period. That the neighborhood channel becomes economically meaningful only after legalization is particularly informative—it suggests that residential exposure to a cannabis retailer generates harm primarily when retail density is high enough that access is truly frictionless, consistent with the peer-network mechanism documented below.

5.4.3 By Peer Environment

Because drug-related suspension records are subject to underreporting—students who consume cannabis near their residences may do so off school premises, and detection depends on surveillance intensity—the GPA estimates capture the combined impact of direct cannabis use and peer externalities without cleanly separating the two margins. To isolate peer spillovers as a distinct mechanism, I ask whether the GPA effect on any given student is amplified by the drug-related behavior of classmates, exploiting variation in the baseline drug environment of each student's school cohort.

To investigate whether peer effects amplify individual exposure, I split the sample by baseline drug-related suspension rates within each school-cohort cell. Figure A12 presents the results for both school and neighborhood exposure channels. Students in school-cohorts with above-median baseline drug suspension rates exhibit substantially larger GPA declines following retailer openings than those in low-suspension cohorts. This is consistent with a broader peer and school-environment reinforcement mechanism: when a retailer opens nearby, the harm to any given student is amplified in school cohorts already characterized by higher

drug-related activity—whether through direct peer influence on individual behavior, shared behavioral norms, or the general disciplinary and academic climate of the school (Sacerdote, 2001; Hoxby, 2000).

5.4.4 Asymmetric Interaction between School and Neighborhood Exposure

I test whether the two exposure channels interact by splitting each exposure sample by the other channel's exposure status. Rather than using the full treatment threshold, I define "also exposed" through the complementary channel using the bottom decile of the distance-to-nearest-retailer distribution in each sample, isolating students with the most intense dual exposure.[4] Figure A13 presents the results for both channel directions. For school exposure (Panels A and B), the GPA effect is nearly identical regardless of whether a retailer also operates near the student's home, indicating that the school channel operates independently of the neighborhood channel: school proximity harms students whether or not they also live near a retailer. For neighborhood exposure (Panels C and D), the pattern differs: the effect is concentrated among students whose schools are not near a retailer (Panel C), while the estimate becomes small and imprecise when the student's school is also exposed (Panel D). This asymmetry suggests that school exposure is the dominant channel and partially substitutes for neighborhood exposure, whereas neighborhood exposure adds harm primarily when school exposure is absent. Because GPA is mechanically recorded from course grades and less susceptible to enforcement variation than disciplinary records, the school-channel dominance documented here provides further evidence that the effects reflect actual changes in substance use and academic engagement rather than differential detection or surveillance intensity across the two channels.

5.4.5 Peer Mediation of the Neighborhood Channel

The asymmetric channel interaction documented above raises a natural question: if the neighborhood effect operates through social rather than purely residential mechanisms, it should be stronger precisely where school-cohort peers are concentrated in the same residential area. I probe this mechanism more directly by asking whether the neighborhood GPA effect is mediated by the density of school peers in the student's residential environment. If

[4] Specifically, the school-exposure sample splits by whether the student's home is within 0.17 miles of a retailer (bottom 10th percentile of home-to-nearest-retailer distances in the school sample), and the neighborhood-exposure sample splits by whether the student's school is within 0.44 miles of a retailer (bottom 10th percentile of school-to-nearest-retailer distances in the neighborhood sample).

residential proximity generates harm in part by amplifying peer-network interactions, the neighborhood effect should be stronger for students who live in neighborhoods where more of their school-cohort peers also reside. I test this by splitting the neighborhood-exposure sample by the median number of school-cohort peers living in the same zipcode: students in the upper half share their residential neighborhood with more school peers, while those in the lower half live in neighborhoods with fewer school peers.

Figure A14 presents the results. Panel A reports event-study estimates for students in low-peer neighborhoods: pre-trends are flat, and post-opening GPA effects are small and statistically indistinguishable from zero. Panel B reports estimates for students in high-peer neighborhoods: the pre-trend remains flat, but the post-opening GPA decline is substantially larger and statistically significant. The contrast indicates that the neighborhood effect is concentrated precisely among students who live in neighborhoods densely populated with school-cohort peers. Residential exposure alone, in neighborhoods with few school peers, generates little measurable academic harm. When a student lives near a retailer in a neighborhood also populated by many school peers, however, the two sources of access reinforce each other through social interaction, peer influence, and shared behavioral norms. These results offer a coherent explanation for the school-channel dominance documented throughout: the school setting is the primary node through which cannabis-retailer externalities propagate, and the smaller neighborhood effect reflects peer spillovers grounded in school-cohort social networks rather than an independent residential pathway.

[Insert Figure A14 Around Here]

5.5 Absence Rate

A natural question is what behavioral mechanism links retailer proximity to GPA decline. Student absence provides a direct answer. An established literature documents that adolescent marijuana use reduces school attendance (Roebuck, French and Dennis, 2004; Chatterji, 2006; Pacula, Ross and Ringel, 2003), and absenteeism is itself one of the strongest proximate predictors of academic underperformance. If retailer proximity increases cannabis access and use, it should raise absence rates—providing a behavioral bridge between the first-stage suspension results and the GPA effects documented above. Figure **??** presents event-study evidence on semester-level absence rates.

[Insert Figure 7 Around Here]

Panel A reports the school-exposure estimates: the pooled ATE is 0.006, statistically significant at the 1% level, against a baseline absence rate of 0.091. This implies that a retailer opening near a student's school raises the absence rate by roughly 6.6%, a magnitude consistent with the GPA effects documented above. Panel B reports the neighborhood-exposure estimates: the post-opening average effect is 0.002, significant at the 5% level, against a baseline mean of 0.100, corresponding to a 2.0% increase. While the post-treatment coefficients in both panels show a clear upward shift relative to the pre-period, the neighborhood estimates warrant some caution because the coefficient at relative semester -5 is elevated above zero, suggesting a possible deviation from parallel trends at the earliest pre-treatment horizon; the remaining pre-treatment coefficients, however, are close to zero and trace out a flat pattern from semesters -4 through -1. Taken together, the absence results reinforce the interpretation that cannabis retailer proximity reduces student engagement, with school exposure producing the larger and more precisely estimated effect.

5.6 Additional Outcomes

5.6.1 Health and Mental Health

I examine health-related outcomes using annual administrative records. Figure A15 and Figure A16 report asthma diagnoses and emotional disturbance flags for school exposure (Panel A) and neighborhood exposure (Panel B). I find no significant impacts on asthma diagnoses for either channel, consistent with recent evidence that cannabis legalization has not worsened respiratory health overall (Jayawardhana, Hou and Maclean, 2025). For mental health flags, I find no significant effects for either channel, though there is suggestive evidence of increased emergency department visits for neighborhood exposure. Recent work on cannabis access and mental health finds mixed results (Markowitz and Leinenbach, 2025), and my null findings on mental health flags are consistent with the possibility that adverse effects may be concentrated in specific subpopulations or manifest in ways not captured by school administrative records.

5.6.2 Safety Perceptions and Local Crime

Using school survey data, I find no significant impacts of retailer proximity on perceived school safety (Figure A17) or perceived neighborhood safety (Figure A17). To directly test

whether retailer openings increase local crime, I use LAPD geocoded 911 call data aggregated to the block-semester level and estimate the same hyperlocal event-study design. Figure A18 presents results for two treatment thresholds: Panel A uses 0.5 miles (consistent with school exposure), and Panel B uses 0.3 miles (consistent with neighborhood exposure). Both panels show null effects on total crime calls, echoing recent findings that recreational marijuana legalization does not increase crime (Sabia et al., 2021). This null result rules out general neighborhood deterioration and increased violent crime as operative mechanisms. Combined with the first-stage drug suspension results, it points to direct cannabis access as the primary pathway through which retailer proximity reduces student academic performance.

5.7 Long-Run Outcomes

Understanding how semester-level GPA effects cumulate into long-run human capital losses is essential for quantifying the welfare cost of retailer proximity at each distance increment. Figure 8 reports distance-decay profiles for three long-run outcomes—cumulative GPA, high school graduation, and college enrollment—separately by exposure channel. The long-run specification is cohort-based. For each distance bin b, channel $c \in \{s, h\}$, and outcome $o \in \{\text{CumGPA}, \text{Grad}, \text{College}\}$, I estimate:

$$O_i^o = \sum_b \beta_b^{c,o} \cdot \mathbf{1}\{\text{dist}_i^c \in b\} \times \text{TreatedCohort}_i + \mu + \varepsilon_i, \tag{3}$$

where $\text{TreatedCohort}_i = 1$ for students who had not yet graduated when the nearest retailer opened, and $= 0$ for cohorts that had already graduated. School-channel regressions ($c = s$) include campus–zipcode fixed effects with race and gender controls; neighborhood-channel regressions ($c = h$) include student–zipcode fixed effects. The coefficients $\{\beta_b^{c,o}\}$ trace the distance-decay profile for each long-run outcome and serve as the key inputs to the welfare calculations in Section 6.

[Insert Figure 8 Around Here]

Across all three outcomes, effects are concentrated in the nearest distance bins and attenuate monotonically, mirroring the distance-decay structure documented for term GPA and drug suspensions. The pooled ATEs displayed in each panel summarize effects across all exposed students. School exposure reduces cumulative GPA by 0.023 points (0.032 SD), high school graduation by 1.2 percentage points (2.8 percent relative to the control mean of 0.431),

and college enrollment by 0.9 percentage points (3.0 percent relative to the control mean of 0.301); all three estimates are significant at the 1 percent level. Neighborhood exposure reduces cumulative GPA by 0.023 points (0.018 SD), high school graduation by 0.4 percentage points (0.64 percent relative to the control mean of 0.624), and college enrollment by 0.5 percentage points (1.2 percent relative to the control mean of 0.422); the GPA and college enrollment effects are significant at the 1 percent level, and the graduation effect at the 5 percent level. Although the two channels yield nearly identical pooled ATEs for cumulative GPA, the school channel produces substantially larger effects on graduation and college enrollment, consistent with its broader spatial reach and peer-mediated mechanisms documented above. These distance-decay profiles are the key inputs to the next section, where I embed them in a spatial entry model to evaluate the net welfare consequences of alternative buffer regulations and Pigouvian taxation.

6 Spatial Entry Model and Policy Counterfactuals

Reduced-form estimates establish that retailer proximity harms youth human capital, but translating them into policy guidance requires confronting a second-best problem. A first-best policy—directly assigning retailer locations to maximize welfare—is not available to regulators; instead, they must choose among implementable instruments that induce equilibrium responses on the supply side. The relevant comparison is therefore among feasible tools evaluated against the status quo 600-foot rule: a minimum-distance buffer that categorically excludes entry near schools, and a distance-varying Pigouvian tax that prices proximity according to its external harm.

Because both instruments shift the equilibrium spatial distribution of retailers rather than simply reducing store counts, a reduced-form calculation that holds firm locations fixed would misstate their welfare consequences. I therefore develop a parsimonious spatial entry model to quantify equilibrium reallocation under each instrument, map simulated retailer locations to student exposure changes through the distance-decay functions estimated in Section 5, and monetize attainment effects to evaluate each policy under the restricted planner's criterion.[5]

6.1 Conceptual Framework

I formalize the narrower policy problem that zoning authorities face in practice. The framework captures a regulator who weighs student harm against fiscal consequences when choosing between two instruments for a retail disamenity near schools: a minimum-distance buffer and a distance-varying corrective tax.

Benchmark. Sandmo (1975) shows that the optimal tax on an externality-generating good decomposes additively into a revenue-raising Ramsey component and an externality-correcting Pigouvian component. Because even the optimal price instrument must balance externality correction against revenue considerations, the planner's objective assigns weight to both $H(P)$ and $R(P)$. Within this criterion, I compare two instruments that address the same un-

[5] I use a parsimonious spatial entry model because the policy variation acts on store siting: the key counterfactual objects are retailer location, student exposure, and tax revenue under alternative zoning rules—the margins at the center of current Los Angeles school-zoning design. A full demand system (Berry, Levinsohn and Pakes, 1995) is suited to product-level pricing and requires census-block-level price and quantity data unavailable here; the empirical design is organized around the location margin. Magnolfi et al. (2024) adopt a comparable approach to study how urgent care center location shapes market structure and patient access, recovering equilibrium entry patterns without modeling the full demand system.

derlying externality: a price instrument—a distance-varying corrective tax that allows retailers near schools subject to a payment calibrated to their harm—and a quantity instrument—a minimum-distance buffer that categorically prohibits operation within a prescribed radius. The two differ not in objective but in how they trade off exposure reduction against preservation of the tax base.

Environment. Let locations be indexed by $j \in \mathcal{J}$, each at distance $s_j \geq 0$ from the nearest school. Under policy P, let $n_j(P) \in \{0, 1\}$ indicate whether a retailer operates at j, and let $N(P) = \{n_j(P)\}_{j \in \mathcal{J}}$ denote the equilibrium retail configuration. The regulator evaluates two objects. First, aggregate student harm:

$$H(P) \equiv \sum_{i \in \mathcal{I}} h_i\big(N(P)\big), \tag{4}$$

where $h_i(\cdot)$ is the present-discounted human-capital loss imposed on student i by the equilibrium pattern of retailer locations—the object identified in Section 5, with effects concentrated at short distances from schools. Second, aggregate fiscal revenue:

$$R(P) \equiv \sum_{j \in \mathcal{J}} \rho_j\big(n_j(P), P\big), \tag{5}$$

where $\rho_j(\cdot)$ collects tax payments from retailer j under policy P, including any distance-varying add-on under a corrective tax.

Restricted planner objective. Sandmo (1975) draws a distinction between polluters—those whose consumption generates the externality—and pollutees—those who bear its costs. In this setting the two groups are categorically distinct. Cannabis retailers and their customers generate the proximity externality through retail demand; students bear it without participating in the market. As legal minors, students cannot contract for cannabis, and they exercise little control over their school or residential location—they are involuntary pollutees in Sandmo's sense. Cannabis consumers, by contrast, are voluntary market participants who can substitute across retailers or forgo purchases altogether. The planner's objective therefore assigns positive weight to the welfare of pollutees—measured by student harm and its fiscal complement—

and zero weight to the surplus of voluntary participants. The planner evaluates policy P by

$$W(P) \;=\; -\,H(P) \;+\; \lambda\, R(P), \tag{6}$$

where $\lambda > 0$ converts revenue into the same units as human-capital harm. In the monetized counterfactuals, this reduces to the sum of human-capital gains and fiscal revenue preserved relative to the status quo.

Buffer regulation. Under a minimum-distance buffer $d \geq 0$, a retailer may operate at j only if $s_j \geq d$. Let $N_{\text{reg}}(d)$ denote the equilibrium configuration. The planner's payoff is

$$V_{\text{reg}}(d) \;=\; -\,H\big(N_{\text{reg}}(d)\big) \;+\; \lambda\, R\big(N_{\text{reg}}(d)\big). \tag{7}$$

A tighter buffer may remove some retailers entirely and displace others to blocks farther from schools, so welfare depends on the full spatial reallocation, not just the number of stores excluded.

Pigouvian tax. Under a distance-dependent tax schedule $\tau(s)$ with $\tau'(s) \leq 0$, retailers closer to schools pay a higher add-on. Let $N_{\text{tax}}(\tau)$ denote the induced configuration. The planner's payoff is

$$V_{\text{tax}}(\tau) \;=\; -\,H\big(N_{\text{tax}}(\tau)\big) \;+\; \lambda\, R\big(N_{\text{tax}}(\tau), \tau\big). \tag{8}$$

The tax preserves locational flexibility by pricing proximity rather than banning it. It need not reduce harm as aggressively as a strict buffer, but it can preserve or increase fiscal revenue through corrective payments from retailers that remain near schools.

Instrument comparison. Each instrument class traces a frontier in (S, R)-space, where $S(P) \equiv -H(P)$ denotes student protection and $R(P)$ denotes fiscal revenue. The planner selects the point maximizing $W = S + \lambda R$. A stricter buffer moves the economy toward higher S but lower R, because it protects students by suppressing entry and shrinking the tax base. A steeper corrective tax raises S more gradually while maintaining higher R, because it internalizes part of the externality without categorically prohibiting operation. Neither instrument is unambiguously superior: the ranking depends on how the planner values student

protection relative to revenue. In the absence of any instrument, retailer location choices are privately optimal but socially excessive in the sense of Mankiw and Whinston (1986): retailers do not internalize the human-capital cost they impose on nearby students.

To evaluate these tradeoffs quantitatively, I need to predict how retailers relocate under each instrument and map their locations to changes in student outcomes. I turn to a spatial entry model that fills in the objects $H(\cdot)$ and $R(\cdot)$.

6.2 Spatial Entry Model

I model the cannabis retail market as a spatial entry game at the census block level, following Seim (2006). The city of Los Angeles contains 2,341 census tracts and 39,637 census blocks, yielding an average choice set of approximately 17 blocks per tract. The city is partitioned into census blocks indexed by j, each belonging to a census tract $\tau(j)$. In each tract-semester, the number of entrants is taken as given from the data—determined by demand conditions and city licensing—while the model governs where within the tract each entrant locates; the outside option of not entering the tract at all is absorbed by the tract-semester fixed effects that scale total entry. Conditional on entering tract τ, the retailer selects the profit-maximizing block $j \in \mathcal{J}_\tau$.

This within-tract choice structure reflects both the economics of cannabis entry and the logic of the reduced-form identification. Cannabis retailers typically select a broad service area first and then choose a specific block based on lease availability, lot dimensions, and idiosyncratic commercial conditions. Conditioning on tract$\times$semester choice sets holds fixed everything that determines which tract a retailer enters, focusing identification on within-tract block selection—analogously to how the reduced-form strategy uses zipcode$\times$semester fixed effects to hold neighborhood-level shocks fixed. By the same logic, within-tract block selection is conditionally exogenous: blocks sharing a tract and semester face similar zoning, demographic composition, and unobserved neighborhood quality, so conditional on tract$\times$semester fixed effects, which specific block receives a retailer is plausibly driven by idiosyncratic factors uncorrelated with block-level student outcomes.

Under buffer regulation $d^{\min}$, the feasible block set is $\mathcal{J}_\tau(d^{\min}) = \{j \in \mathcal{J}_\tau : \text{dist}_j^{\text{school}} >$

$d^{\min}\}$, and the retailer chooses $j^* \in \mathcal{J}_\tau(d^{\min})$ to maximize:

$$\Pi_{jt} = X_{jt}^d \beta_d - X_{jt}^c \beta_c + \sum_b \delta_b N_{jt}^b + \lambda_{\tau(j),t} + \xi_{jt}, \tag{9}$$

where X_{jt}^d is a vector of demand shifters, X_{jt}^c cost shifters, N_{jt}^b the number of incumbents in distance band b, $\lambda_{\tau(j),t}$ tract×semester fixed effects, and ξ_{jt} a block-semester profitability shock. Counterfactual simulations vary $d^{\min}$, changing which blocks are feasible while profit parameters remain unchanged."

Following Seim (2006), competitive effects vary with distance. I define two bands: N_{jt}^1 counts incumbents within 0.25 miles, and N_{jt}^2 counts those between 0.25 and 0.5 miles. The parameters δ_1 and δ_2 capture how competitive pressure attenuates with distance; δ_2 may differ in sign from δ_1 if moderately distant retailers generate agglomeration spillovers.

The demand vector X_{jt}^d includes commercial activity measures from the broader retail sector (SafeGraph credit-card spending, customer counts, and median spend per customer; Advan visitor counts and total visits), together with analogous measures across all POI categories. Because cannabis customers frequently transact in cash, foot-traffic measures are particularly valuable. I also include ACS demographic controls (population and per-capita income at the block-group level). The cost vector X_{jt}^c includes median gross rent, local parking density, and a commercial land-use indicator. To address the standard concern that unobserved location quality biases competition parameters (Seim, 2006), I follow Caoui, Hollenbeck and Osborne (2022) and include total POI points in the block-semester as a proxy for unobserved commercial attractiveness.

6.3 Estimation

I estimate the location-choice parameters using a conditional-logit framework with tract-semester choice sets. Let $g = (\tau, t)$ index a tract-semester cell, and let N_{jg} be the observed number of entrants in block j in that cell. The model-implied conditional choice probability is

$$P(j \mid g) = \frac{\exp\left(X_{jg}^d \beta_d - X_{jg}^c \beta_c + \sum_b \delta_b N_{jg}^b\right)}{\displaystyle\sum_{k \in \mathcal{J}_g} \exp\left(X_{kg}^d \beta_d - X_{kg}^c \beta_c + \sum_b \delta_b N_{kg}^b\right)}. \tag{10}$$

Given observed totals $N_g = \sum_j N_{jg}$, estimation maximizes the grouped conditional-logit likelihood $\sum_g \sum_{j \in \mathcal{J}_g} N_{jg} \log P(j \mid g)$, identifying $(\beta_d, \beta_c, \delta)$ from within-choice-set variation in block characteristics and competition.

Table A1 reports summary statistics for the estimation sample. Table 4 presents five conditional-logit specifications, progressively adding broad-retail demand shifters (1), total POI demand (2), distance-banded competition (3), the log POI proxy for unobserved location quality (4), and lagged store stock, school distance, and demographic and cost controls (5).

6.4 Estimation Results

Table 4 reports the conditional-logit estimates. In column (5), log broad retail spending enters positively and significantly, while log POI points and log lagged store stock are both positive and highly significant, indicating that retailers prefer commercially vibrant blocks and locations with demonstrated prior entry. The competition coefficients are small: $N_1 = 0.0151$ (s.e. 0.0448) and $N_2 = 0.0750$ (s.e. 0.0398, $p < 0.10$). A Wald test cannot reject equality of the two coefficients ($p = 0.15$), and a one-sided test for distance-based attenuation is not significant (Table A2). The modest competition effects suggest that spatial crowding is not a first-order determinant of entry, consistent with retailers differentiating on product variety, service quality, or local reputation rather than competing primarily on proximity.

