Legal Segregation, Ethno-Religious Boundaries, and Urban Inequality in India's Disturbed Areas Act*

Shreya Dutt†

August 19, 2026

Abstract

This paper studies how legal control over ordinary property transfers can generate residential segregation and reshape urban opportunity. I examine Gujarat's Disturbed Areas Act (DAA), an active law requiring administrative approval for property transfers in Hindu–Muslim riot-affected areas and giving the state transaction-by-transaction discretion over local exchange. I construct a novel panel of DAA designations linked to polling-booth residential composition and pincode-level housing, school, and public-safety outcomes. To my knowledge, this provides the first causal evidence on how an ongoing legal restriction on ordinary property transactions reshapes religious segregation and urban amenities in India. Exploiting the Act's staggered expansion, I show that designation disproportionately targets Muslim-majority pockets embedded within Hindu-majority surroundings and hardens residential boundaries: Hindu presence falls by 4.8 percentage points, and treated areas contribute more to within-pincode Muslim isolation. Downstream effects depend on initial legal depth. In high-intensity pincodes, where multiple localities are covered at entry, housing prices and construction rise even as school infrastructure deteriorates and exam failure rates increase; in low-intensity pincodes, housing activity weakens. This divergence is not explained by stronger segregation in high-intensity areas. Instead, it is consistent with a critical-mass mechanism: larger, multi-locality DAA markets may sustain a segmented housing submarket in which constrained buyers continue searching locally and developers respond to concentrated demand. Communal violence does not decline robustly after designation. The results show how a law intended to protect minorities can harden group boundaries, segment housing markets, and decouple rising property values from improvements in local public goods.

Keywords: Segregation, Housing, Neighborhood Amenities, Urban India

JEL Codes: J15, R23, O18, K11, R31, H41

*I am deeply grateful to Martin Fiszbein, David Lagakos, Benjamin Marx, and Dilip Mookherjee for their guidance and support. I thank faculty and graduate students at Boston University, as well as seminar and conference participants at the 4th Political Economy Summer School of the Association for Comparative Economic Studies, the CSWEP Mentoring Workshop for Economics PhD Students, the 10th Urbanization and Development Conference at the World Bank, the 5th CISEPS Annual Workshop at the University of Milano-Bicocca, and the 4th Annual AEA CSQIEP Mentoring Conference, for helpful comments and suggestions. I gratefully acknowledge financial support from the Weiss Fund for Research in Development Economics and the Institute for Humane Studies. All errors are my own.

†Department of Economics, Boston University. 270 Bay State Road, Boston, MA 02215. Corresponding author email: dshrey@bu.edu.

1 Introduction

Who can live where is shaped not only by preferences and prices, but also by law. Neighborhoods determine access to schools, public services, safety, and social networks, with lasting consequences for economic mobility and well-being (Chetty et al., 2014, 2016; Chetty and Hendren, 2018a,b; Chetty et al., 2026). A large literature explains residential segregation through decentralized sorting over preferences, prices, amenities, and demographic composition (Schelling, 1971; Cutler et al., 1999; Bayer et al., 2007). Another studies the persistent effects of historical formal barriers such as segregation ordinances, redlining, racial covenants, zoning, and school boundaries (Troesken and Walsh, 2019; Aaronson et al., 2021; Fishback et al., 2024; Monarrez, 2023; Sood and Ehrman-Solberg, 2023; Almagro and Sood, 2025). Research on active restrictions is much thinner. Singapore's ethnic housing quotas show that limits on permissible counterparties can affect residential sorting and housing-market outcomes (Wong, 2013, 2014). Yet those restrictions are designed to preserve integration. We know much less about how active legal rules that limit residential exchange may instead harden group boundaries and reshape the distribution of urban opportunity while they remain in force.