[Insert Table 4 Around Here]

Figure A19 plots predicted versus observed entry at the block-semester level. The model captures a meaningful share of within-market variation; the level correlation is 0.610.

6.5 Counterfactual Design

I evaluate two policy counterfactuals relative to Los Angeles' current 600-foot (≈0.114 miles) school buffer: a uniform minimum-distance buffer on a grid from 0.0 to 1.0 miles in 0.1-mile increments, and a distance-varying Pigouvian tax with add-on rates (in 0.1-mile bins) calibrated to equate marginal private cost with marginal social cost.

Spatial reallocation. For each $d^{\min}$, the feasible block set for tract τ in semester t is $\mathcal{J}_{\tau t}(d^{\min}) = \{j \in \mathcal{J}_{\tau t} : \text{dist}_j^{\text{school}} > d^{\min}\}$. I use an iterative pseudo-equilibrium to capture

endogenous competitive crowding. At iteration k, competition counts update as

$$N_{jt}^{1,(k)} = \sum_{l \in \mathcal{J}_{\tau t}(d^{\min})} \mathbf{1}\{\text{dist}_{jl} \leq 0.25\} \hat{N}_{lt}^{(k-1)}, \tag{11}$$
$$N_{jt}^{2,(k)} = \sum_{l \in \mathcal{J}_{\tau t}(d^{\min})} \mathbf{1}\{0.25 < \text{dist}_{jl} \leq 0.5\} \hat{N}_{lt}^{(k-1)},$$

and block probabilities follow

$$P^{(k)}(j \mid \tau, t, d^{\min}) = \frac{\exp\left(\hat{V}_{jt}^{(k)}\right) \mathbf{1}[j \in \mathcal{J}_{\tau t}(d^{\min})]}{\sum_{\ell \in \mathcal{J}_{\tau t}(d^{\min})} \exp\left(\hat{V}_{\ell t}^{(k)}\right)}, \quad \hat{V}_{jt}^{(k)} = X_{jt}\hat{\beta} + \hat{\delta}_1 N_{jt}^{1,(k)} + \hat{\delta}_2 N_{jt}^{2,(k)}. \tag{12}$$

Tract-semester market size adjusts via an inclusive-value ratio:

$$\hat{M}_{\tau t}^{cf,(k)} = M_{\tau t}^{obs} \times \left(\frac{IV_{\tau t}^{(k)}(d^{\min})}{IV_{\tau t}(d^{\text{sq}})} \right)^{\eta}, \quad IV_{\tau t}^{(k)}(d) = \exp\left(\ln \sum_{\ell \in \mathcal{J}_{\tau t}(d)} \exp\left(\hat{V}_{\ell t}^{(k)}\right) \right), \tag{13}$$

and block-level predicted entry is $\hat{N}_{jt}^{(k)} = \hat{M}_{\tau t}^{cf,(k)} \times P^{(k)}(j \mid \tau, t, d^{\min})$. I iterate until convergence (sup-norm below 10^{-5}). The parameter $\eta \in (0, 1)$ governs the extensive-margin response; I set $\eta = 0.60$ following Eaton, Kortum and Kramarz (2004). Appendix Table A3 reports robustness over $\eta \in \{0.55, 0.60, 0.65\}$.

Mapping to student exposure. For each student i, let d_i^s and d_i^h denote the closest simulated retailer distance from school and home. I apply piecewise-constant dose-response functions $g^{c,o}(d) = \hat{\beta}_b^{c,o}$ for $d \in b$, where $\hat{\beta}_b^{c,o}$ is the estimated distance-bin coefficient for channel $c \in \{h, s\}$ and outcome $o \in \{\text{grad}, \text{col}\}$, and the student-level attainment change is

$$\Delta O_i(d^{\min}) = \sum_{c \in \{h,s\}} \left[g^{c,o}\left(d_i^c(d^{\min})\right) - g^{c,o}(d_i^c(d^{\text{sq}})) \right]. \tag{14}$$

The two channels are not additive. Section 5.4.4 shows that the home-channel effect is null when $d_i^s < 0.44$ miles: school proximity crowds out the residential channel. I therefore apply the home-channel dose response only when $d_i^s \geq 0.44$ miles, preventing double-counting for doubly-exposed students.

Welfare. The human-capital loss at buffer d is

$$L^{HC}(d) = -\sum_i \left[G_i(d) \times \rho \bar{E} + C_i(d) \times 2\rho \bar{E} \right], \tag{15}$$

where G_i and C_i are graduation and college effects, $\rho = 0.125$ is the return to schooling (Angrist and Krueger, 1991; Card, 1999; Oreopoulos and Salvanes, 2011), $\bar{E}$ is baseline present-value lifetime earnings, and the 3 percent discount rate follows Office of Management and Budget (2003). The human-capital gain from tightening relative to status quo is $\Delta W^{HC}(d^{\min}) = L^{HC}(d^{\text{sq}}) - L^{HC}(d^{\min})$.

The fiscal cost is $\Delta W^{\text{Rev}}(d^{\min}) = -\Delta M(d^{\min}) \times \overline{\text{Rev}}_{\text{sem}}$, where $\overline{\text{Rev}}_{\text{sem}}$ is average tax revenue per retailer-semester. Because cannabis transactions are predominantly cash (Cannabis Business Times, 2022), I calibrate a credit-to-total-sales multiplier by comparing SafeGraph credit-card spending against CDTFA taxable sales (California Department of Tax and Fee Administration, 2023); the median multiplier is approximately $4\times$. Applying Los Angeles' combined 34.5 percent effective tax rate (10% local, 15% excise, 9.5% sales) yields $\overline{\text{Rev}}_{\text{sem}}$. Net welfare is:

$$\Delta W(d^{\min}) = \Delta W^{\text{HC}}(d^{\min}) - \Delta W^{\text{Rev}}(d^{\min}). \tag{16}$$

This implements the restricted planner's criterion from Section 6.1: each policy is evaluated by the combination of human-capital gains and fiscal revenue it produces.

Pigouvian alternative. A school buffer imposes an implicit infinite tax within the restricted radius. A Pigouvian alternative allows retailers anywhere provided they internalize the externality via distance-varying add-on rates $\tau_b^{\text{add}} = L_b^{HC}/Q_b$, where L_b^{HC} is the share of status-quo HC loss allocated to bin b and Q_b is the calibrated tax base for stores in that bin. No direct empirical estimate of a cannabis retailer entry semi-elasticity with respect to tax rates exists in the literature: Hansen, Miller and Weber (2017) estimate price pass-through and the consumer demand response to Washington's 2015 tax restructuring (finding a short-run price elasticity of approximately -0.4), while Hollenbeck and Uetake (2021) study tax incidence and market power in Washington's capped-entry market—where retail licenses are fixed by regulation at 550, precluding identification of an entry margin. I therefore calibrate the entry response from demand-side bounds. Under full pass-through and free entry, the long-run entry semi-elasticity would equal the demand price elasticity in magni-

tude; because retailers in this market exercise local market power (Hollenbeck and Uetake, 2021)—implying add-on taxes compress profit margins by less than the competitive benchmark, dampening exit incentives—I set $\varepsilon = 0.50$, below the competitive demand-elasticity bound, so that $M_{\tau t}^{\text{pig}} = M_{\tau t}^{\text{obs}} \times (1 - \varepsilon\, \bar{\tau}_{\tau t}^{\text{add}})$, where $\bar{\tau}_{\tau t}^{\text{add}}$ is the logit-share-weighted average add-on rate.[6]

6.6 Counterfactual Results

6.6.1 School Buffer Results

Table 5 reports outcomes across 11 buffer distances from 0.0 to 1.0 miles. Relative to the 600-foot status quo, predicted entry falls sharply as buffers tighten: -138.2 stores at 0.3 miles, -381.6 at 0.5 miles, and -591.3 at 1.0 miles. Figure A20 shows entry declines rapidly once the buffer exceeds 0.2–0.3 miles, then flattens as feasible locations become heavily constrained.

[Insert Table 5 Around Here]

Human-capital gains and revenue losses move in opposite directions as the buffer tightens. HC gains rise steeply to $251.0 million at 0.5 miles, then saturate ($276.6M at 1.0 miles), while revenue losses increase monotonically to $214.2 million at 1.0 miles. Panel B of Figure 9 plots net welfare ΔW, which is hump-shaped with an interior maximum at $d^* = 0.5$ miles, yielding net gains of $112.7 million. Beyond 0.5 miles, marginal fiscal costs exceed marginal human-capital gains. The optimum is computed with human-capital losses discounted at 3 percent annually (Office of Management and Budget, 2003), while fiscal revenue accrues contemporaneously. A regulator who applies a higher implicit discount rate to future attainment gains, or who disproportionately weights near-term tax revenue, would favor a shorter optimal buffer—one interpretation of why the status quo 600-foot rule falls well short of the model-optimal distance.

[Insert Figure 9 Around Here]

6.6.2 Policy Comparison: Buffer versus Pigouvian Tax

Table 6 compares the optimal school buffer against the Pigouvian alternative. The school buffer at $d^* = 0.5$ miles achieves large human-capital gains ($251.0M) by aggressively re-

[6] Unlike the buffer counterfactual, the Pigouvian simulation does not iterate on competitive crowding: since no locations are excluded, the spatial choice set is unchanged from baseline and competitive pressure does not shift across distance bins. Given the small estimated competition parameters ($\hat{\delta}_1 = 0.015$, $\hat{\delta}_2 = 0.075$; Table 4), endogenous crowding is a second-order effect relative to the market-size adjustment.

stricting entry (-381.6 stores), at substantial fiscal cost ($138.2M foregone revenue), yielding net welfare of $112.7M. The Pigouvian tax allows retailers near schools provided they pay externality-correcting add-on rates, producing a moderate entry reduction (-131.8 stores) and smaller human-capital gains ($60.3M), but remaining retailers pay higher tax rates that generate a fiscal revenue gain of $150.3M, so net welfare reaches $210.5M—nearly twice the buffer's net gain. The choice between instruments depends on whether policymakers prioritize maximizing protection of the most-exposed students (favoring buffers) or maximizing net social welfare inclusive of fiscal consequences (favoring Pigouvian taxation). Appendix Table A3 reports sensitivity of all buffer results to $\eta \in \{0.55, 0.60, 0.65\}$: the optimal buffer distance remains $d^* = 0.5$ miles throughout. Appendix Table A4 reports sensitivity of the Pigouvian results to $\varepsilon \in \{0.33, 0.50, 0.67\}$: Pigouvian taxation yields higher net welfare than the optimal buffer across the full range.

[Insert Table 6 Around Here]

The model captures a margin that reduced-form estimates alone cannot recover: by embedding the estimated distance-decay functions in a spatial equilibrium, it traces how buffer-induced reallocation affects exposure through both channels, including whether protecting students near schools inadvertently increases residential exposure elsewhere. The model is static and treats markets as geographically independent, in the tradition of the spatial entry literature (Caoui, Hollenbeck and Osborne, 2022; Magnolfi et al., 2024). The welfare calculation similarly omits non-pecuniary harms and general-equilibrium responses in housing or school choice. On the scope of the welfare criterion: the restricted planner's objective assigns zero weight to adult consumers' surplus because cannabis consumers are voluntary market participants who can substitute across retailers or forgo purchases, while students bear proximity costs without any market participation—a distinction Sandmo (1975) formalizes between polluters and pollutees. A planner who also valued consumers' surplus would find smaller net gains from tighter buffers, since stricter siting restrictions reduce retail access for voluntary consumers. The counterfactuals should therefore be read as policy comparisons on the child-protection and fiscal margin, not as complete accounts of total social surplus across all agents. One additional scope condition is the distinction between tightening and relaxing existing restrictions. The dose-response functions are identified from variation in retailer proximity at distances that Los Angeles retailers actually occupy; the framework recovers welfare consequences of pushing buffers further than the status quo—the question of how far—but not of

permitting retailers closer than they currently locate in contexts where no within-buffer variation exists. Applying the framework to evaluate buffer relaxation in a city with a substantially stricter incumbent rule would require extrapolating the harm function beyond the range of identified variation.

7 Conclusion

This paper asks a question that confronts every city regulating the proximity of commercial disamenities to children: how close is too close? Using cannabis retailers in Los Angeles as a case study, I develop and apply a framework for evidence-based calibration of zoning regulations, combining reduced-form causal estimates of hyperlocal exposure effects on youth human capital with a spatial entry model that predicts firm reallocation under counterfactual policies.

Three sets of findings emerge. First, the opening of a cannabis retailer causes immediate and sustained harm to nearby students: drug-related suspensions rise sharply, term GPA declines by 0.012–0.028 points, and these effects cumulate into reductions in high school graduation (1.2 percentage points for school exposure, 0.4 percentage points for neighborhood exposure) and college enrollment (0.9 percentage points for school exposure, 0.5 percentage points for neighborhood exposure). Second, these harms are not borne equally: effects are larger for low-income families and students with lower baseline academic ability, as well as those in peer environments with high baseline suspension rates, implying that cannabis retailers amplify existing educational inequalities. Third, the effects are highly localized, concentrated within 0.5 miles for school exposure and 0.3 miles for neighborhood exposure, with school proximity operating as the dominant channel.

These findings carry direct policy implications. The current 600-foot school buffer in Los Angeles is far narrower than the 0.5-mile radius over which school-channel effects remain significant, leaving students substantially exposed. Embedding the causal estimates in a spatial entry model—which captures the reallocation of retailers when entry near schools is restricted—reveals that the optimal buffer is 0.5 miles, yielding net welfare gains of $112.7 million relative to the status quo. A distance-varying Pigouvian tax calibrated to internalize the human-capital externality achieves nearly twice the net welfare gain ($210.5 million) by generating fiscal revenue from remaining retailers rather than suppressing the tax base, though at the cost of less aggressive student protection. The choice between instruments depends on whether policymakers prioritize maximizing protection of exposed students or maximizing net social welfare inclusive of fiscal consequences.

Several limitations warrant mention. My estimates are identified from the Los Angeles market and may differ in other settings. I measure human capital through administrative

records—GPA, graduation, and college enrollment—which may not capture the full range of consequences. The spatial entry model is static and abstracts from dynamic strategic interactions, and the partial-equilibrium welfare calculation omits general-equilibrium responses in housing markets, residential sorting, and school choice. These limitations suggest the results should be interpreted as informative first-order approximations rather than precise welfare statements.

More broadly, this paper provides a portable framework for evidence-based zoning policy. The approach requires only a credibly estimated distance-decay of harm and a model of firm location choice, making it applicable to any establishment that generates localized externalities on children: alcohol outlets (Livingston, 2008), fast-food restaurants (Currie et al., 2010), gambling venues, vaping shops (Henriksen et al., 2008), and payday lenders (Melzer, 2011). The welfare frontier maps any proposed regulation onto its human-capital and fiscal consequences, enabling policymakers to calibrate buffer distances even when they weigh these objectives differently. The comparison between uniform buffers and Pigouvian taxes illustrates a general lesson for place-based regulation: price-based instruments that internalize externalities while preserving locational choice can dominate categorical prohibitions on net welfare grounds, particularly when the regulated activity generates tax revenue that policymakers must also consider.

References

Aizer, Anna, and Janet Currie. 2019. "Lead and Juvenile Delinquency: New Evidence from Linked Birth, School, and Juvenile Detention Records." *Review of Economics and Statistics*, 101(4): 575–587.

Almond, Douglas, Janet Currie, and Valentina Duque. 2018. "Childhood Circumstances and Adult Outcomes: Act II." *Journal of Economic Literature*, 56(4): 1360–1446.

Angrist, Joshua D., and Alan B. Krueger. 1991. "Does Compulsory School Attendance Affect Schooling and Earnings?" *The Quarterly Journal of Economics*, 106(4): 979–1014.

Associated Press. 2020. "State Mounts Largest Crackdown on Illegal Pot Shops in LA." *ABC News*, February 14, 2020. Available at https://abcnews.go.com/US/wireStory/state-mounts-largest-crackdown-illegal-pot-shops-la-67725936.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890.

Beverly Press and Park La Brea News. 2022. "LAUSD Targets Marijuana Dispensaries near Campuses." *Beverly Press, June 1*, Reporting on LAUSD Superintendent Alberto M. Carvalho's initiative to close or relocate dispensaries near schools.

Billings, Stephen B, and Kevin T Schnepel. 2018. "Life after Lead: Effects of Early Interventions for Children Exposed to Lead." *American Economic Journal: Applied Economics*, 10(3): 315–344.

California Department of Tax and Fee Administration. 2023. "Cannabis Tax Revenues: Taxable Sales by County." *https://cdtfa.ca.gov/dataportal/dataset.htm?url=CannabisSalesByCounty*, Quarterly taxable sales data for licensed cannabis businesses.

Callaway, Brantly, and Pedro H.C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics*, 225(2): 200–230.

Cannabis Business Times. 2022. "Cannabis Dispensary Transactions, Tickets, and Banking Trends." *Industry survey report*, Approximately 70–80 percent of cannabis dispensary transactions by volume are conducted in cash.

Caoui, El Hadi, Brett Hollenbeck, and Matthew Osborne. 2022. "The Impact of Dollar Store Expansion on Local Market Structure and Food Access." *SSRN Electronic Journal*.

Card, David. 1999. "The Causal Effect of Education on Earnings." In *Handbook of Labor Economics*. Vol. 3, , ed. Orley Ashenfelter and David Card, 1801–1863. Elsevier.

CBS News Los Angeles. 2021. "California Cannabis Dispensary Regulations Prohibiting Minors Not Working as Hoped, According to Study." *CBS News Los Angeles, August 31*, Reporting on Shi and Pacula (2021).

Chatterji, Pinka. 2006. "Illicit Drug Use and Educational Attainment." *Health Economics*, 15(5): 489–511.

Chetty, Raj, and Nathaniel Hendren. 2018. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." *Quarterly Journal of Economics*, 133(3): 1107–1162.

Chetty, Raj, John N Friedman, and Jonah E Rockoff. 2014. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review*, 104(9): 2633–2679.

Chetty, Raj, Nathaniel Hendren, and Lawrence 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.

Currie, Janet, and Reed Walker. 2011. "Traffic Congestion and Infant Health: Evidence from E-ZPass." *American Economic Journal: Applied Economics*, 3(1): 65–90.

Currie, Janet, Joshua Graff Zivin, Jamie Mullins, and Matthew Neidell. 2014. "What Do We Know about Short- and Long-Term Effects of Early-Life Exposure to Pollution?" *Annual Review of Resource Economics*, 6(1): 217–247.

Currie, Janet, Lucas Davis, Michael Greenstone, and Reed Walker. 2015. "Environmental Health Risks and Housing Values: Evidence from 1,600 Toxic Plant Openings and Closings." *American Economic Review*, 105(2): 678–709.

Currie, Janet, Matthew Neidell, and Johannes F. Schmieder. 2009. "Air Pollution and Infant Health: Lessons from New Jersey." *Journal of Health Economics*, 28(3): 688–703.

Currie, Janet, Stefano DellaVigna, Enrico Moretti, and Vikram Pathania. 2010. "The Effect of Fast Food Restaurants on Obesity and Weight Gain." *American Economic Journal: Economic Policy*, 2(3): 32–63.

Davis, Lucas W. 2011. "The Effect of Power Plants on Local Housing Values and Rents." *Review of Economics and Statistics*, 93(4): 1391–1402.

DeSimone, Jeff. 2002. "Illegal Drug Use and Employment." *Journal of Labor Economics*, 20(4): 952–977.

Dills, Angela K., Sietse Goffard, and Jeffrey Miron. 2017. "The Effects of Marijuana Liberalizations: Evidence from Monitoring the Future." National Bureau of Economic Research NBER Working Paper 23779.

Eaton, Jonathan, Samuel Kortum, and Francis Kramarz. 2004. "Dissecting Trade: Firms, Industries, and Export Destinations." *American Economic Review*, 94(2): 150–154.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics*, 225(2): 254–277.

Greenstone, Michael, and Justin Gallagher. 2008. "Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program." *The Quarterly Journal of Economics*, 123(3): 951–1003.

Hansen, Benjamin, Keaton Miller, and Caroline Weber. 2017. "Getting into the Weeds of Tax Invariance." National Bureau of Economic Research NBER Working Paper 23632.

Henriksen, Lisa, Ellen C. Feighery, Nicole C. Schleicher, David W. Cowling, R. Sam Kline, and Stephen P. Fortmann. 2008. "Is Adolescent Smoking Related to the Density and Proximity of Tobacco Outlets and Retail Cigarette Advertising Near Schools?" *Preventive Medicine*, 47(2): 210–214.

Hollenbeck, Brett, and Kosuke Uetake. 2021. "Taxation and Market Power in the Legal Marijuana Industry." *RAND Journal of Economics*, 52(3): 559–595.

Hoxby, Caroline M. 2000. "Peer Effects in the Classroom: Learning from Gender and Race Variation." National Bureau of Economic Research NBER Working Paper 7867.

Jayawardhana, Jayani, Jialin Hou, and Johanna Catherine Maclean. 2025. "Cannabis and Respiratory Health." National Bureau of Economic Research NBER Working Paper 34435.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard. 2019. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics*, 11(4): 181–209.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz. 2007. "Experimental Analysis of Neighborhood Effects." *Econometrica*, 75(1): 83–166.

Linden, Leigh, and Jonah E. Rockoff. 2008. "Estimates of the Impact of Crime Risk on Property Values from Megan's Laws." *American Economic Review*, 98(3): 1103–1127.

Livingston, Michael. 2008. "A Longitudinal Analysis of Alcohol Outlet Density and Assault." *Alcoholism: Clinical and Experimental Research*, 32(6): 1074–1079.