Evidence from developing countries shows that state-created spatial institutions can have persistent consequences, but it comes from settings that differ from ongoing restrictions on ordinary neighborhood transactions. Research on South Africa and Malaysia studies historical forced resettlement and the long-run effects of imposed racial or ethnic geography (Selod and Zenou, 2001; Abel, 2019; Carrillo et al., 2023; Kok and Lim, 2024). Work on Indonesia once again examines how historical state-directed settlement patterns shape intergroup contact and national integration (Bazzi et al., 2019). Studies of China and Kenya show how administrative status, informal property rights, and local discretion affect housing and access to urban opportunity (Chen et al., 2015; Liao and Zhang, 2021; Marx et al., 2019). Related work in Latin America examines socioeconomic segregation and public-housing policy rather than group-specific legal barriers (Larrañaga Jiménez and Sanhueza Riveros, 2007; Lambiri and Vargas, 2011; García-López and Moreno-Monroy, 2018). India is a particu-

larly important setting because it combines a vast urban system with multiple salient social cleavages, including caste and religion. Recent research documents substantial caste and religious segregation, unequal access to local public services, and religious separation in rural areas (Bharathi et al., 2019, 2021, 2022; Asher et al., 2026; Arora, 2024). Yet this evidence is largely descriptive or cross-sectional and provides limited guidance on how an active legal institution changes neighborhood sorting over time in a large democracy. This paper asks how an ongoing neighborhood-specific restriction on property transfers reshapes Hindu–Muslim residential separation and the distribution of urban opportunity.

I study this question through the Disturbed Areas Act (DAA), an active law originating in the western state of Gujarat in India in 1991. The DAA requires prior administrative approval before a property transfer can proceed in a notified police station jurisdiction in cities. The state justifies this requirement as protection against distress sales and transactions arising from Hindu-Muslim communal disorder. In practice, however, the approval process gives the state authority over whether a proposed buyer–seller transaction is permitted and can therefore restrict exchange across community boundaries. Because the Act emerged from Hindu–Muslim conflict, the analysis focuses on residential separation along this margin. The central finding of the paper is that the DAA hardens Hindu–Muslim residential boundaries, while its downstream market consequences depend on the initial depth of legal coverage. In pincodes with multiple DAA localities (defined as "high-intensity"), housing prices and development strengthen even as school infrastructure and learning outcomes deteriorate; in pincodes with one DAA locality (defined as "low-intensity"), housing activity weakens instead. Because the residential response is not systematically stronger in high-intensity areas, this divergence appears to reflect differences in how segmented housing markets adjust rather than a stronger segregation first stage. In particular, the joint rise in housing prices and quantities in high-intensity areas is consistent with stronger local demand accompanied by a developer response, although the data do not separately identify shifts in housing demand and supply. Communal violence, meanwhile, does not decline in a broad or robust way.

To measure the law's spatial expansion, I construct a novel locality-year panel of DAA designations. I collect official notifications through fieldwork at the Gujarat Revenue Department and supplement them with newspaper records used to identify additional candidate designations. I standardize locality names and cross-validate designation dates and boundaries against the official legal record. I then geocode each notified locality and map it to two nested spatial units: polling booths, electoral catchment areas which capture fine religious composition, and pincodes, postal zones which approximate local housing markets. To measure Hindu–Muslim composition, I combine polling-booth aggregates from the 2014 electoral rolls (Susewind, 2016), parsed individual records from the 2017 rolls (Sood et al., 2018), and newly processed Election Commission electoral-roll PDFs from 2025. These data allow me to observe both the expansion of the legal regime and changes in religious composition within narrowly defined neighborhoods.

I next link legal exposure to three sources of data on downstream neighborhood outcomes. This design provides causal evidence on how an ongoing legal regime changes both residential sorting and urban amenities. Housing project-level records from the Gujarat Real Estate Regulatory Authority (RERA) measure registered housing development, unit counts, and prices from 2005 to 2024. I aggregate these outcomes to the pincode-year level. Annual school records at the pincode level come from the Unified District Information System for Education Plus (UDISE+) from 2005 to 2022 that provide school-level measures of infrastructure, resources, institutional status, and student performance. Finally, geocoded events from Global Database of Events, Language, and Tone (GDELT) between 1980 and 2022 provide local measures of public safety such as communal violence and civilian targeting. Taken together, these linkages create a spatially granular, multi-year panel that follows how an active legal regime affects residential composition and margins of local urban opportunity as measured by housing activity, school environments, and public safety.

The empirical strategy uses the staggered expansion of the DAA to compare changes in newly designated areas with changes in areas that have not yet been designated or are never designated. Treatment is defined along both the extensive margin and the initial intensive

margin. A polling booth or pincode becomes treated in the first year that its pincode contains any DAA-notified locality. Baseline DAA intensity is measured at the pincode level as the number of distinct notified localities covered when the pincode first enters the DAA regime. Treated pincodes are divided at the median, which is one DAA locality, of this baseline distribution for heterogeneity analyses, and polling booths inherit the intensity classification of their pincode. High-intensity pincodes are pincodes with multiple DAA localities and low-intensity pincodes are pincodes with exactly one DAA locality. The baseline measure is fixed thereafter, separating initial differences in the depth of legal exposure from subsequent, potentially endogenous expansions of DAA coverage. For annually observed outcomes, I estimate dynamic treatment effects separately by designation cohort and event time, allowing the effects of the law to vary across places and over time (Sun and Abraham, 2021; Callaway and Sant'Anna, 2021). For neighborhood composition, which is observed in 2014, 2017, and 2025, I use stacked difference-in-differences designs that compare changes across electoral-roll intervals around designation. All specifications absorb permanent differences across spatial units and common changes over time. The identifying assumption is that, absent designation, treated and comparison areas would have continued on similar trends. Because nearby untreated areas may themselves adjust when the law is imposed, the main specifications are tested for robustness by excluding potentially exposed neighbors from the comparison group. Complementary spillover specifications treat proximity to an active DAA area as the exposure of interest.

I begin by examining where the DAA is likely to be imposed. Muslim-majority booths located within otherwise Hindu-majority pincodes are 13.7 percentage points more likely to be designated, relative to a mean of 38.6 percent. These patterns indicate that the Act disproportionately enters boundary areas where cross-community residence and property transactions remain active margins. In separate event-study analyses, GDELT measures of communal violence and civilian targeting do not increase systematically before designation. The geography of designation is therefore more closely associated with religious composition than with a visible pre-designation escalation in violence, complicating a narrow interpretation of the Act as a response to recent communal disorder.

The first main result is that designation causes religious boundaries to harden within neighborhoods. In the stacked specification, DAA exposure reduces the Hindu share of treated polling booths by 4.8 percentage points from a baseline mean of 86.6 percent. The Muslim share changes little and is imprecisely estimated. The log Muslim-to-Hindu ratio rises by 0.849 log points, equivalent to an increase of approximately 134 percent relative to the comparison group. Treated booths' contribution to within-pincode Muslim isolation increases by 0.002, relative to a baseline mean of 0.004. The central pattern is therefore a decline in Hindu presence while the remaining Muslim population becomes more concentrated within particular parts of the pincode. Nearby untreated booths within five kilometers show a 2.8 percentage-point increase in Muslim share, but their Muslim-to-Hindu ratio changes little and is imprecisely estimated. The clearest evidence of boundary-hardening therefore remains concentrated within formally treated areas.

The housing-market effects depend sharply on the initial depth of legal exposure. In high-intensity pincodes, which contain multiple notified localities at DAA entry, prices rise gradually after designation and reach approximately Rs. 22,000–25,000 more per square foot eight to ten years out, while unit density rises by about 100–150 units per square kilometer. In low-intensity pincodes, which contain one notified locality at entry, prices remain flat or decline and unit density falls by roughly 50–100 units per square kilometer. Importantly, the residential-composition estimates do not show an analogous intensity gradient: the incremental high-intensity effects on Muslim share, Hindu share, the Muslim-to-Hindu ratio, and isolation are imprecisely estimated. The housing divergence therefore does not appear to arise from substantially stronger boundary hardening in high-intensity places.