Magnolfi, Lorenzo, Corina Mommaerts, Natalia Serna, and Christopher Sullivan. 2024. "The Rise of Urgent Care Centers: Implications for Competition and Access to Health Care." *Journal of Political Economy Microeconomics*, 2(2): 201–243.

Mankiw, N Gregory, and Michael D Whinston. 1986. "Free Entry and Social Inefficiency." *RAND Journal of Economics*, 17(1): 48–58.

Markowitz, Sara, and Katie E. Leinenbach. 2025. "Marijuana Legalization and Suicides Among Older Adults." National Bureau of Economic Research NBER Working Paper 34519.

Mathur, Neeraj K., and Christopher J. Ruhm. 2023. "Marijuana Legalization and Opioid Deaths." *Journal of Health Economics*, 88: 102728.

Melzer, Brian T. 2011. "The Real Costs of Credit Access: Evidence from the Payday Lending Market." *The Quarterly Journal of Economics*, 126(1): 517–555.

NBC Los Angeles. 2017. "Pot Shop Near School." NBC Los Angeles, April 2017.

Office of Management and Budget. 2003. "Circular A-4: Regulatory Analysis." *Executive Office of the President*, Recommends reporting central discount rates of 3 percent and 7 percent.

Office of the City Controller, City of Los Angeles. 2021. "High Maintenance: Review of the City's Cannabis Regulation Efforts." *Los Angeles City Controller Audit Report*, Available at https://controller.lacity.gov/audits/review-of-the-citys-cannabis-regulation-efforts.

Oreopoulos, Philip, and Kjell G Salvanes. 2011. "Priceless: The Nonpecuniary Benefits of Schooling." *Journal of Economic Perspectives*, 25(1): 159–184.

Pacula, Rosalie Liccardo, Karen E. Ross, and Jeanne Ringel. 2003. "Does Marijuana Use Impair Human Capital Formation?" , (9963).

Powell, David, Rosalie Liccardo Pacula, and Mireille Jacobson. 2018. "Do Medical Marijuana Laws Lead Younger People to Use Marijuana? A Review of the Evidence." *Journal of Health Economics*, 58: 58–72.

Roebuck, M. Christopher, Michael T. French, and Mallie L. Dennis. 2004. "Adolescent Marijuana Use and School Attendance." *Economics of Education Review*, 23(2): 133–141.

Sabia, Joseph J., Dhaval M. Dave, Fawaz Alotaibi, and Daniel I. Rees. 2021. "Is Recreational Marijuana a Gateway to Harder Drug Use and Crime?" National Bureau of Economic Research NBER Working Paper 29038.

Sacerdote, Bruce. 2001. "Peer Effects with Random Assignment: Results for Dartmouth Roommates." *The Quarterly Journal of Economics*, 116(2): 681–704.

Sandmo, Agnar. 1975. "Optimal Taxation in the Presence of Externalities." *The Swedish Journal of Economics*, 77(1): 86–98.

San Fernando Sun. 2017. "Parents Protest Marijuana Dispensary Near Vaughn Schools in Pacoima." San Fernando Sun, April 5, 2017.

Seim, Katja. 2006. "An Empirical Model of Firm Entry with Endogenous Product-Type Choices." *RAND Journal of Economics*, 37(3): 619–640.

Sharkey, Patrick. 2010. "The Acute Effect of Local Homicides on Children's Cognitive Performance." *Proceedings of the National Academy of Sciences*, 107(26): 11733–11738.

Shi, Yuyan, and Rosalie Liccardo Pacula. 2021. "Assessment of Recreational Cannabis Dispensaries' Compliance with Underage Access and Marketing Restrictions in California." *JAMA Pediatrics*, 175(11): 1178–1180.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics*, 225(2): 175–199.

Turner, Matthew A, Andrew Haughwout, and Wilbert van der Klaauw. 2014. "Land Use Regulation and Welfare." *Econometrica*, 82(4): 1341–1403.

Yamada, Tetsuji, Maw-Cherng Kendrix, and Tadashi Yamada. 1996. "The Impact of Alcohol Consumption and Marijuana Use on High School Graduation." *Health Economics*, 5(1): 77–92.

Figures

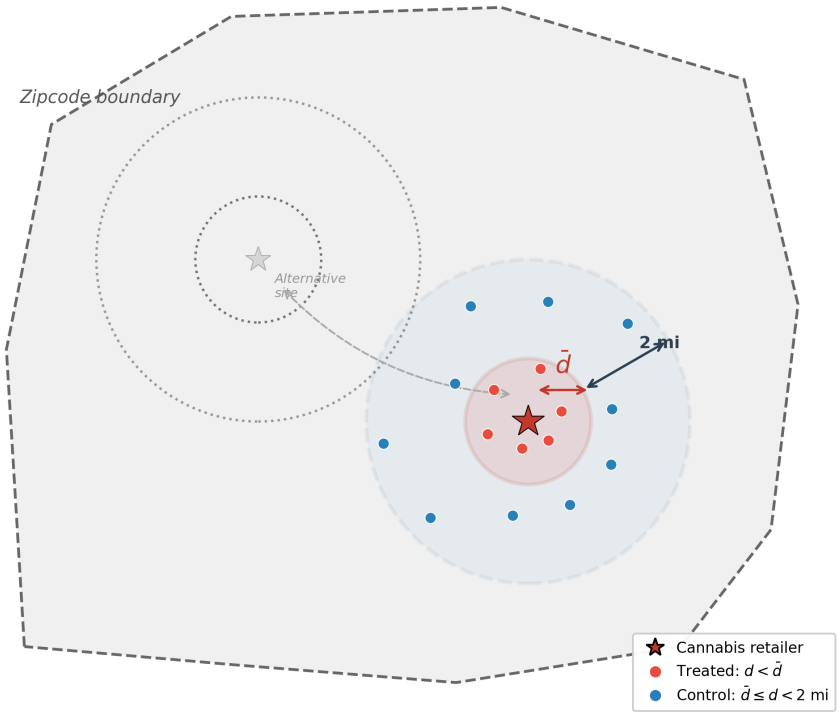
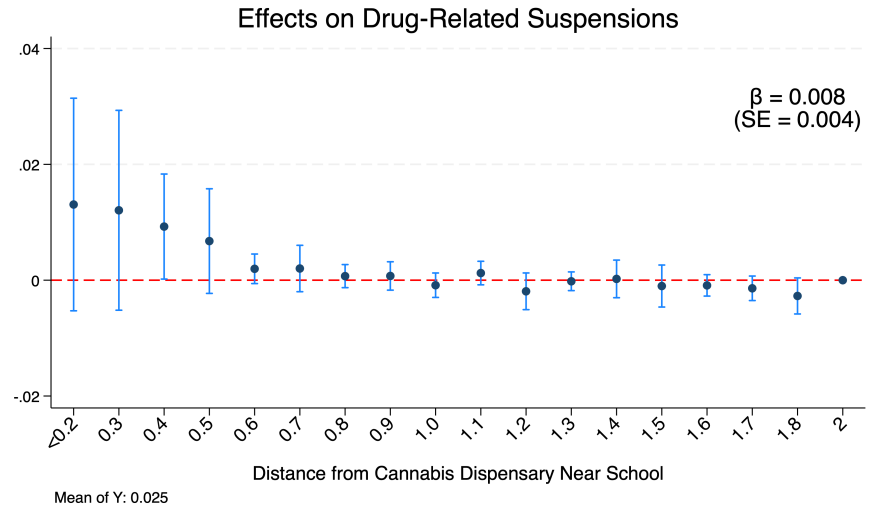


Figure 1 Identification Strategy: Hyperlocal Variation within Zipcode

Notes: This figure provides a schematic illustration of the identification strategy. The inner circle (solid, radius $\bar{d}$) delineates the treated zone: students whose homes or schools fall within $\bar{d}$ of the retailer. The outer ring (dashed, $\bar{d} \leq d < 2$ miles) delineates the control zone within the same zipcode. The faded alternative site illustrates that any feasible location within the zipcode would yield a comparable treated–control comparison. The threshold $\bar{d}$ is determined by data-driven distance-decay estimation (Equation (1)). Zipcode-by-semester fixed effects absorb all common shocks, and student fixed effects absorb all time-invariant individual characteristics.

[Back to Page 16]

Panel A: School Exposure



Panel B: Neighborhood Exposure

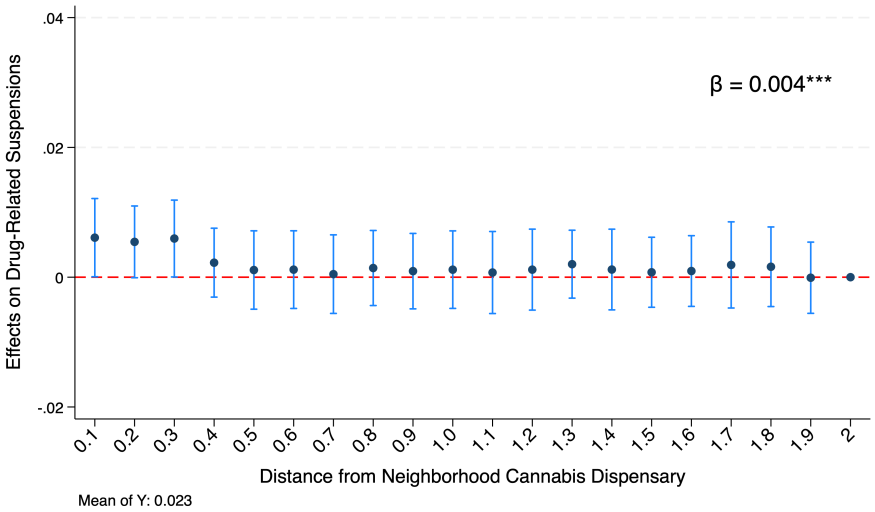
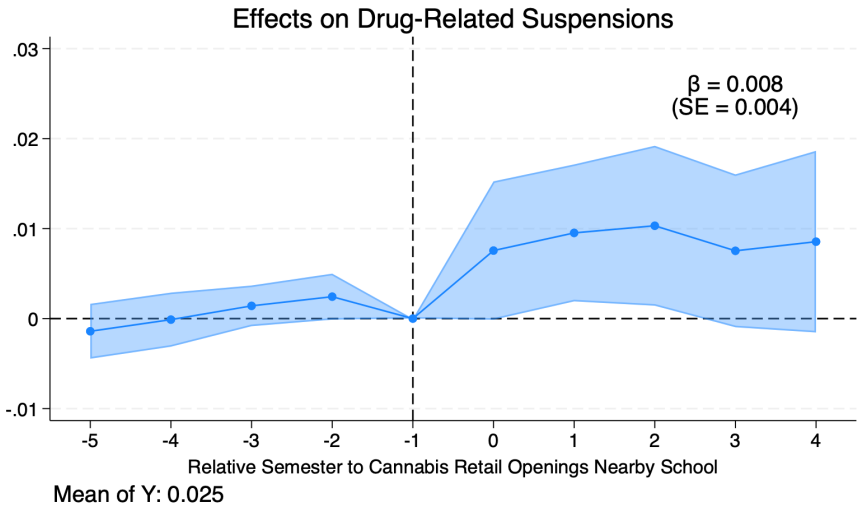


Figure 2 Distance-Decay in Drug-Related Suspensions

Notes: This figure estimates how the effect of cannabis retailer proximity on drug-related suspensions varies with distance. Each point shows the estimated coefficient $\hat{\beta}_d$ on the interaction of a 0.1-mile distance bin with a post-opening indicator, from the specification in Equation (1). The dependent variable is an indicator for drug-related disciplinary infractions. The omitted reference category is [1.9, 2.0) miles. Panel A measures distance from the student's enrolled school; Panel B measures distance from the student's home address. The sample includes LAUSD students within 2 miles of a retailer, 2009–2023. All specifications include student fixed effects and zipcode-by-semester fixed effects. Vertical bars represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 22]

Panel A: School Exposure



Panel B: Neighborhood Exposure

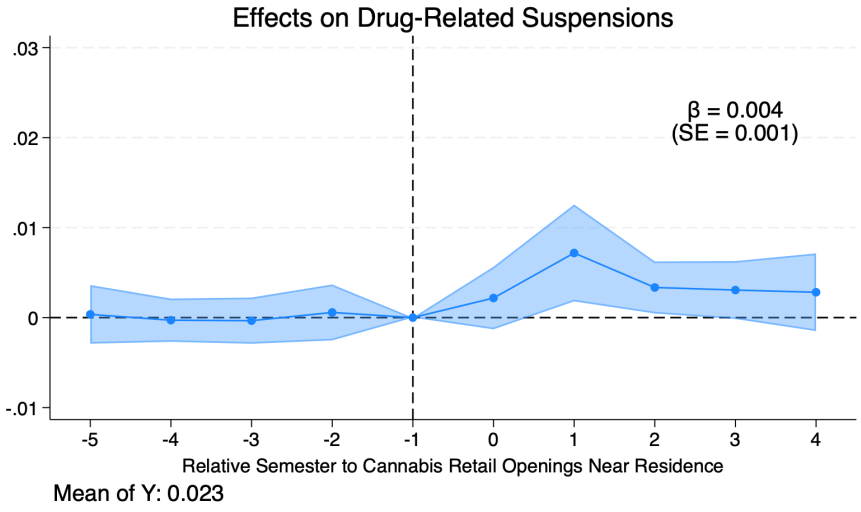
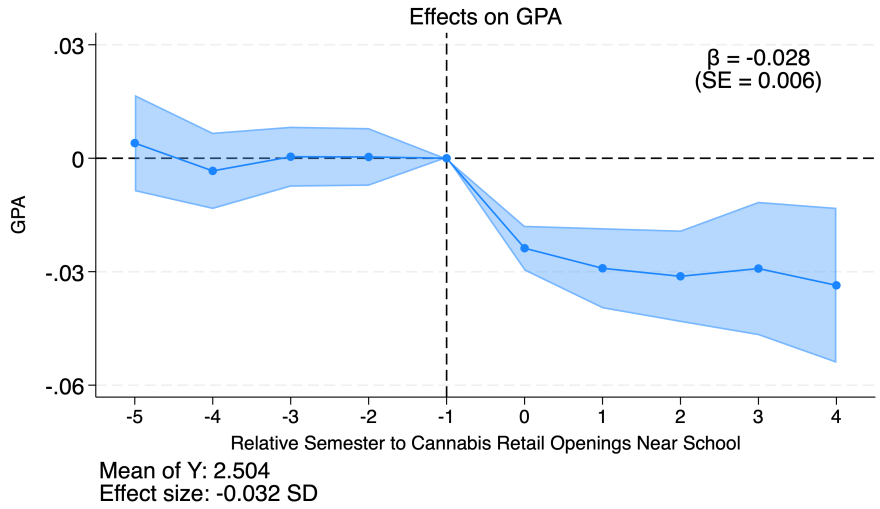


Figure 3 Event-Study Estimates: Effect of Retailer Opening on Drug-Related Suspensions

Notes: This figure presents event-study estimates of the effect of a nearby cannabis retailer opening on drug-related student suspensions. The dependent variable is an indicator for drug-related disciplinary infractions. The specification follows Equation (2). The x-axis shows semesters relative to the retailer opening; $\ell = -1$ is normalized to zero. Panel A shows school exposure (threshold: 0.5 miles); Panel B shows neighborhood exposure (threshold: 0.3 miles). Data are from LAUSD, 2009–2023. All specifications include student fixed effects and zipcode-by-semester fixed effects. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 23]

Panel A: School Exposure



Panel B: Neighborhood Exposure

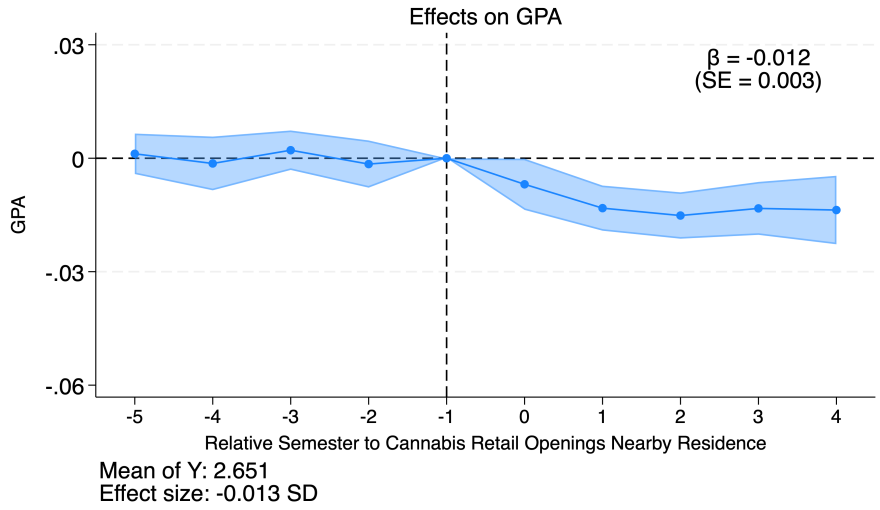
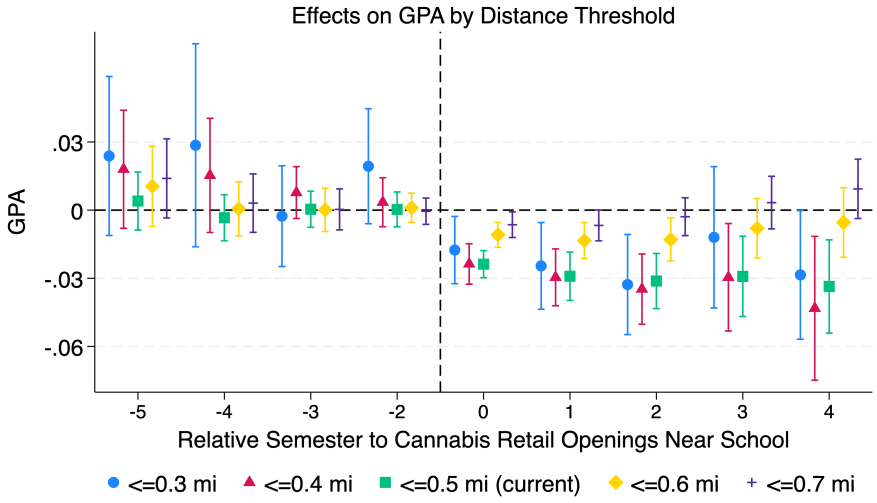


Figure 4 Event-Study Estimates: Effect of Retailer Opening on Student GPA

Notes: This figure presents event-study estimates of the effect of a nearby cannabis retailer opening on student term GPA (4.0 scale). The specification follows Equation (2). The x-axis shows semesters relative to the retailer opening; $\ell = -1$ is normalized to zero. Panel A shows school exposure (threshold: 0.5 miles); Panel B shows neighborhood exposure (threshold: 0.3 miles). Data are from LAUSD, 2009–2023. All specifications include student fixed effects and zipcode-by-semester fixed effects. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 25]

Panel A: School Exposure



Panel B: Neighborhood Exposure

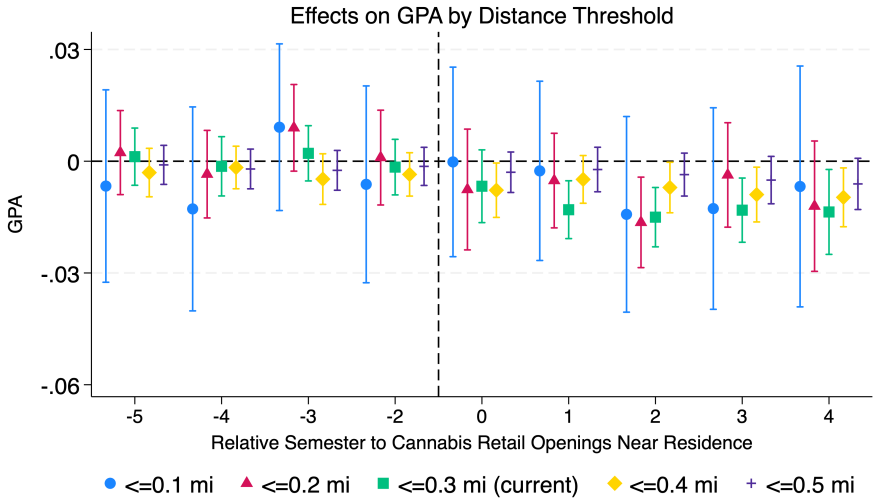


Figure 5 Sensitivity of GPA Event-Study Estimates to Treatment Distance Threshold
Notes: This figure overlays GPA event-study estimates under five alternative treatment distance thresholds. Panel A (school exposure) uses thresholds of 0.3, 0.4, 0.5 (current baseline, green), 0.6, and 0.7 miles from the student's enrolled school. Panel B (neighborhood exposure) uses thresholds of 0.1, 0.2, 0.3 (current baseline, green), 0.4, and 0.5 miles from the student's home address. The x-axis shows semesters relative to the retailer opening; $\ell = -1$ is normalized to zero. For a given threshold, the treated group consists of students within the specified radius and the control group consists of students located between that radius and 2 miles. Changing the threshold mechanically shifts which distance bands are classified as treated versus control: a smaller radius assigns students in the outer band to the control group, while a larger radius extends the treated group to more distant students. All specifications include student fixed effects and zipcode-by-semester fixed effects. Vertical bars denote 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 25]