The school results show that stronger housing-market activity does not translate into broader improvements in neighborhood amenities. The school-infrastructure index falls by approximately 0.05–0.06 standard deviations from three to ten years relative to designation under DAA. The decline is concentrated in high-intensity pincodes, where later estimates range from approximately 0.03 to 0.07 standard deviations below the comparison group. Primary-school failure rates rise by 1–2 percentage points within five to seven years of designation

and by 3–4 percentage points in later years, with similar increases in secondary failure rates. These changes are concentrated in high-intensity pincodes. Grants received fall by as much as Rs. 140–160 per student, grants spent fall by Rs. 110–135 per student, and the private-school share declines by approximately 3–4.5 percentage points in later years. Taken together, the estimates indicate weaker infrastructure, reduced funding, and worse student performance in high-intensity pincodes. The contrast with the housing results is important: the same places in which housing prices and development strengthen experience deterioration rather than improvement in local school environments. Price growth therefore does not appear to reflect broad-based improvement in neighborhood amenities.

Finally, I find little evidence that designation produces a broad or statistically precise improvement in public safety. Pooled GDELT estimates for communal violence are negative after designation, but the confidence intervals are wide and generally include zero. In high-intensity pincodes, the point estimates become more negative after treatment, reaching approximately three to five fewer communal events per pincode-year in the later event years. Estimates for low-intensity pincodes remain closer to zero. Moreover, pre-treatment coefficients are broadly flat, providing little evidence that designation systematically follows a visible escalation in communal violence as stated in the law.

The joint pattern is consistent with the DAA acting as a state-run sorting mechanism segmenting housing markets. The approval rule raises the cost and uncertainty of cross-community exchange, and the residential results show that this friction hardens Hindu–Muslim boundaries broadly. Yet these compositional responses are not systematically stronger in high-intensity areas, suggesting that the housing divergence operates through a distinct market adjustment. One interpretation is a critical-mass mechanism. Low-intensity pincodes with a single DAA locality may be too small to sustain a substantial alternative housing market once transactions become harder, causing affected buyers to postpone purchases or search elsewhere. When several localities are covered within a larger residential market, however, they may form a sufficiently large segmented submarket for households whose neighborhood choices have become constrained to continue finding feasible sellers, projects, and locations

locally. This does not mean that greater DAA exposure makes transactions easier in absolute terms; rather, conditional on cross-community constraints, a larger segmented market may offer more feasible alternatives. A sufficiently large pool of buyers can in turn make new development profitable. The simultaneous rise in prices and housing quantities in high-intensity areas is consistent with such a demand-and-development response, rather than a pure contraction in housing supply. Because I observe equilibrium prices and quantities, I cannot identify demand and supply shifts separately; nor do I observe buyer–seller religion or DAA approval decisions needed to directly test the proposed transaction channel. But I do observe that high-intensity pincodes are larger electorates, including a larger number of Muslim registered voters, which I use as a proxy for local residential-market scale. This lends credence to the suggested mechanism. This also explains why rising property values need not indicate improving neighborhood amenities: under segmented housing-market access, prices can rise even as school quality deteriorates. Standard hedonic models interpret property prices as reflecting housing characteristics and local amenities, and existing evidence shows that school quality, school facilities, and urban infrastructure are capitalized into housing values (Rosen, 1974; Black, 1999; Cellini et al., 2010; Gonzalez-Navarro and Quintana-Domeque, 2016; Kuminoff et al., 2013). Under the DAA, however, the evidence points to a state-created market-segmentation logic in which identity-linked approval frictions restrict cross-neighborhood movements, so prices reflect scarcity within legally constrained market segments rather than local amenities alone.

The paper's central contribution is to provide, to my knowledge, the first causal evidence on how a contemporary legal restriction on property exchange reshapes religious segregation and urban amenities in India, the world's largest democracy. The institution is distinctive because it does not formally prohibit members of a religious group from entering particular neighborhoods. Instead, it operates through transaction-by-transaction administrative discretion over ordinary property exchange. By tracing the staggered expansion of this regime over space and time, the paper moves beyond measuring existing segregation or documenting the persistent legacy of historical restrictions. It shows how legal segregation can be produced dynamically through procedural control over who is able to transact across com-

munity boundaries. The findings therefore identify a contemporary institutional source of Hindu–Muslim residential separation and show that its effects extend beyond composition to housing markets, schools, and public safety.

Second, the paper contributes to research on housing-market sorting, constrained neighborhood choice, and the capitalization of local amenities. The hedonic and equilibrium-sorting literature uses housing prices to recover households' valuation of schools, infrastructure, environmental quality, and other place-based attributes (Rosen, 1974; Black, 1999; Cellini et al., 2010; Kuminoff et al., 2013). Related work shows that discrimination and steering can restrict the neighborhoods households are effectively able to choose (Christensen and Timmins, 2022). This paper connects these ideas by showing how state scrutiny of individual transactions can create group-specific constraints on housing-market access. Under such a regime, prices need not reflect amenities and conventional supply constraints alone; they can also reflect equilibrium demand within a segmented market in which households face a restricted set of feasible transactions. The joint rise in housing prices and new construction in multiple-locality DAA markets is consistent with a sufficiently large segmented buyer market inducing a developer response, even as school infrastructure and learning outcomes deteriorate. More generally, when legal rules constrain neighborhood choice, rising property values need not signal improving local opportunity.

Third, the paper contributes to the political economy of social divisions, local public goods, and intergroup contact. Existing work examines how ethnic heterogeneity and political organization affect the provision and allocation of schools, infrastructure, and other public goods (Alesina et al., 1999; Miguel and Gugerty, 2005; Banerjee and Somanathan, 2007). A related literature shows that sustained and collaborative contact can reduce discrimination and strengthen cross-group social and economic ties (Rao, 2019; Lowe, 2021). This paper adds an institutional margin to both literatures: the state can endogenously alter local composition and opportunities for contact through the regulation of property exchange. The combination of sharper residential separation and weaker school environments suggests that legal segregation may change not only who lives near whom, but also the public-good en-

vironment available within the resulting neighborhoods. At the same time, the absence of a robust decline in Hindu–Muslim violence provides little evidence that DAA designation generates measurable public-safety benefits despite the accompanying increase in residential separation. The paper demonstrates that local religious composition and local urban amenities may evolve together following the imposition of an ordinary law governing housing market transactions.

The remainder of the paper proceeds as follows. Section 2 describes the Disturbed Areas Act and develops the conceptual framework. Section 3 presents the treatment panel and linked neighborhood data, and Section 4 sets out the empirical strategy. Sections 5–7 present the results on designation, residential composition, and neighborhood outcomes, and Section 8 concludes.

2 Institutional Background: The Disturbed Areas Act

The Disturbed Areas Act emerged from Gujarat's history of Hindu–Muslim violence and displacement, especially in Ahmedabad, where major episodes in 1969, 1985, 1991, and 2002 pushed Muslims from mixed or Hindu-dominated neighborhoods into safer Muslim-majority areas (Laliwala et al., 2021; Menon and Joji, 2023; Tejani, 2023). First introduced as a temporary law in 1986 and made permanent in 1991, the Act was formally designed to prevent distress sales, coerced transfers, and tenant displacement in riot-affected localities (Government of Gujarat, 1991; Shrivastava, 2022; Tejani, 2023). It operates by requiring prior district-collector approval for property transfers in notified areas, with officials assessing consent and fair value; unapproved transfers can be treated as void (Government of Gujarat, 1991; Shrivastava, 2022). Although not written as a religion-specific ban, the law is most salient for Hindu–Muslim transactions in segregated areas or near religious boundaries, where administrative scrutiny, delay, denial risk, and political pressure raise the cost of cross-community exchange (Tejani, 2023; Shrivastava, 2022; Menon and Joji, 2023). Later amendments strengthened this regulatory role: a 2009 amendment expanded the collector's

authority, while a 2020 amendment introduced concerns about polarization, demographic equilibrium, and improper clustering, shifting the law from past violence toward prospective demographic control (Government of Gujarat, 1991, 2020; Tejani, 2023; Shrivastava, 2022). Although the Gujarat High Court stayed fresh notifications under the amended provisions in 2021, the broader law remained in force (PRS Legislative Research, 2022); for this paper, the key point is that the DAA evolved from a protective response to riot-affected space into an ongoing instrument of neighborhood regulation (Tejani, 2023; Laliwala et al., 2021).