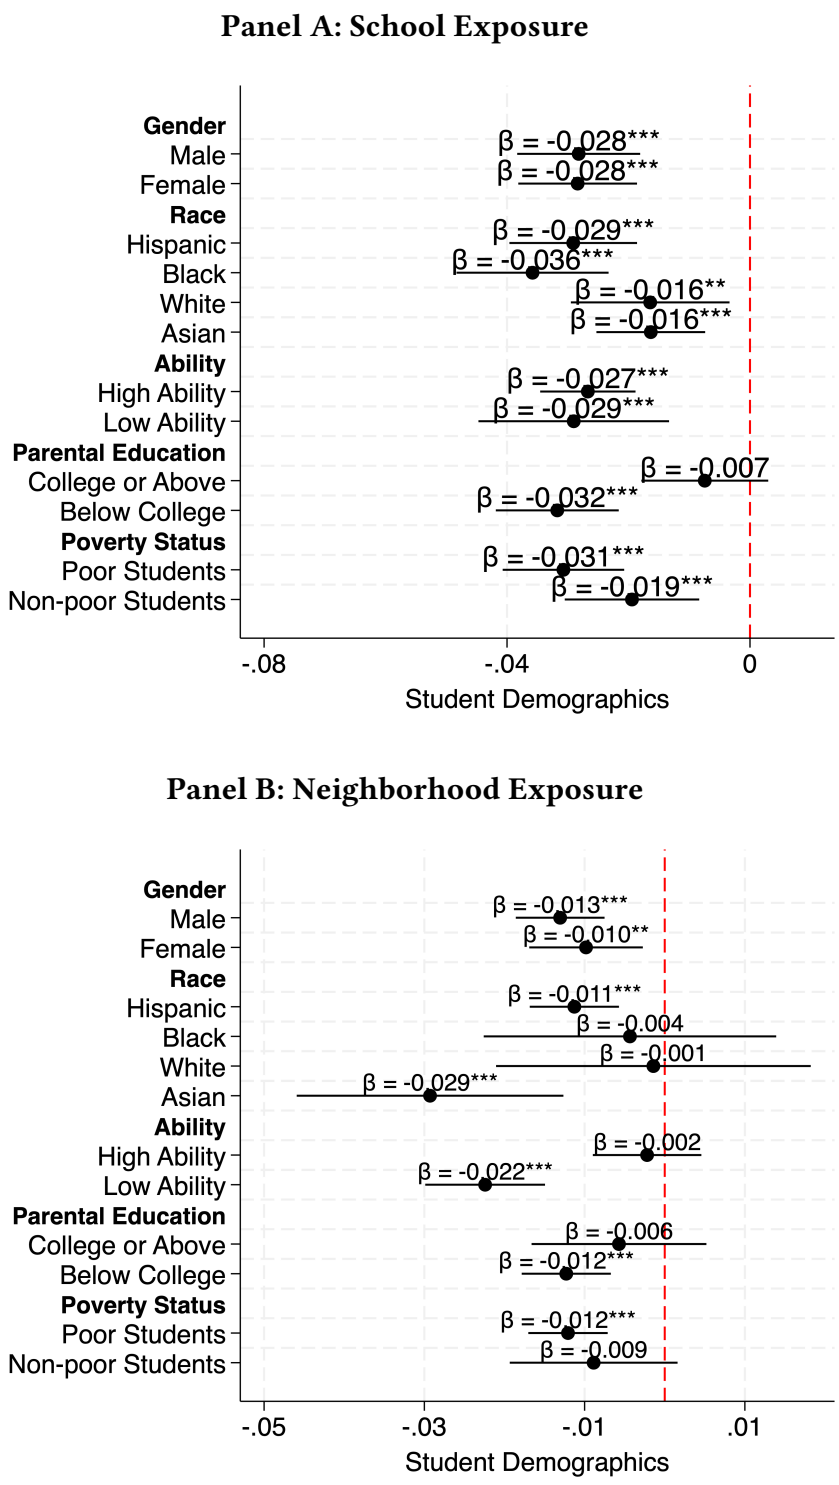
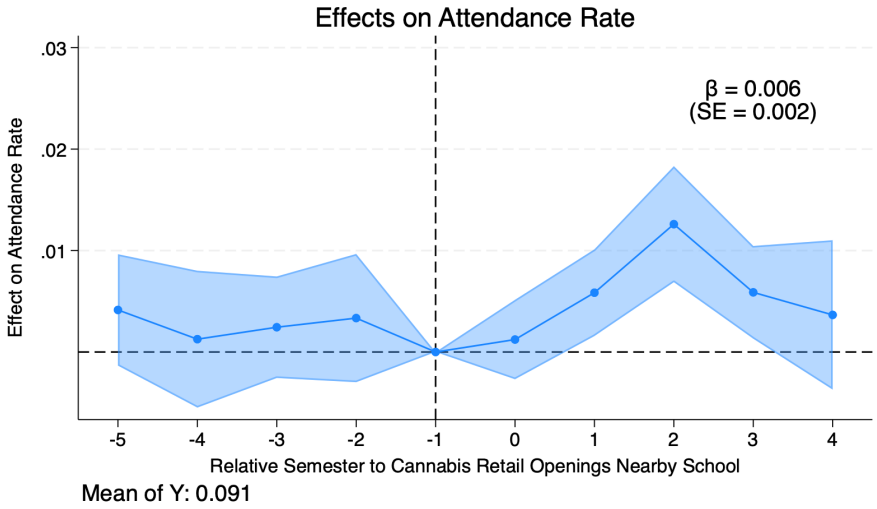


Figure 6 Heterogeneity in GPA Effects by Student Demographics

Notes: This figure presents the estimated effect of cannabis retailer proximity on GPA from the stacked difference-in-differences specification, estimated separately for demographic subgroups. Panel A shows school exposure (threshold: 0.5 miles); Panel B shows neighborhood exposure (threshold: 0.3 miles). Subgroups are defined by gender, race/ethnicity, baseline academic ability (above- vs. below-median pre-treatment GPA), parental education (college or above vs. below college), and poverty status (free/reduced-price lunch eligibility). Each point is the estimated coefficient; horizontal bars represent 90% confidence intervals. Data are from LAUSD, 2009–2023. All specifications include student fixed effects and zipcode-by-semester fixed effects. Standard errors are clustered at the zipcode level.

[Back to Page 27]

Panel A: School Exposure



Panel B: Neighborhood Exposure

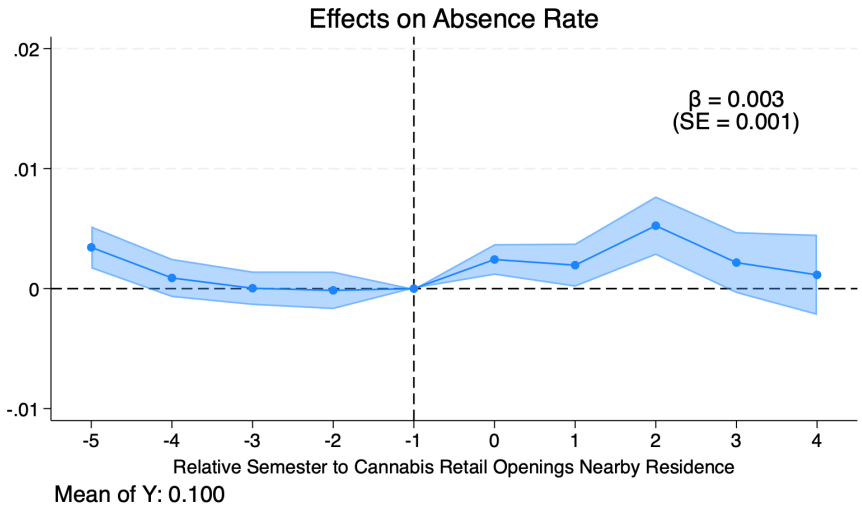


Figure 7 Event-Study Estimates: Effect of Retailer Opening on Student Absence Rate

Notes: This figure presents event-study estimates for student absence rates around cannabis retailer openings. The specification follows Equation (2). Panel A shows school exposure (threshold: 0.5 miles); Panel B shows neighborhood exposure (threshold: 0.3 miles). Data are from LAUSD, 2009–2023. All specifications include student fixed effects and zipcode-by-semester fixed effects. Standard errors are clustered at the zipcode level.

[Back to Page 30]

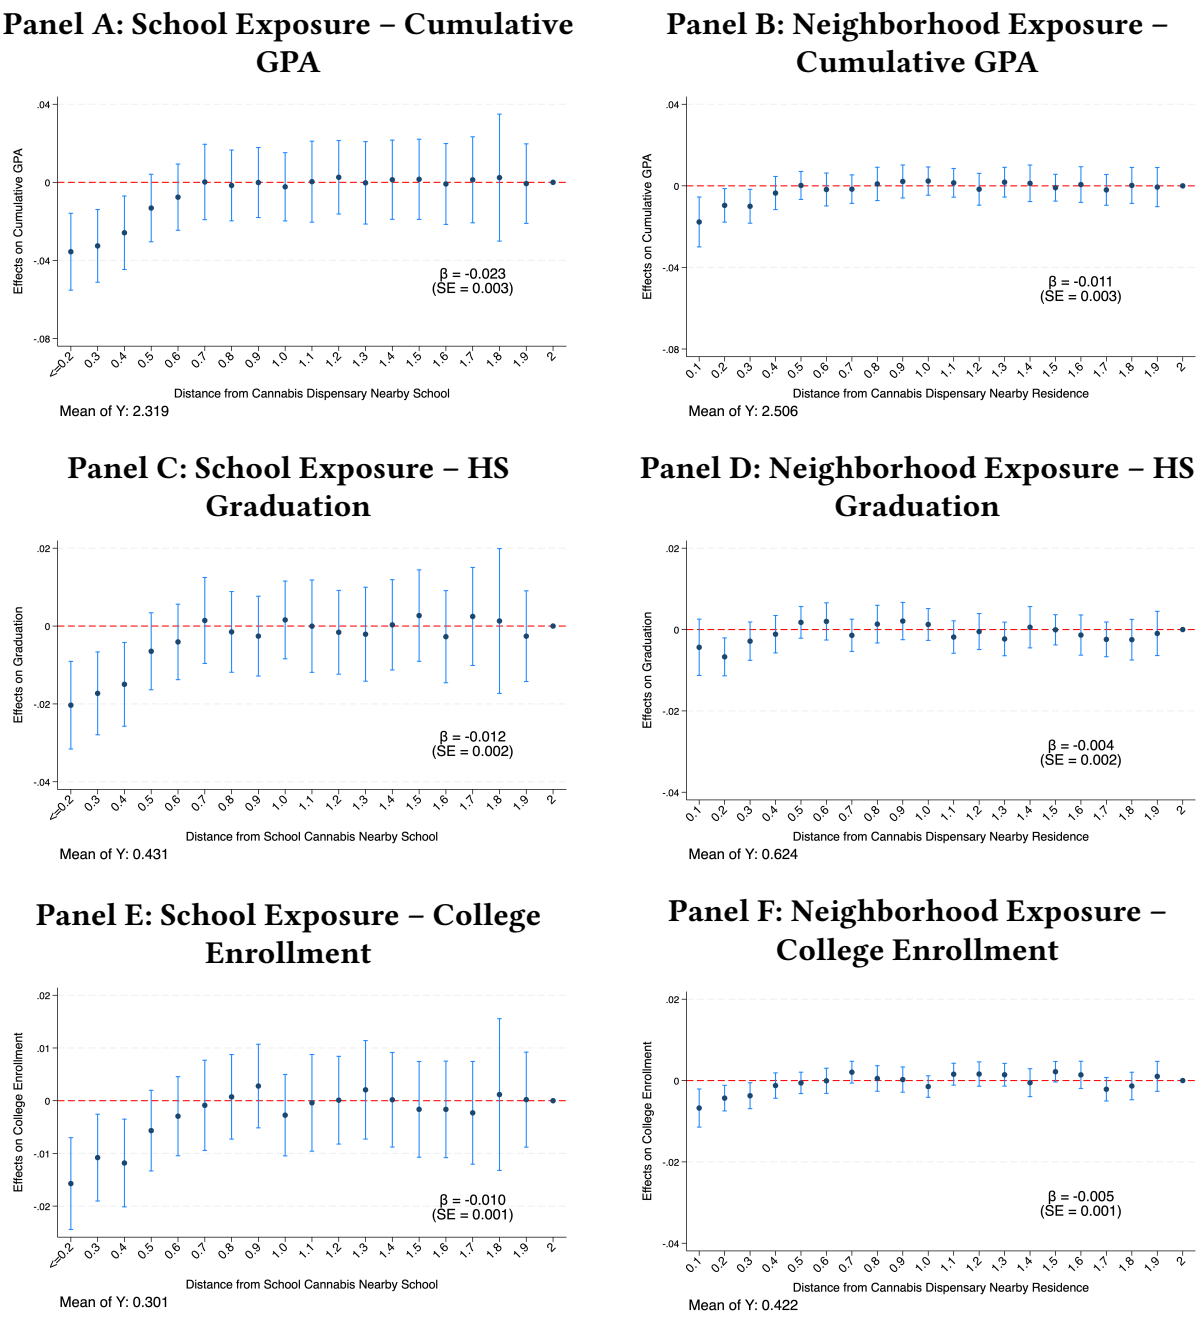


Figure 8 Distance-Decay Effects on Long-Run Human Capital Outcomes

Notes: This figure reports long-run distance-decay estimates by exposure channel. Panels A, C, and E use distance from the student's enrolled school; Panels B, D, and F use distance from the student's home address. Outcomes are cumulative GPA, on-time high school graduation, and first-fall college enrollment. Each coefficient $\hat{\beta}_b$ is estimated by interacting a 0.1-mile distance-bin indicator with a treated-cohort indicator ($TreatedCohort_i$), defined as one for students whose expected graduation cohort had not yet graduated at the time of the nearest retailer's opening, and zero for cohorts whose expected graduation predated the opening. School-channel specifications include campus-zipcode fixed effects, race controls, and gender controls, with standard errors clustered at the campus-zipcode level. Neighborhood-channel specifications include residence-zipcode fixed effects and cluster at the residence-zipcode level. Vertical bars show 90% confidence intervals.

[Back to Page 32]

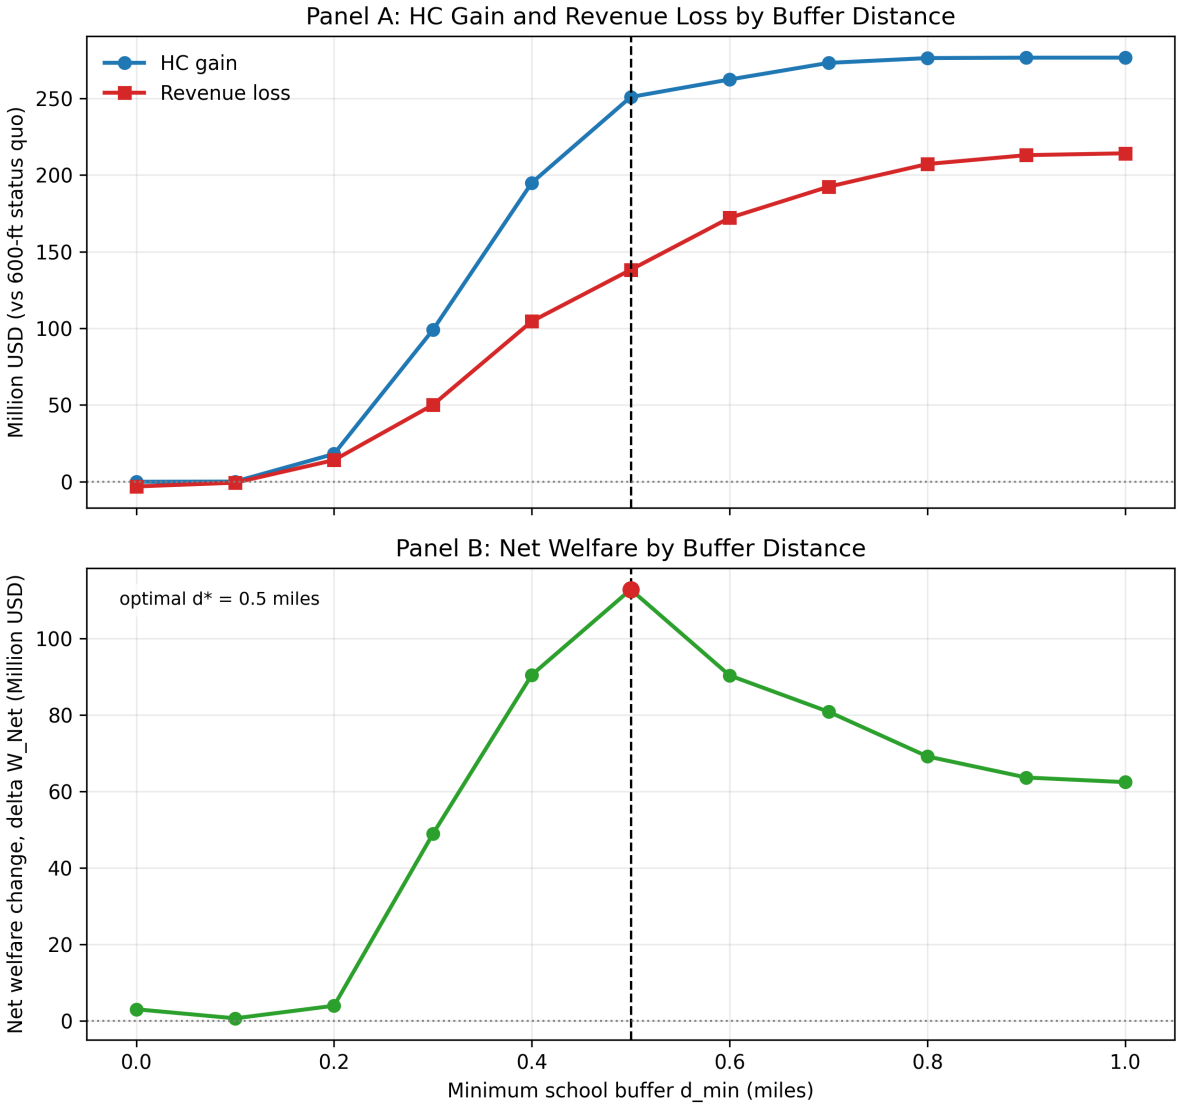


Figure 9 Welfare Components and Net Welfare by School Buffer Distance

Notes: Panel A plots human-capital loss reduction $\Delta W^{HC}(d^{\min})$ and revenue loss $\Delta W^{Rev}(d^{\min})$ against buffer distance. Panel B plots net welfare $\Delta W^{Net}(d^{\min}) = \Delta W^{HC}(d^{\min}) - \Delta W^{Rev}(d^{\min})$. Human-capital effects are computed from direct distance-decay coefficients on graduation and college enrollment (school and neighborhood channels), then monetized into discounted lifetime earnings losses. All values are relative to the current Los Angeles 600-foot rule ($d^{\min} \approx 0.114$ miles) over the grid $d^{\min} \in \{0.0, 0.1, \ldots, 1.0\}$. The vertical line marks the model-implied optimum on this grid.

[Back to Page 42]

Tables

Table 1 Balance of Pre-Treatment Covariates: Treated vs. Control Students

	School Exposure ($\bar{d} = 0.5$ mi)		**Neighborhood Exposure ($\bar{d} = 0.3$ mi)**	
	T−C	T−C | Zip FE	T−C	T−C | Zip FE
Age	0.062 (0.070)	0.074 (0.066)	0.010 (0.010)	0.008 (0.013)
Male	−0.013 (0.009)	−0.015* (0.009)	−0.001 (0.003)	−0.001 (0.005)
Hispanic	−0.019 (0.013)	−0.002 (0.012)	0.007 (0.006)	0.003 (0.005)
Black	0.007 (0.009)	−0.007 (0.007)	0.011*** (0.004)	0.003 (0.004)
White	0.003 (0.008)	0.006 (0.004)	−0.020*** (0.005)	−0.003 (0.007)
Asian	0.004 (0.005)	0.000 (0.004)	0.003 (0.002)	−0.001 (0.003)
Parent College	0.005 (0.008)	0.007 (0.010)	−0.015*** (0.004)	−0.002 (0.005)
Poverty	−0.033** (0.016)	−0.018 (0.016)	0.011** (0.005)	0.002 (0.006)
English at Home	−0.004 (0.013)	−0.013 (0.011)	−0.022*** (0.006)	−0.006 (0.008)

Notes: This table compares pre-treatment characteristics of treated and control students. "T−C" reports the unconditional mean difference. "T−C | Zip FE" reports the conditional mean difference from a regression with zipcode fixed effects. The left two columns use school exposure (treated = school within 0.5 miles; control = 0.5–2.0 miles). The right two columns use neighborhood exposure (treated = residence within 0.3 miles; control = 0.3–2.0 miles). "Parent College" indicates whether at least one parent has a college degree. "Poverty" indicates free/reduced-price lunch eligibility. "English at Home" indicates English as the primary home language. Data are from LAUSD administrative records. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[Back to Page 19]

Table 2 First-Stage Magnitudes: Cannabis Retailer Proximity and Drug-Related Student Outcomes

	(1) Drug-Related Suspensions	(2) First-Time Marijuana Use	(3) Non-First Marijuana Use
Panel A: School Exposure			
Post × Treat Dist	0.008*	0.002*	0.003*
	(0.004)	(0.001)	(0.002)
Mean	0.025	0.007	0.012
Observations	10,267,868	10,267,868	10,267,868
Panel B: Neighborhood Exposure			
Post × Treat Dist	0.004***	0.001***	0.002***
	(0.001)	(0.000)	(0.000)
Mean	0.023	0.008	0.010
Observations	11,931,197	11,931,197	11,931,197
Student FEs	YES	YES	YES
Zipcode×Semester FE	YES	YES	YES

Notes: Each coefficient is from a stacked difference-in-differences regression of the outcome on Post × Treat Dist with student fixed effects and zipcode-by-semester fixed effects. Panel A defines treatment as being within 0.5 miles of the student's school; Panel B defines treatment as being within 0.3 miles of the student's residence. Column (1) uses any drug-related suspension; Column (2) restricts to first-time marijuana possession offenses; Column (3) uses non-first marijuana offenses. Data are from LAUSD, 2009–2023. Standard errors clustered at the zipcode level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[Back to Page 23]

Table 3 Effect of Cannabis Retailer Proximity on Student GPA: Robustness across Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: School Exposure (threshold: 0.5 miles from school)								
Mean GPA: 2.504								
Post × Treat Dist	-0.028***	-0.023***	-0.023***	-0.029***	-0.023***	-0.017***	-0.029***	-0.027***
	(0.006)	(0.003)	(0.003)	(0.006)	(0.003)	(0.004)	(0.006)	(0.006)
Observations	10,267,868	10,267,868	10,267,868	10,255,909	9,540,414	10,267,848	9,962,100	10,267,868
Panel B: Neighborhood Exposure (threshold: 0.3 miles from residence)								
Mean GPA: 2.651								
Post × Treat Dist	-0.012***	-0.012***	-0.012***	-0.013***	-0.012***	-0.013***	-0.012***	-0.011***
	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)
Observations	11,931,197	11,931,197	11,204,725	11,202,488	10,331,775	11,930,349	11,534,871	11,931,197
Student FEs	✓	✓	✓	✓	✓	✓	✓	✓
Zipcode×Semester FE	✓	✓	✓	✓	✓		✓	✓
School FEs		✓	✓	✓	✓	✓	✓	✓
Year×Cohort FE			✓	✓	✓	✓	✓	✓
Stack (Event) FEs				✓	✓	✓	✓	✓
Violence Rate (CBG)					✓	✓	✓	✓
Tract×Semester FE						✓		
Drop Late Entrants (Gr 10-12)							✓	
Equal Stack Weights								✓

Notes: Each cell reports the coefficient on Post × Treat Dist from a difference-in-differences regression of term GPA (4.0 scale). Panel A measures school exposure (treated $<$ 0.5 miles; control $=$ 0.5–2.0 miles). Panel B measures neighborhood exposure (treated $<$ 0.3 miles; control $=$ 0.3–2.0 miles). Column (1) is the baseline with student and zipcode-by-semester FEs. Column (2) adds school FEs. Column (3) adds year-by-cohort FEs. Column (4) adds stack (event) FEs. Column (5) controls for CBG-level violence rates. Column (6) replaces zipcode-by-semester FEs with census tract-by-semester FEs. Column (7) drops students who first entered the District in grades 10–12. Column (8) re-weights so each event stack contributes equally. Data are from LAUSD, 2009–2023. Standard errors clustered at the zipcode level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[Back to Page 26]

Table 4 Spatial Entry Estimates

	(1) Broad	(2) Broad+Full	(3) +Comp	(4) +Quality	(5) Full
Log broad spend	0.1314 (0.7450)	-0.0456 (0.7556)	-0.0328 (0.7446)	1.6044** (0.7561)	1.4671* (0.8235)
Log broad customers	-0.0011 (0.9095)	0.2335 (0.9226)	0.2153 (0.9102)	-1.4977 (0.9200)	-1.3774 (0.9856)
Log broad spend/customer	-0.2023 (0.6584)	-0.1098 (0.6699)	-0.1241 (0.6616)	-1.7215** (0.7827)	-1.5610* (0.8493)
Log broad visitors	0.2809 (0.3583)	0.0407 (0.3239)	0.0834 (0.3266)	0.0713 (0.3954)	0.1513 (0.4679)
Log broad visits	-0.1410 (0.3466)	0.0584 (0.3153)	0.0136 (0.3189)	-0.0268 (0.3876)	-0.1232 (0.4592)
Log block-semester total spend		0.0737 (0.2031)	0.0653 (0.1973)	0.0771 (0.2585)	0.1129 (0.2406)
Log block-semester total customers		-0.1248 (0.2297)	-0.1132 (0.2228)	-0.1661 (0.2895)	-0.1616 (0.2610)
Log block-semester spend/customer		0.1446 (0.1913)	0.1429 (0.1877)	0.0523 (0.2428)	-0.0112 (0.2333)
Log block-semester visitors		0.1526 (0.2564)	0.1798 (0.2559)	-0.1677 (0.2707)	-0.2084 (0.2502)
Log block-semester visits		-0.0305 (0.2491)	-0.0600 (0.2477)	0.0301 (0.2565)	0.0600 (0.2406)
Incumbents within 0.25 mi (N_1)			0.2228*** (0.0785)	0.0639 (0.0414)	0.0151 (0.0448)
Incumbents in (0.25, 0.50] mi (N_2)			0.0748 (0.0718)	0.0808* (0.0432)	0.0750* (0.0398)
Log POI points				0.8831*** (0.0515)	0.8279*** (0.0552)
Log lagged store stock					1.0359*** (0.2260)
Distance to nearest school (miles)					0.3411 (0.5511)
Log population					-0.1766 (0.2731)
Log per-capita income					-0.0369 (0.4340)
Log median rent					0.1691 (0.5468)
Commercial land-use indicator					-0.2460 (0.2632)
Log parking count (0.25 mi)					0.0121 (0.1786)
Tract × semester FE	Yes	Yes	Yes	Yes	Yes
Observations	7,422	7,422	7,422	7,422	7,422

Notes: This table reports conditional-logit estimates of cannabis retailer location choice at the census block level. The dependent variable is block-level entry; tract×semester fixed effects define choice sets. Column (1) includes broad-retail demand shifters (spending and foot traffic from non-cannabis establishments). Column (2) adds block-semester total demand (all POI activity). Column (3) adds distance-banded competition (N_1: incumbents within 0.25 miles; N_2: incumbents 0.25–0.5 miles). Column (4) adds log POI points as a proxy for unobserved commercial attractiveness. Column (5) is the full specification with lagged store stock, school distance, and demographic/cost controls. Standard errors clustered at the tract level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[Back to Page 39]

Table 5 Counterfactual Welfare Analysis: Alternative School Buffer Distances

Buffer (miles)	Pred. Entry	Δ Stores	HC Gain ($)	Revenue Loss ($)	Net Welfare ($)
0.0	625	8.9	-191,368	-3,212,106	3,020,738
0.1	618	2.1	-54,554	-747,667	693,113
0.2	577	-38.7	22,123,963	14,005,159	8,118,804
0.3	478	-138.2	112,101,286	50,061,353	62,039,933
0.4	327	-288.6	216,850,147	104,560,457	112,289,690
0.5	234	-381.6	275,716,590	138,222,112	137,494,479
0.6	141	-475.1	288,363,591	172,093,038	116,270,553
0.7	85	-531.1	299,865,900	192,398,131	107,467,768
0.8	44	-572.0	303,127,476	207,225,629	95,901,847
0.9	28	-588.0	303,462,108	213,004,911	90,457,198
1.0	25	-591.3	303,468,027	214,220,432	89,247,594

Notes: Each row reports outcomes under a counterfactual minimum school buffer distance $d^{\min}$. "Pred. Entry" is the total predicted number of cannabis retailers across all census blocks. "Δ Stores" is the change in total entry relative to the current Los Angeles 600-foot rule ($d^{\min} = 600/5280 \approx 0.114$ miles). "HC Gain" is the reduction in discounted human-capital loss relative to the status quo, where distance-decay coefficients for graduation and college enrollment are mapped to lifetime earnings using a standard returns-to-schooling calibration (baseline $\rho = 12.5\%$, Card, 1999; Oreopoulos and Salvanes, 2011), a policy-evaluation present-value mapping (Chetty, Friedman and Rockoff, 2014), baseline PV earnings $1.2M, and a 3% discount rate (Office of Management and Budget, 2003). "Revenue Loss" is the fiscal term from market-size changes using a CDTFA-calibrated per-store-semester back-of-envelope tax mapping; negative values denote revenue gains. "Net Welfare" is HC Gain minus Revenue Loss. Monetary columns are aggregate present values across all evaluated students and all simulated markets, not per-student or per-semester welfare values. Results use the full conditional-logit specification from Table 4, column (5), with semester-level reallocation and inclusive-value scaling elasticity $\eta = 0.60$ (Eaton, Kortum and Kramarz, 2004 benchmark).

[Back to Page 42]

Table 6 Policy Comparison: School Buffer versus Pigouvian Distance Tax

Policy	Pred. Entry	Δ Stores	HC Gain	Revenue Term	Net Welfare
Status quo (600-ft)	616.0	0.0	0.0	0.0	0.0
Best school buffer ($d = 0.5$ mi)	234.4	−381.6	251.0	138.2	112.7
Pigouvian distance tax	484.2	−131.8	60.3	−150.3	210.5

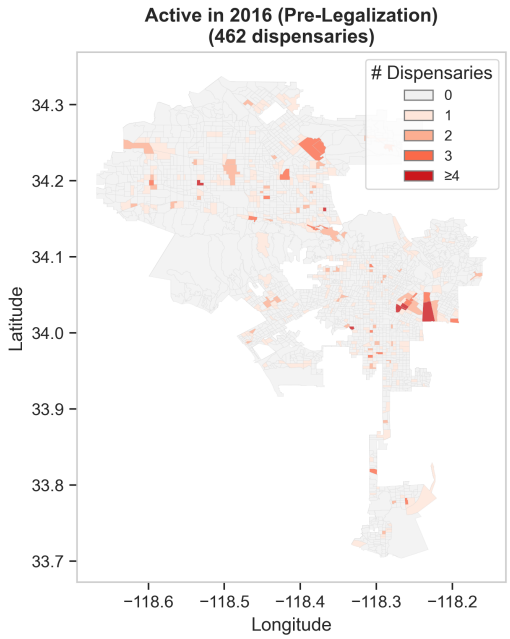
Notes: This table compares two policy alternatives to the current Los Angeles 600-foot school buffer. The school-buffer policy sets a uniform minimum distance $d^{\min}$ to all schools; the optimal choice on the grid $\{0.0, 0.1, \ldots, 1.0\}$ miles is $d^* = 0.5$. The Pigouvian policy imposes distance-varying add-on tax rates τ_b^{add} (in 0.1-mile bins) calibrated to internalize the human-capital externality, allowing retailers to locate anywhere if they pay the externality-correcting add-on. The Pigouvian entry response is simulated using a tax-entry semi-elasticity $\varepsilon = 0.50$, calibrated from Hollenbeck and Uetake (2021), who estimate that a one-percentage-point increase in the effective cannabis tax rate reduces the number of dispensaries by approximately 0.5 percent. This approach is analogous to the extensive-margin calibration (η) and avoids the need for an identified profit model. "HC Gain" measures the reduction in discounted human-capital losses (from improved graduation and college enrollment) relative to status quo, incorporating the asymmetric channel interaction (school exposure crowds out neighborhood exposure). "Revenue Term" captures foregone tax revenue from reduced entry (positive values) or additional revenue from the add-on tax (negative values). "Net Welfare" is HC Gain minus Revenue Term. All monetary values are aggregate present values in millions of 2026 USD across all simulated markets and evaluated students, computed using the full spatial entry model (Table 4, column 5) with extensive-margin elasticity $\eta = 0.60$.

[Back to Page 43]

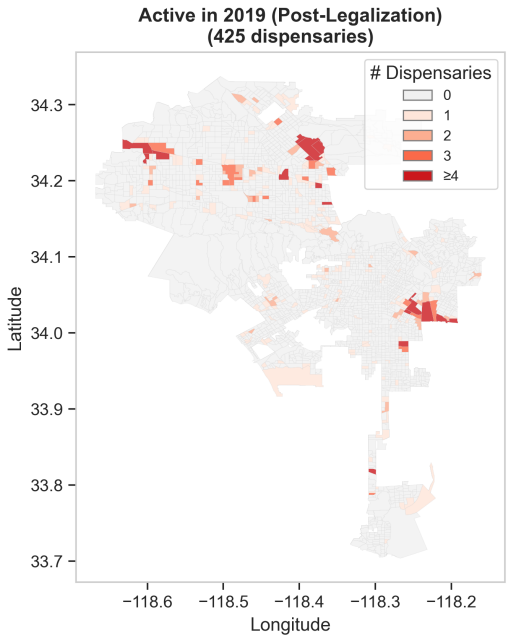
Appendix

Appendix Figures

Panel A: 2016 (Pre-Legalization)



Panel B: 2019 (Post-Legalization)



Panel C: 2023 (Current)

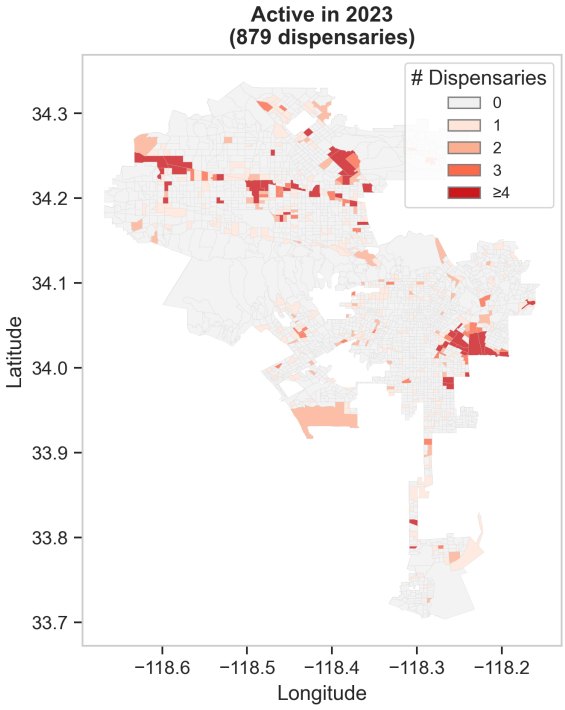
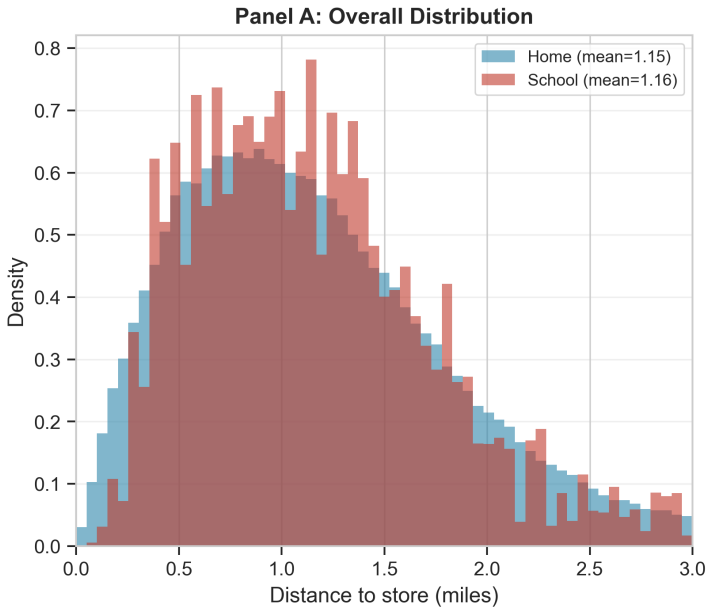


Figure A1 Annual Snapshots of the Cannabis Retail Market in Los Angeles

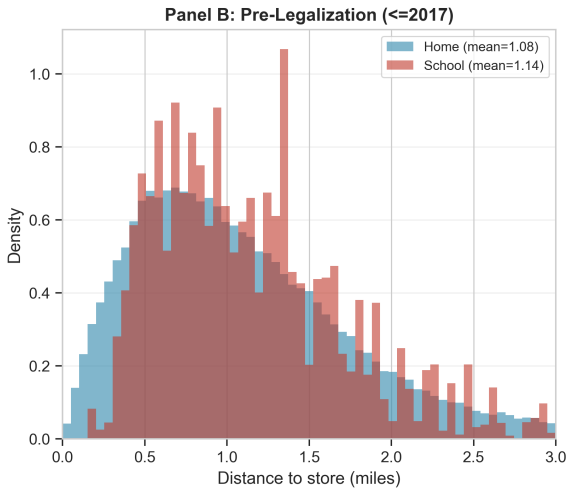
Notes: This figure presents cross-sectional snapshots of active licensed cannabis retailers (dispensaries) in the City of Los Angeles. Data are from the City of Los Angeles Office of Finance. Because retailers enter and exit over time, each panel captures only those outlets operating at the end of the indicated year. Panel A shows the retail market in 2016, prior to Proposition 64 (which legalized recreational sales in November 2016); Panel B shows the market in 2019, two years after recreational legalization took effect; Panel C shows the market in 2023, the final year of the sample. Each dot represents one licensed retailer location.

[Back to Page 11]

Panel A: Full Sample (2009-2023)



Panel B: Pre-Legalization (2010–2016)



Panel C: Post-Legalization (2017–2024)

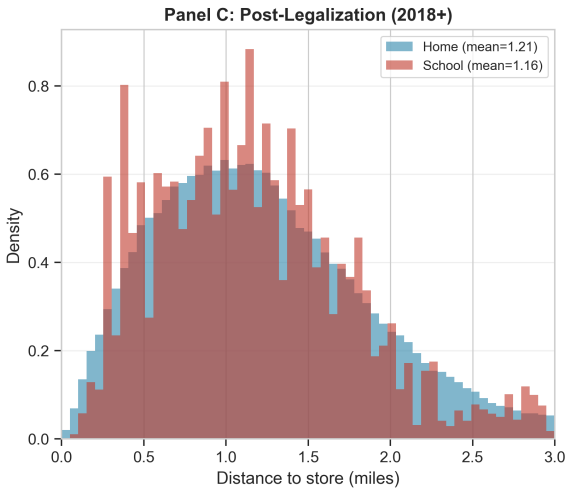


Figure A2 Distribution of Distance from Students to Nearest Cannabis Retailer

Notes: This figure shows the distribution of straight-line distances (in miles) from LAUSD students' home addresses to the nearest licensed cannabis retailer. Panel A pools all student-semester observations from 2009 to 2023. Panel B restricts to the medical-only era (before Proposition 64 in November 2016). Panel C shows the post-recreational-legalization period (2017–2024).

[Back to Page 14]

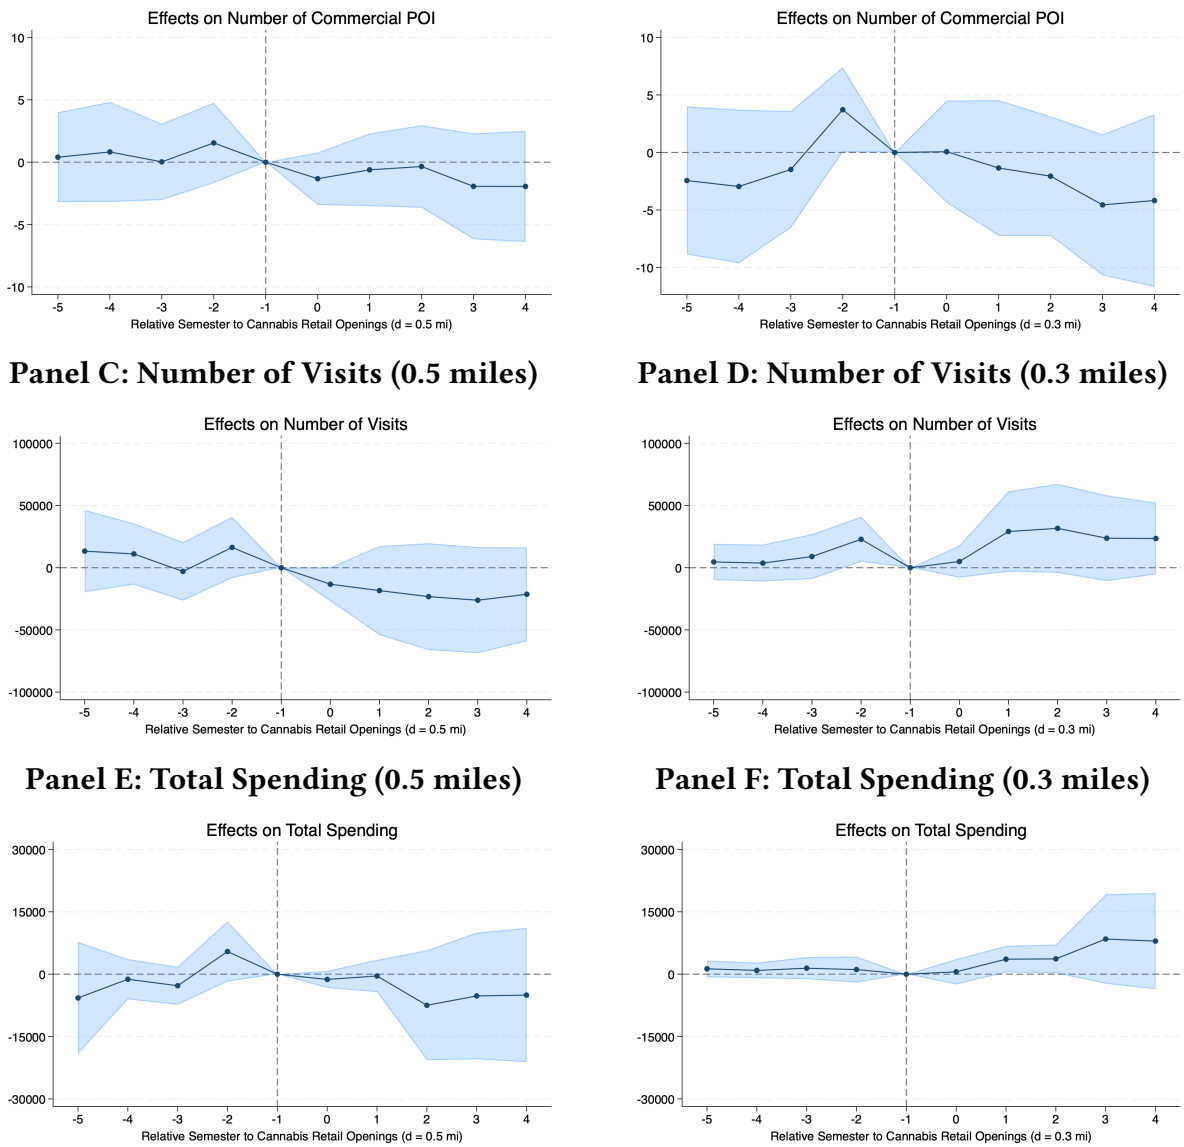
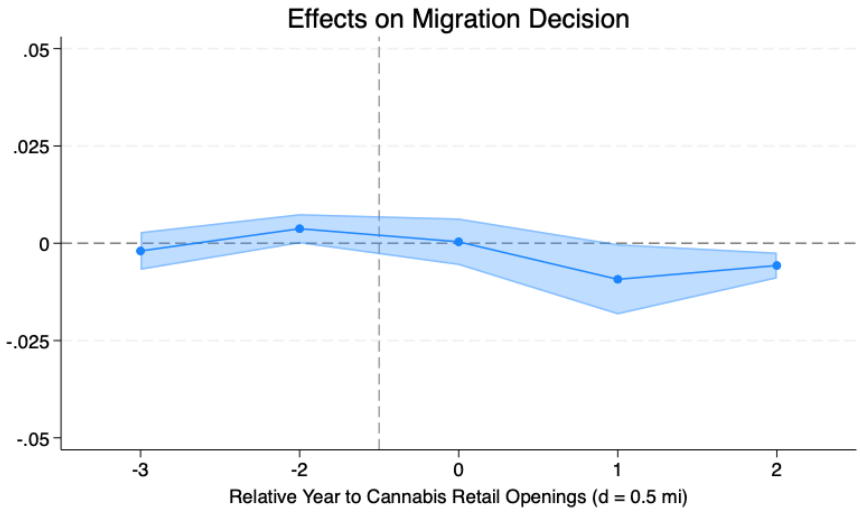


Figure A3 Placebo Check: Event-Study Effects on Broad Commercial Activity Around Cannabis Openings

Notes: This figure presents event-study estimates for commercial-environment outcomes around cannabis retailer openings. The six panels estimate the same stacked hyperlocal event-study design using block-semester commercial-environment measures from SafeGraph spending and Advan foot-traffic data. Panels A–B use total commercial POI counts (excluding cannabis), Panels C–D use total visit counts, and Panels E–F use total spending; odd panels use the 0.5-mile treatment threshold and even panels use 0.3 miles. The x-axis shows semesters relative to opening with $\ell = -1$ normalized to zero. All specifications include block fixed effects and tract-by-semester fixed effects. Shaded areas represent 90% confidence intervals based on standard errors clustered at the census tract level.