I therefore treat the Act as a legal friction on property transfers across community boundaries. By increasing delay, uncertainty, scrutiny, and the risk of denial, it can narrow the set of transactions households can readily complete and thereby potentially increase residential segregation along religious lines. Its housing-market consequences, however, are theoretically ambiguous. Where only a small local pocket is covered, transaction friction may primarily deter exchange and lead households to postpone purchases or search elsewhere. Where several localities are covered within a sufficiently large residential market, the restriction may instead help create a larger segmented submarket in which constrained households can still find feasible sellers, projects, and locations locally. Such a market may retain enough buyer demand to induce a developer response. Treatment intensity may therefore generate different housing-market responses even if the underlying segregation effect does not increase monotonically with intensity. The law may also affect public goods and public safety as residential composition and local economic activity adjust.

3 Data

3.1 DAA Notifications and Treatment Construction

A central contribution of the paper is a new dataset tracing the expansion of the Disturbed Areas Act across urban Gujarat over time. The Act's operative legal geography is the notified disturbed area: a specific locality, neighborhood, society, street, village, ward, or

set of localities within a police-station jurisdiction[1] named in a government notification. Because no complete public archive of notified areas exists, I combine Gujarat Revenue Department notifications, collected through field visits and legal contacts, with newspaper reporting on new designations and disputes related to the law. The official notifications form the backbone of the dataset: they identify newly covered localities within police-station jurisdictions but are often scanned, non-standardized, or semi-structured, so I standardize them into a unified panel of notified localities, police-station jurisdictions, cities, and years of first notification. Newspaper sources supplement incomplete official records and provide an independent source that can be cross-checked against the legal record. I process these articles in two stages: first extracting whether an article mentions a disturbed area and the relevant locality, and then recovering the month and year of designation when available, or the year otherwise. I reconcile repeated locality references across articles and merge the resulting candidate locality-year records with the official notification database. Online Appendix A.1 describes the extraction procedure in the newspaper pipeline. I then geocode disturbed localities, recover candidate pincodes through reverse-geocoding, and translate this legal geography into two empirical units: polling booths, which approximate fine neighborhoods and are used for residential composition, and pincodes, which are broader postal areas used for housing, school, and violence outcomes[2]. A unit is coded as treated once it contains any locality within a notified disturbed area, and its treatment year is the earliest year in which any part of the unit is brought under the Act.

To capture variation in the local reach of the law, I also construct a time-varying intensity

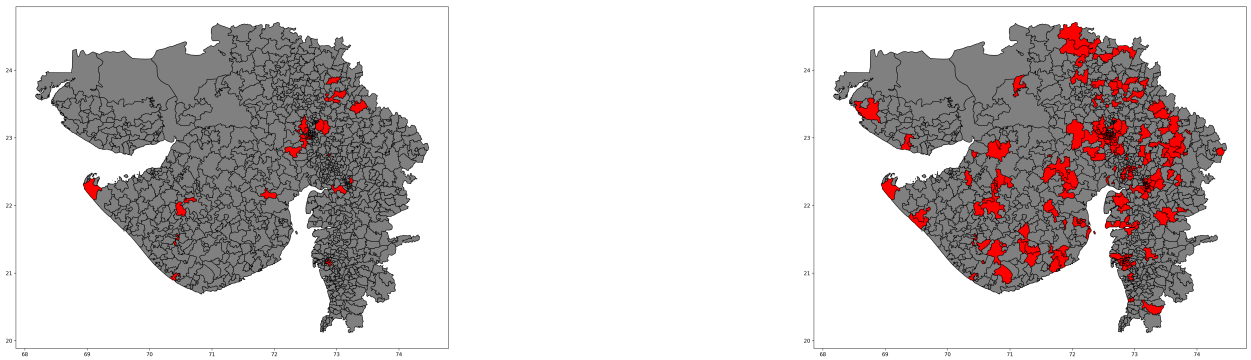
[1] In Gujarat, the India Justice Report reports roughly 77,000 people per rural police station and 283,000 people per urban police station, using National Commission on Population data and BPR&D police-station counts (India Justice Report, 2025a,b,c). These figures should be interpreted as approximate scale benchmarks, since police-station jurisdictions vary across rural and urban areas.