[Back to Page 20]

Panel A: 0.5-Mile Threshold



Panel B: 0.3-Mile Threshold

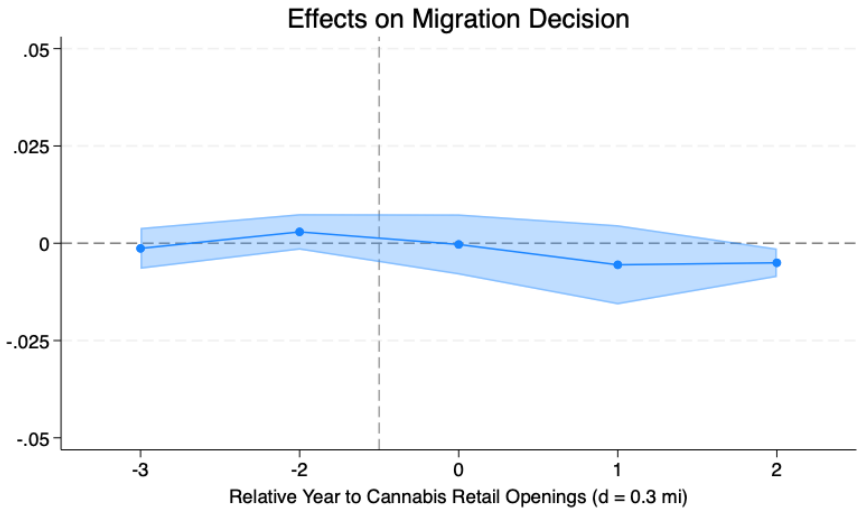


Figure A4 Addressing Selective Migration: Individual Residential Mobility from L2 Voter File Data

Notes: This figure presents event-study estimates of individual-level residential migration using L2 voter file data spanning 2014–2022. The dependent variable is an indicator for whether a registered voter changes their home address between consecutive years. The event-study specification follows the stacked design from Equation (2), restricting to registered voters within 2 miles of a retailer, and is indexed relative to the nearest cannabis retailer opening; $\ell = -1$ is normalized to zero. Panel A applies a 0.5-mile treatment threshold; Panel B applies a 0.3-mile threshold. All specifications include individual fixed effects and tract-by-semester fixed effects. Shaded areas represent 90% confidence intervals based on standard errors clustered at the census tract level.

[Back to Page 20]

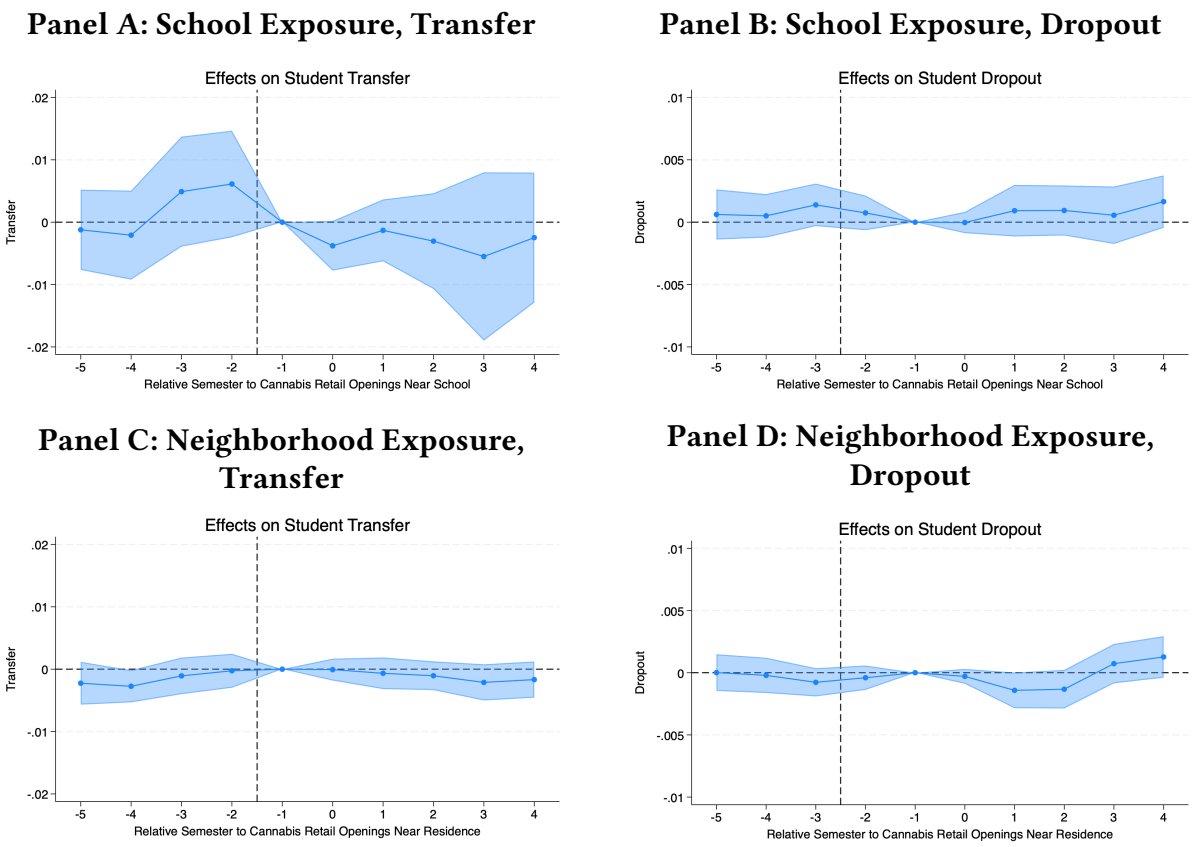


Figure A5 Addressing Selective Migration: Student Transfers and Dropouts (Full Sample)

Notes: This figure reports event-study estimates for transfer-out and dropout outcomes in the full sample. Panels A–B use school exposure (threshold: 0.5 miles); Panels C–D use neighborhood exposure (threshold: 0.3 miles). All specifications include student fixed effects and zipcode-by-semester fixed effects. The x-axis shows semesters relative to retailer opening; $\ell = -1$ is normalized to zero. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 21]

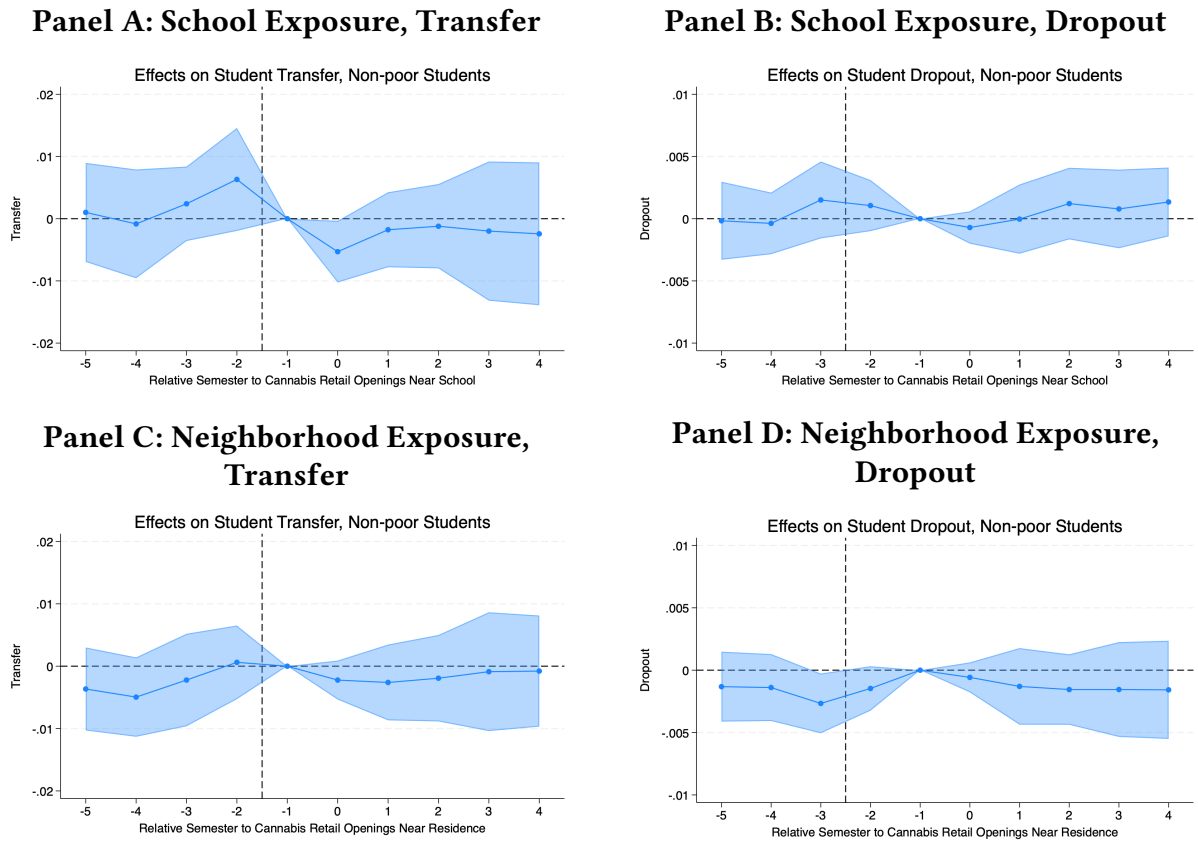


Figure A6 Addressing Selective Migration: Student Transfers and Dropouts (Non-poor Students)

Notes: This figure replicates Figure A5 for the non-poor subsample. Panels A–B use school exposure (threshold: 0.5 miles); Panels C–D use neighborhood exposure (threshold: 0.3 miles). Estimation and fixed effects are identical to the baseline migration specification. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 21]

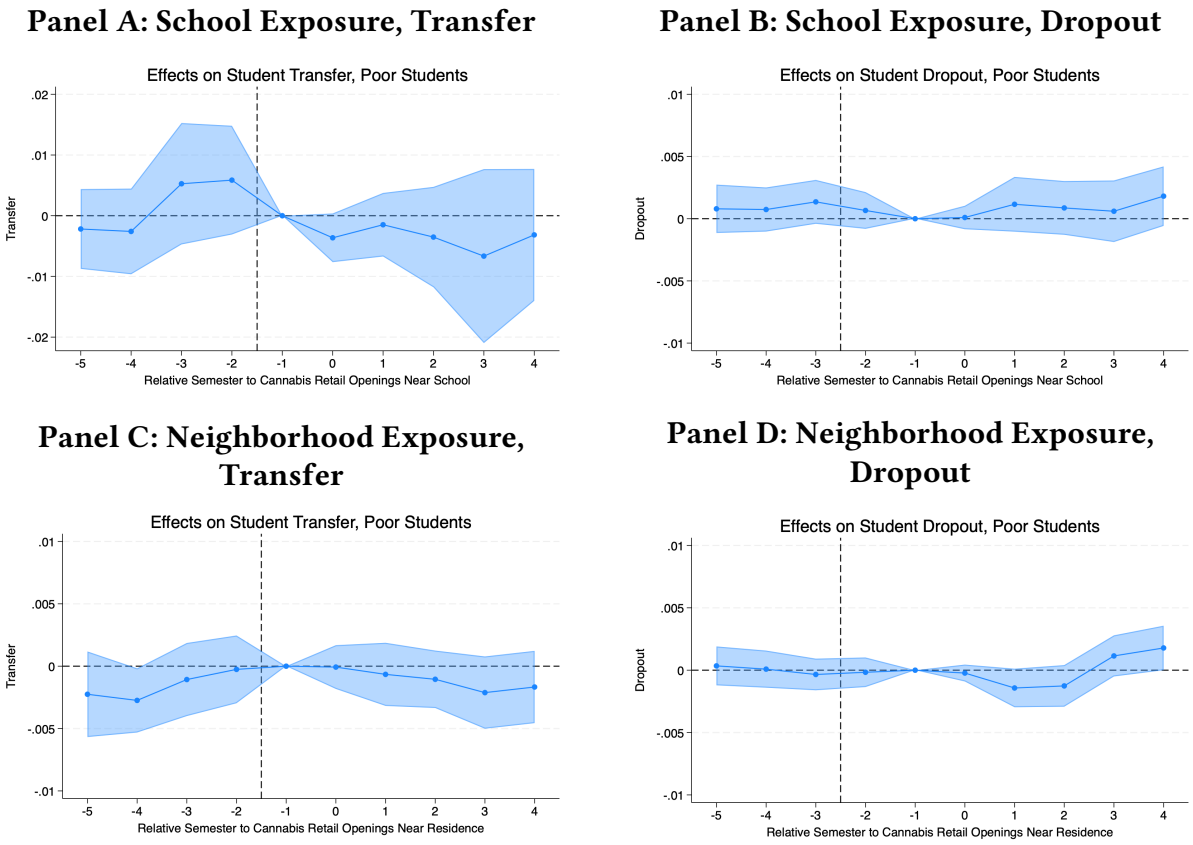
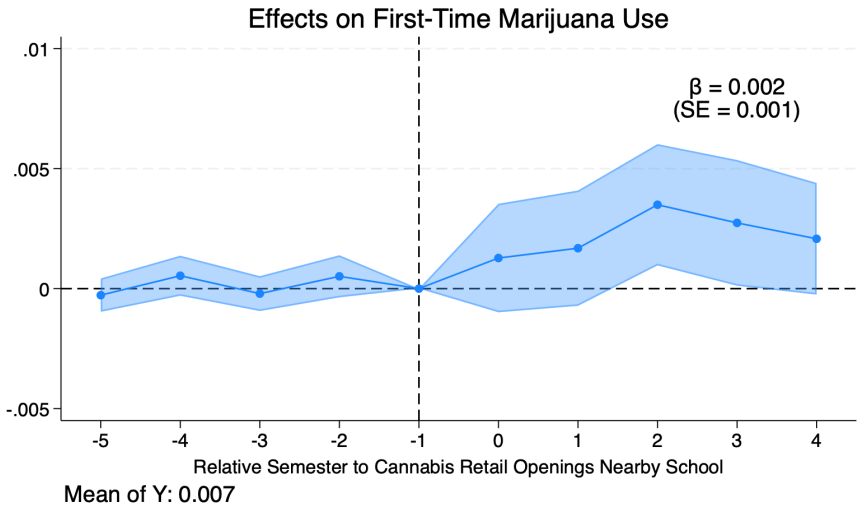


Figure A7 Addressing Selective Migration: Student Transfers and Dropouts (Poor Students)

Notes: This figure replicates Figure A5 for the poor subsample. Panels A–B use school exposure (threshold: 0.5 miles); Panels C–D use neighborhood exposure (threshold: 0.3 miles). Estimation and fixed effects are identical to the baseline migration specification. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 21]

Panel A: School Exposure



Panel B: Neighborhood Exposure

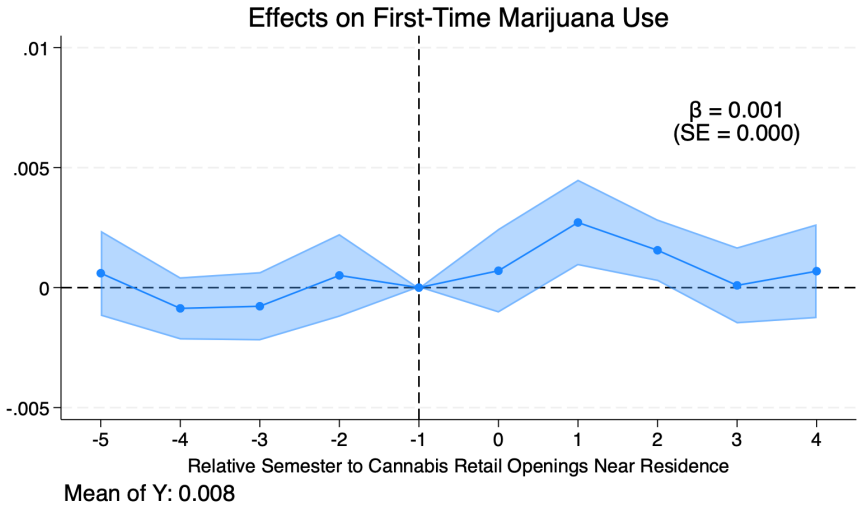
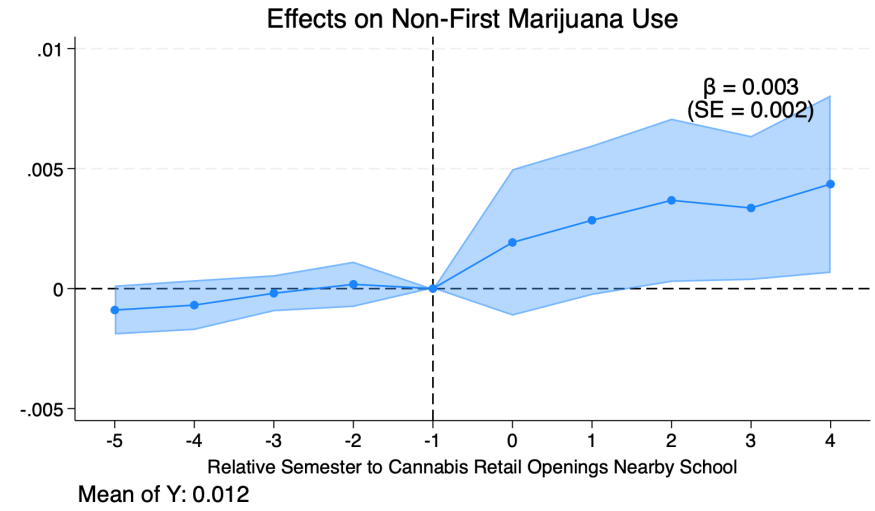


Figure A8 Event-Study Estimates: First-Time Marijuana Possession Suspensions

Notes: This figure shows event-study estimates for first-time marijuana possession offenses (possession of not more than 1 ounce). The specification, sample, and fixed effects are identical to Figure 3. Panel A shows school exposure (threshold: 0.5 miles); Panel B shows neighborhood exposure (threshold: 0.3 miles). The x-axis shows semesters relative to retailer opening; $\ell = -1$ is normalized to zero. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 23]

Panel A: School Exposure



Panel B: Neighborhood Exposure

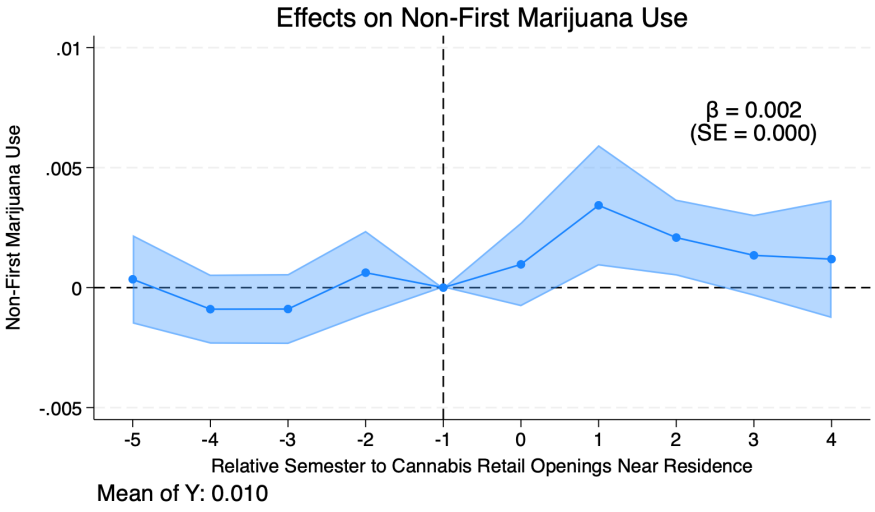
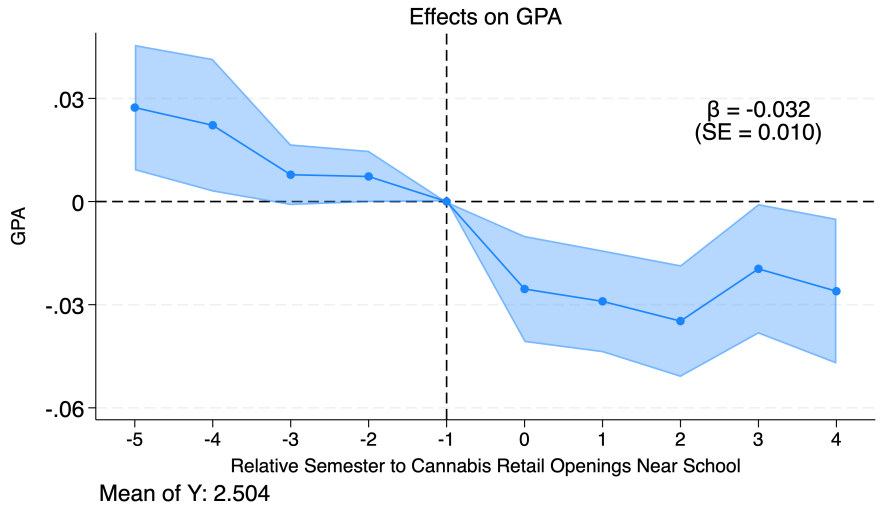


Figure A9 Event-Study Estimates: Non-First Marijuana Possession

Notes: This figure shows event-study estimates for non-first-marijuana drug-related suspensions, including controlled substance violations, prescription drug offenses, drug sales, and repeat marijuana offenses. The specification, sample, and fixed effects are identical to Figure 3. Panel A shows school exposure (threshold: 0.5 miles); Panel B shows neighborhood exposure (threshold: 0.3 miles). Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 23]

Panel A: School Exposure



Panel B: Neighborhood Exposure

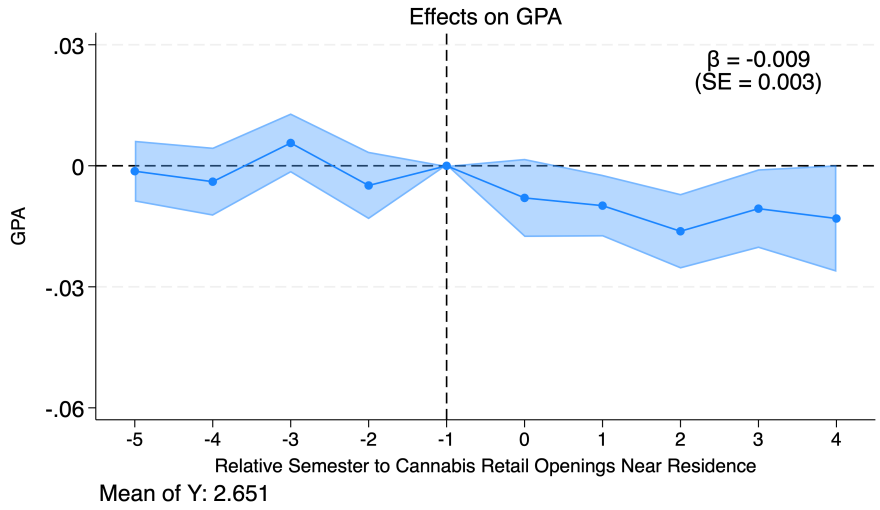


Figure A10 Robustness: Event-Study Estimates Using Walking Distance

Notes: This figure replicates the GPA event-study analysis from Figure 4 using actual walking distances calculated via OpenRouteService rather than geodesic (straight-line) distances. Panel A shows school exposure (threshold: 0.5 miles); Panel B shows neighborhood exposure (threshold: 0.3 miles). The specification includes student fixed effects and zipcode-by-semester fixed effects. The x-axis shows semesters relative to retailer opening; $\ell = -1$ is normalized to zero. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 25]

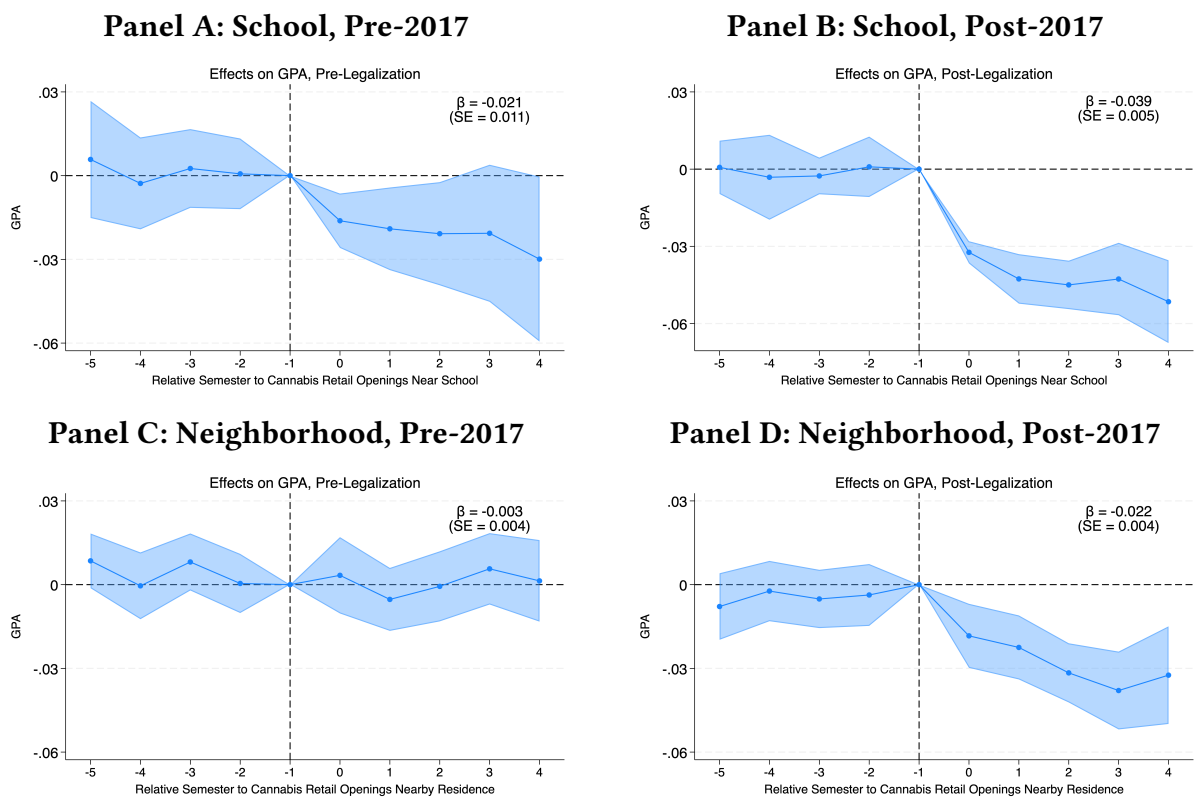
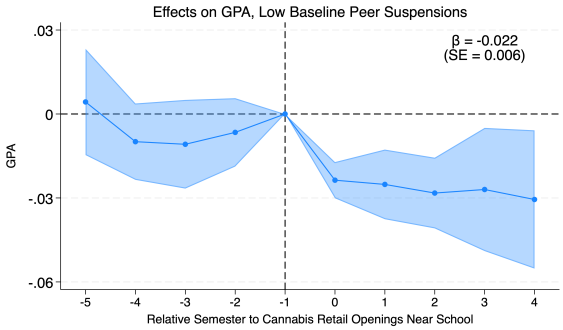


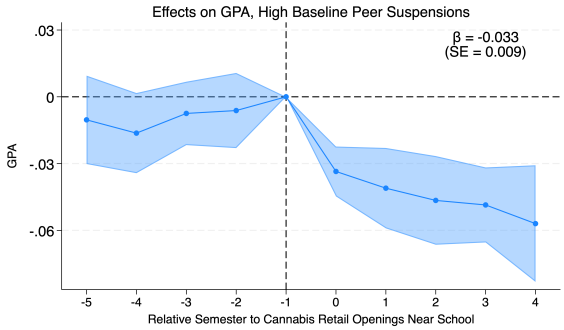
Figure A11 GPA Event Studies: Medical Era (Pre-2017) vs. Recreational Era (Post-2017)
Notes: This figure replicates the GPA event-study analysis from Figure 4, estimated separately by regulatory era. Panels A and C restrict to retailers opening during the medical-only era (pre-2017). Panels B and D restrict to retailers opening after recreational legalization (post-2017). Panels A–B show school exposure (threshold: 0.5 miles); Panels C–D show neighborhood exposure (threshold: 0.3 miles). All specifications include student fixed effects and zipcode-by-semester fixed effects. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 27]

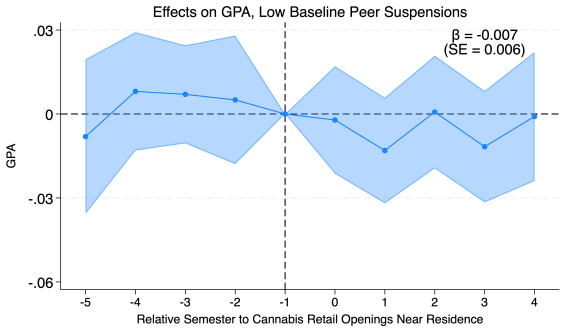
Panel A: School Exposure, Low Baseline Suspension



Panel B: School Exposure, High Baseline Suspension



Panel C: Neighborhood Exposure, Low Baseline Suspension



Panel D: Neighborhood Exposure, High Baseline Suspension

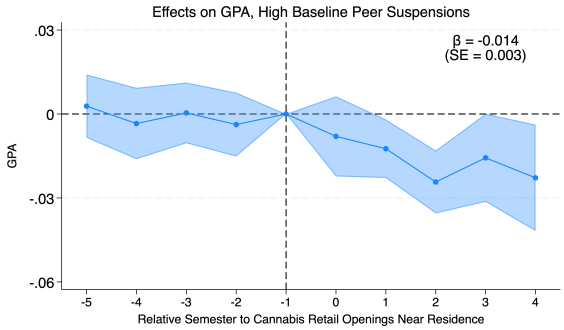
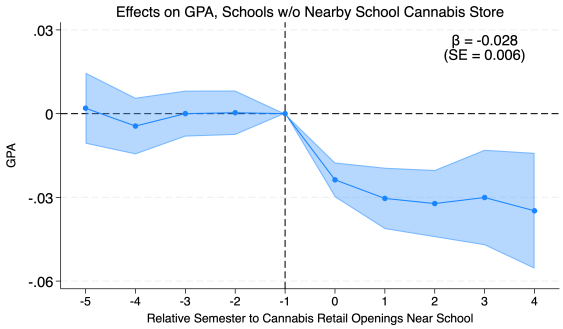


Figure A12 Peer Amplification: GPA Effects by Baseline School-Cohort Suspension Rate

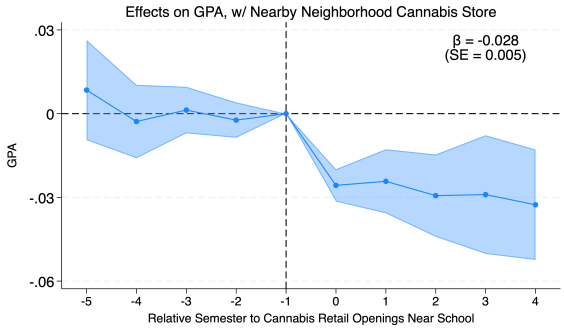
Notes: This figure examines whether peer environments amplify the effect of cannabis retailer proximity on GPA. Panels A and B report school exposure (threshold: 0.5 miles), split by baseline (pre-opening) drug suspension rates within school-by-cohort cells. Panels C and D report the analogous split for neighborhood exposure (threshold: 0.3 miles). "Low" and "High" indicate below- and above-median baseline suspension rates. The specification is identical to Figure 4. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 28]

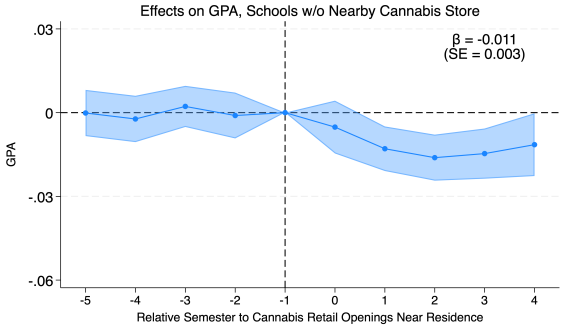
Panel A: School Exposure, No Retailer Near Residence



Panel B: School Exposure, Retailer Also Near Residence



Panel C: Neighborhood Exposure, No Retailer Near School



Panel D: Neighborhood Exposure, Retailer Also Near School

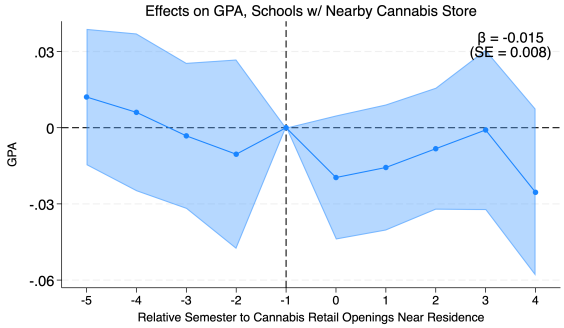
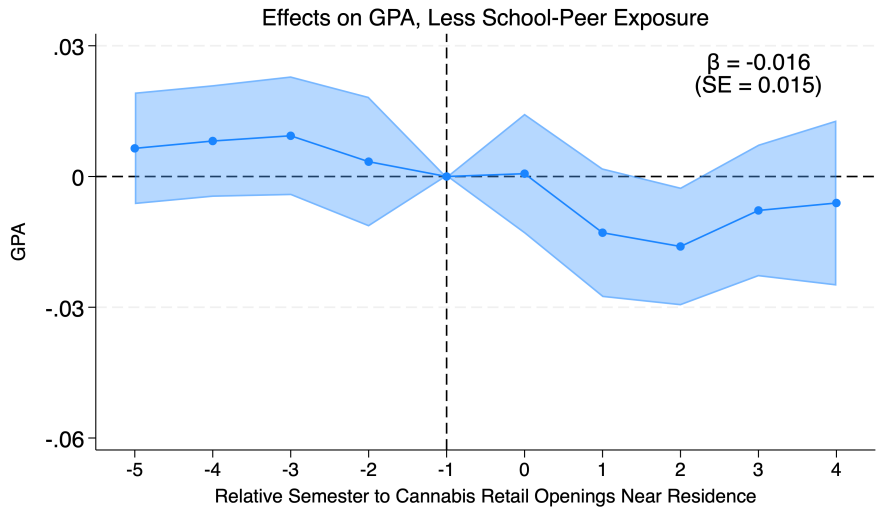


Figure A13 Interaction between School and Neighborhood Exposure Channels

Notes: This figure tests whether school and neighborhood exposure channels interact by splitting each sample according to the other channel. Panels A and B split the school-exposure sample (threshold: 0.5 miles) by whether the student's residence also has a retailer within 0.17 miles (bottom 10% of distance distribution). Panels C and D split the neighborhood-exposure sample (threshold: 0.3 miles) by whether the student's school also has a retailer within 0.44 miles (bottom 10% of distance distribution).

[Back to Page 29]

Panel A: Low School-Peer Exposure in ZIP



Panel B: High School-Peer Exposure in ZIP

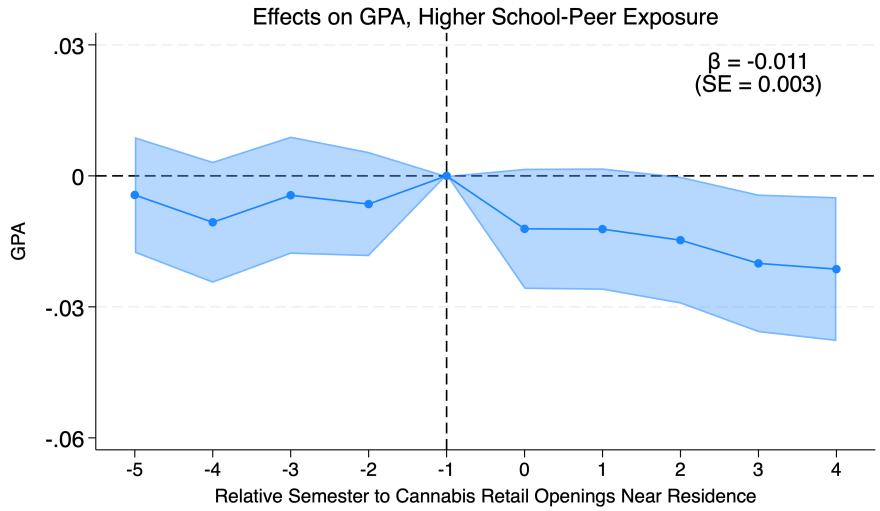
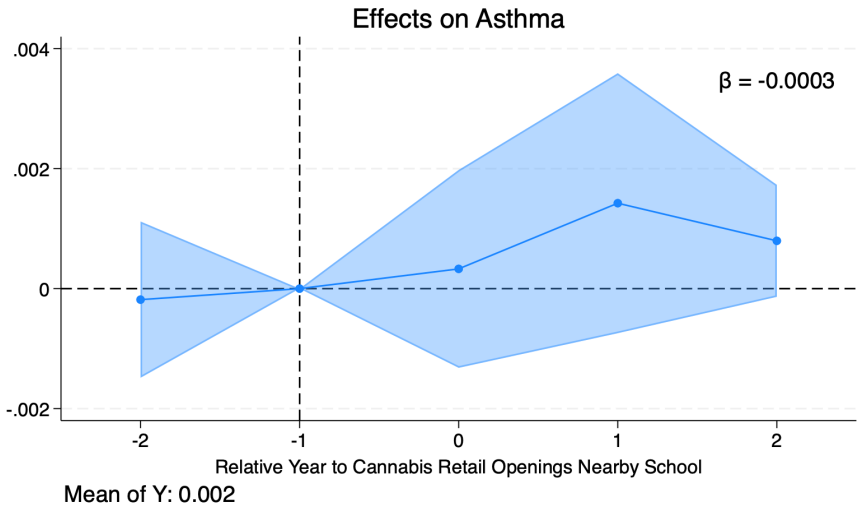


Figure A14 School-Cohort Peer Exposure and the Neighborhood Channel

Notes: This figure examines whether the neighborhood-channel GPA effect is mediated by the school-peer environment. The neighborhood-exposure sample (threshold: 0.3 miles) is split by the median level of school-channel exposure among students in the same school-cohort and ZIP code. Panel A reports event-study estimates for students whose school-cohort peers have below-median school-channel exposure; Panel B reports estimates for those with above-median school-channel exposure. School-cohort peer exposure is measured as the fraction of same-school-cohort peers in the student's ZIP code who attend a school within 0.5 miles of a cannabis retailer. The specification is identical to Figure 4, with student fixed effects and zipcode-by-semester fixed effects. The x-axis shows semesters relative to retailer opening; $\ell = -1$ is normalized to zero. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 30]

Panel A: School Exposure



Panel B: Neighborhood Exposure

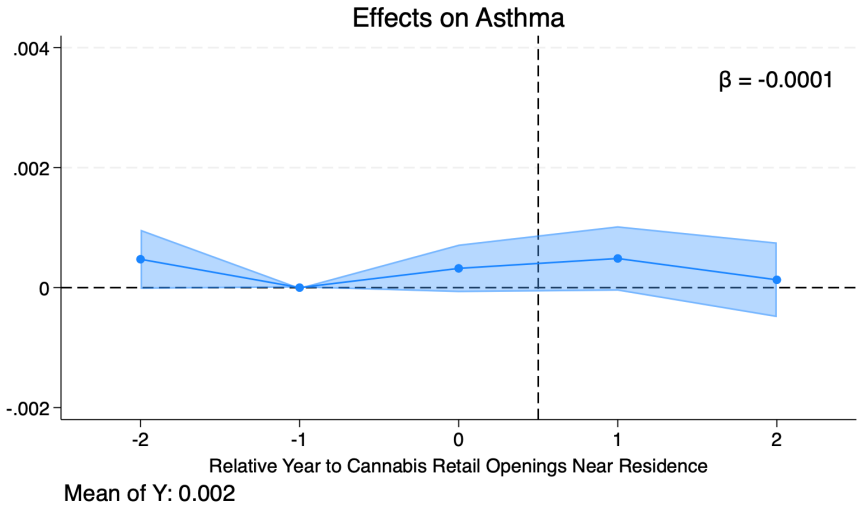
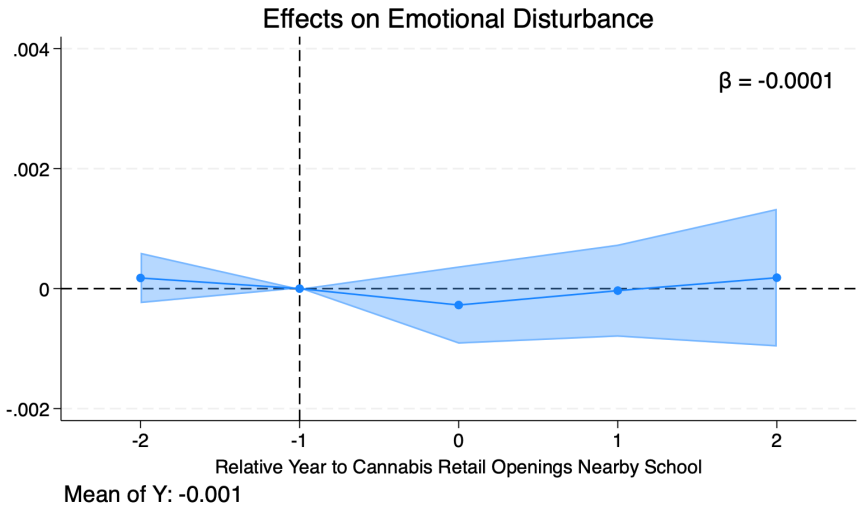


Figure A15 Event-Study Estimates: Effect of Retailer Opening on Student Asthma Diagnoses

Notes: This figure presents event-study estimates with an indicator for annual asthma diagnosis as the dependent variable. The x-axis shows years (not semesters) relative to the retailer opening, reflecting the annual frequency of health records. Panel A shows school exposure (threshold: 0.5 miles); Panel B shows neighborhood exposure (threshold: 0.3 miles). The specification includes student fixed effects and zipcode-by-year fixed effects. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 31]

Panel A: School Exposure



Panel B: Neighborhood Exposure

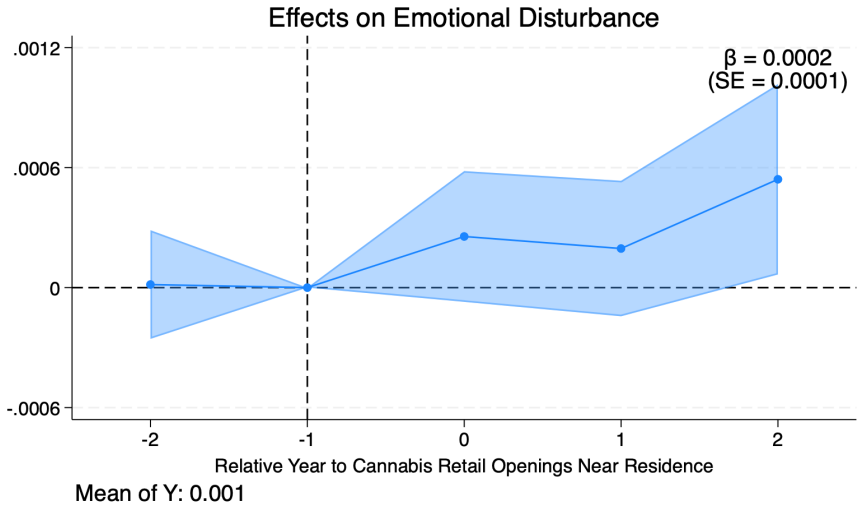


Figure A16 Event-Study Estimates: Effect of Retailer Opening on Emotional Disturbance Flags

Notes: This figure presents event-study estimates with an indicator for emotional disturbance designation (a school-recorded mental health flag) as the dependent variable. The x-axis shows years relative to the retailer opening. Panel A shows school exposure (threshold: 0.5 miles); Panel B shows neighborhood exposure (threshold: 0.3 miles). The specification includes student fixed effects and zipcode-by-year fixed effects. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 31]

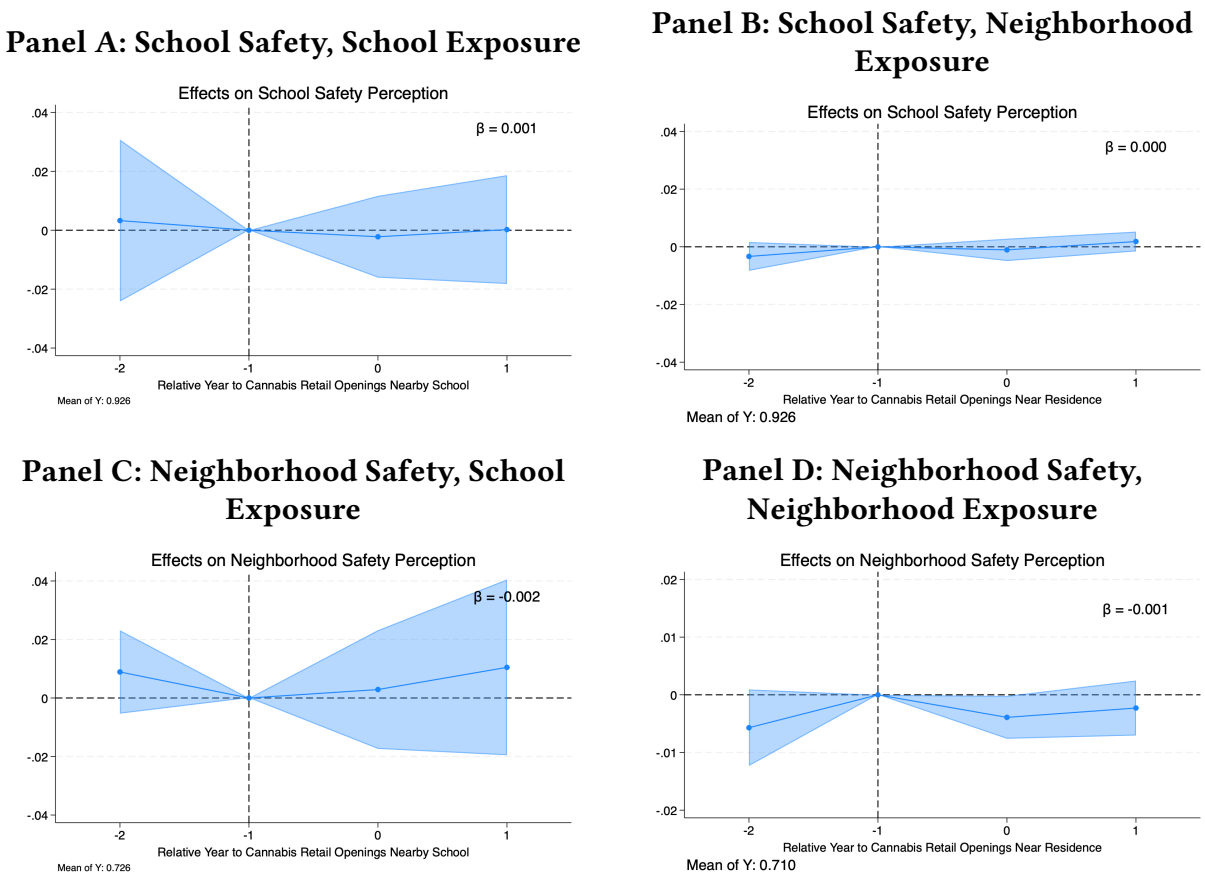
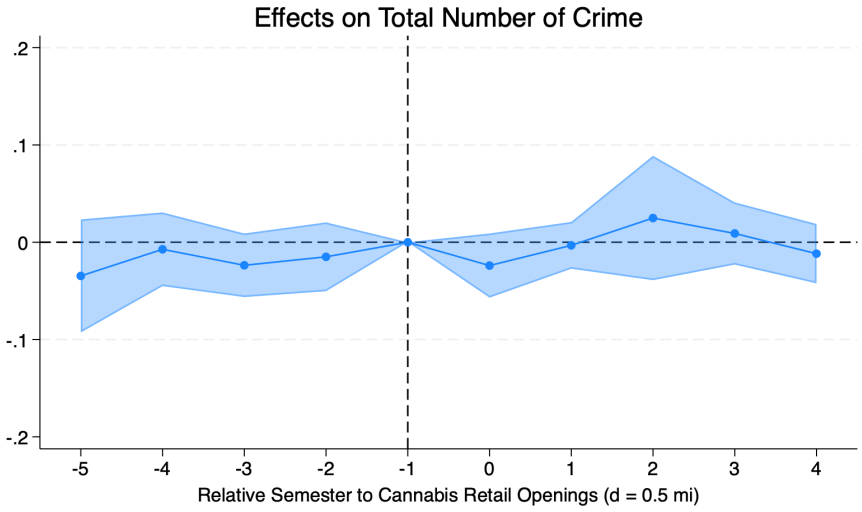


Figure A17 Event-Study Estimates: Effect of Retailer Opening on Safety Perceptions

Notes: This figure presents event-study estimates with perceived safety (from annual LAUSD student surveys) as the dependent variable. Higher values indicate greater perceived safety. Panels A and B report perceived school safety; Panels C and D report perceived neighborhood safety. Panels A and C show school exposure (threshold: 0.5 miles); Panels B and D show neighborhood exposure (threshold: 0.3 miles). The specification includes student fixed effects and zipcode-by-semester fixed effects. Shaded areas represent 90% confidence intervals based on standard errors clustered at the zipcode level.

[Back to Page 31]

Panel A: 0.5 Mile Treatment Threshold



Panel B: 0.3 Mile Treatment Threshold

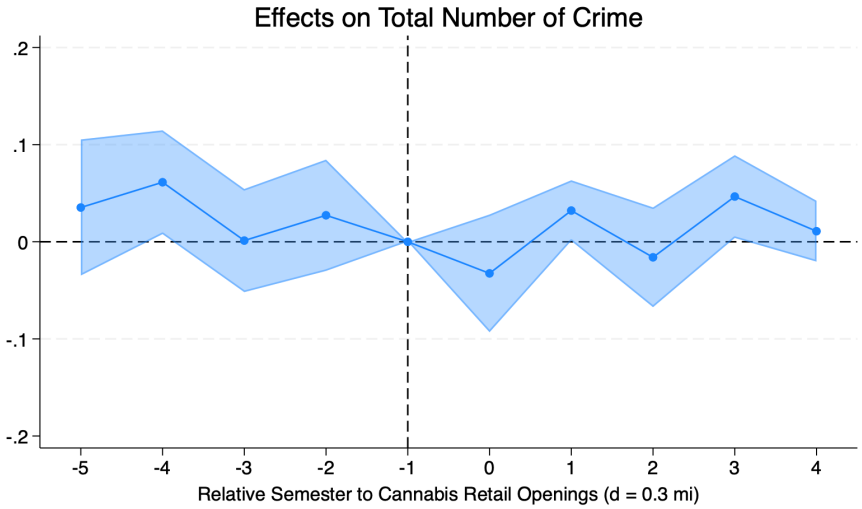


Figure A18 Event-Study Estimates: Effect of Retailer Opening on Total 911 Crime Calls

Notes: This figure presents event-study estimates with the total number of 911 crime calls as the dependent variable, aggregated to the semester-by-block level around cannabis retailer locations. Panel A defines treatment as blocks within 0.5 miles of a retailer opening (consistent with school exposure); Panel B uses a 0.3 mile threshold (consistent with neighborhood exposure). The x-axis shows semesters relative to the retailer opening; $\ell = -1$ is normalized to zero. The specification includes block fixed effects and semester fixed effects. Shaded areas represent 90% confidence intervals. Data are from LAPD 911 call records, 2009–2023.

[Back to Page 31]

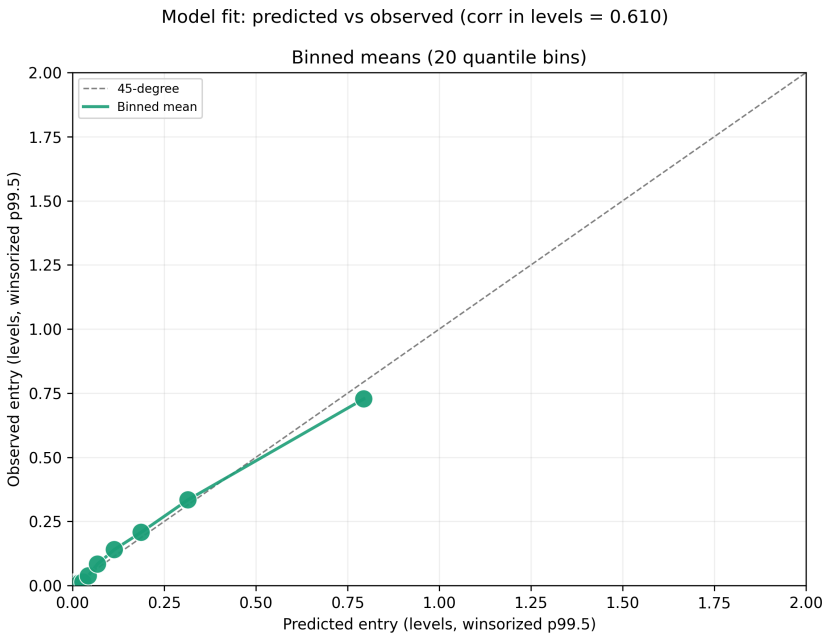


Figure A19 Model Fit: Predicted vs. Observed Entry at the Block-Semester Level

Notes: This figure plots predicted versus observed entry at the block-semester level. Predictions are from the full conditional-logit specification (Table 4, column 5). The figure displays binned means: observations are grouped into 20 quantile bins of predicted entry (winsorized at the 99.5th percentile), and each marker shows the mean predicted and observed entry within that bin. Marker size is proportional to the number of observations per bin. The dashed 45-degree line represents perfect prediction.

[Back to Page 39]

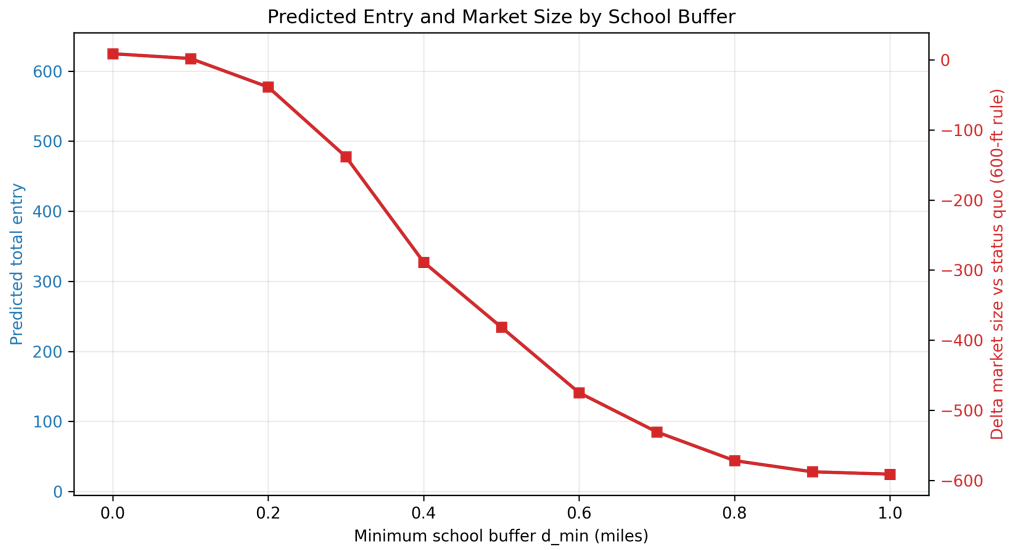


Figure A20 Predicted Total Entry and Market Size Change by School Buffer Distance

Notes: The left axis (blue line) plots total predicted entry across all census blocks for each buffer $d^{\min} \in \{0.0, 0.1, \ldots, 1.0\}$ miles. Entry is computed by re-normalizing block choice probabilities over feasible locations (those beyond the buffer from any school) within each tract-semester, then scaling tract-semester totals by the inclusive-value ratio using extensive-margin elasticity $\eta = 0.60$ (following Eaton, Kortum and Kramarz, 2004). The right axis (red line) plots the change in total entry $\Delta M(d^{\min})$ relative to the status quo ($d^{\min} \approx 0.114$ miles). Appendix Table A3 reports sensitivity to $\eta \in \{0.55, 0.60, 0.65\}$.

[Back to Page 42]

Appendix Tables

Table A1 Summary Statistics: Spatial Entry Model Estimation Sample

Variable	Mean	SD	Min	Max	N
New cannabis store entries	0.083	0.329	0.000	6.000	7,422
Incumbents within 0.25 mi	1.779	3.299	0.000	18.000	7,422
Incumbents 0.25–0.50 mi	3.612	4.939	0.000	21.000	7,422
Log broad retail spending	0.090	0.787	0.000	10.006	7,422
Log broad retail customers	0.040	0.375	0.000	6.488	7,422
Log median spend per customer	0.048	0.421	0.000	7.588	7,422
Log foot traffic visitors	1.605	3.523	0.000	15.387	7,422
Log foot traffic visits	1.641	3.595	0.000	15.452	7,422
Log POI points	1.232	1.971	0.000	7.011	7,422
Log existing cannabis stores	0.126	0.328	0.000	2.398	7,422
Log lagged existing cannabis stores	0.092	0.285	0.000	2.197	7,422
Distance to nearest school (mi)	0.388	0.270	0.014	4.199	7,422
Log population (block group)	7.171	0.802	0.000	8.377	7,422
Log per-capita income	10.394	0.527	8.372	12.484	7,422
Log median rent	7.435	0.289	6.213	8.161	7,422
Commercial land-use indicator	0.102	0.302	0.000	1.000	7,422
Log parking count (0.25 mi)	1.336	1.350	0.000	4.431	7,422

Notes: This table reports summary statistics for the spatial entry model estimation sample. The unit of observation is census block $\times$ semester. The sample covers all census blocks in the City of Los Angeles from 2019–2023 (semester-level panel), yielding 7,422 block-semester observations. "New cannabis store entries" is the number of retailers entering each block-semester. Competition variables (N_1, N_2) count active cannabis retailers within the specified distance band at the start of each semester. Demand shifters are derived from SafeGraph (credit-card spending at non-cannabis retail) and Advan (foot traffic). Log POI points is the log of the total count of all points of interest in the block-semester. Cost and demographic controls come from ACS 5-year block-group data (2019–2023) and OpenStreetMap land-use classifications.

[Back to Page 38]

Table A2 Competition Distance Attenuation Test

Test	Value
$\hat{\delta}_1$ (incumbents $\leq$ 0.25 mi)	0.0151
$\hat{\delta}_2$ (incumbents 0.25–0.50 mi)	0.0750
Wald test: $H_0 : \delta_1 = \delta_2$	
χ^2 statistic	2.109
p-value	0.1464
One-sided test: $H_1 : \lvert\delta_1\rvert > \lvert\delta_2\rvert$	
z-statistic	-1.452
p-value	0.9268

Notes: This table reports tests of distance-banded competition effects from the full conditional-logit specification (Table 4, column 5). $\hat{\delta}_1 = 0.0151$ is the estimated coefficient on incumbents within 0.25 miles; $\hat{\delta}_2 = 0.0750$ is the coefficient on incumbents 0.25–0.50 miles away. The Wald test evaluates equality ($H_0 : \delta_1 = \delta_2$). The one-sided test evaluates whether nearby competition exceeds distant competition in absolute magnitude ($H_1 : |\delta_1| > |\delta_2|$). Standard errors are clustered at the tract level.

[Back to Page 39]

Table A3 Extensive-Margin Elasticity Sensitivity ($\eta \pm 0.05$)

η	Optimal d^*	Net Welfare ($)	HC Gain ($)	Revenue Loss ($)
0.55	0.5	134,884,286	270,211,678	135,327,391
0.60	0.5	137,494,479	275,716,590	138,222,112
0.65	0.5	135,416,207	276,343,276	140,927,069

Notes: This table reports the optimal school buffer and associated welfare components under alternative values of the extensive-margin elasticity parameter η. The parameter η governs how changes in inclusive value (driven by buffer restrictions) translate into changes in total entry: $\eta = 0$ fixes total entry within each tract-semester, while $\eta = 1$ applies the full inclusive-value elasticity. The baseline calibration is $\eta = 0.60$, following Eaton, Kortum and Kramarz (2004), who estimate that roughly 60 percent of market-size variation transmits to firm entry on the extensive margin. Robustness checks use $\eta = 0.55$ and $\eta = 0.65$. Each row reports the optimal buffer d^* on the grid $\{0.0, 0.1, \ldots, 1.0\}$ miles and the associated net welfare, human-capital gain, and revenue loss (all in aggregate present-value dollars relative to the 600-foot status quo).

[Back to Page 42]

Table A4 Pigouvian Tax Entry Semi-Elasticity Sensitivity (ε)

ε	Pred. Entry	Δ Stores	HC Gain ($M)	Revenue Term ($M)	Net Welfare ($M)
0.33	529.0	-87.0	28.0	-191.6	219.6
0.50	484.2	-131.8	60.3	-150.3	210.5
0.67	439.4	-176.6	80.1	-108.9	189.1

Notes: This table reports sensitivity of the Pigouvian-tax counterfactual to the tax-entry semi-elasticity ε, which governs the percentage reduction in dispensary count per one-percentage-point increase in the effective tax rate. The baseline $\varepsilon = 0.50$ is calibrated from Hollenbeck and Uetake (2021); the bounds $\{0.33, 0.67\}$ span the range implied by Hansen, Miller and Weber (2017) and Hollenbeck and Uetake (2021). This parameter replaces the unidentified profit-index incidence parameter in favor of a transparent, externally anchored calibration analogous to the extensive-margin elasticity η. All values are in million USD relative to the current 600-foot rule. A negative revenue term denotes a fiscal gain.

[Back to Page 43]