[2] For scale, polling booths in Gujarat are very small electoral units: Election Commission of India norms state that a polling station should generally not exceed 1,500 electors and should ordinarily be located so that voters do not travel more than two kilometers to vote (Election Commission of India, 2024b). In the 2024 electoral roll, Gujarat had 50,677 polling stations and 49.48 million electors, implying about 976 electors per polling station (Election Commission of India, 2024a). Pincodes are larger postal units: using Gujarat's 12.18 million Census households and the Department of Posts pincode directory, a rough statewide benchmark is about 12,000 households per pincode (Office of the Registrar General and Census Commissioner, India, 2011a; Department of Posts, Government of India, 2026). Because pincodes are postal rather than Census units, this figure should be interpreted as an approximate scale benchmark rather than an official administrative average.

measure based on the number of distinct disturbed localities mapped into a pincode in each year. This distinguishes lightly exposed from heavily exposed places and is important because average treatment effects may mask substantial heterogeneity in the depth of legal exposure. The final treatment panel is, to my knowledge, the first systematic dataset on the spread of the Disturbed Areas Act across Gujarat over time. It combines legal notifications, field-collected archival material, and newspaper text to recover where and when the law was applied, and it translates these legal designations into neighborhood units that can be matched to electoral rolls, school records, housing data, and violence data.

The dataset shows that the DAA expanded both extensively and intensively across urban Gujarat. Figure 1 shows the extensive margin: between 2002 and 2024, the number of treated pincodes rises from 35 to 258, indicating that the law broadened beyond a narrow set of historically affected neighborhoods.

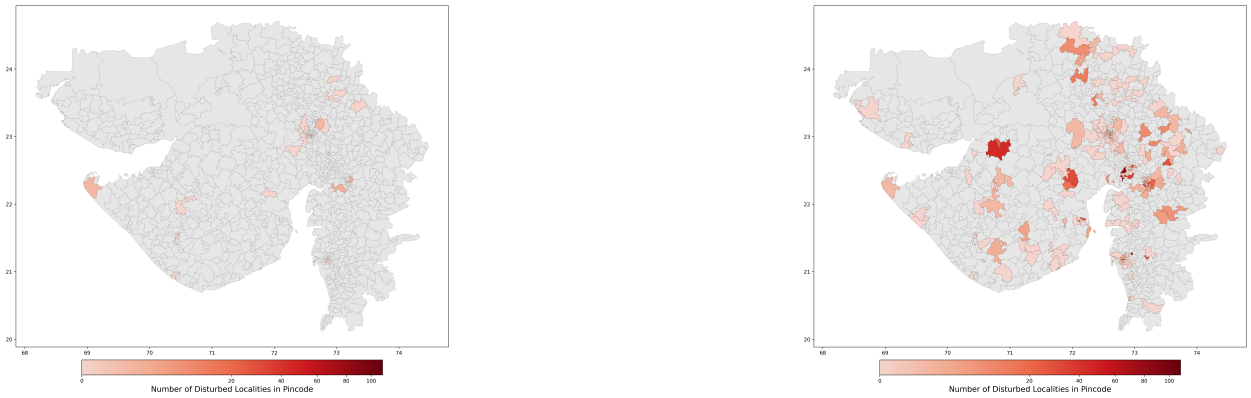
Figure 1: Expansion of DAA Exposure across Gujarat Pincodes



(a) 2002 ($N_{treated} = 35$) (b) 2024 ($N_{treated} = 258$)

Notes: The figure maps the extensive margin of DAA exposure across Gujarat at the pincode level. A pincode is coded as treated in a given year if it contains at least one locality notified under the Disturbed Areas Act. Red pincodes are treated; grey pincodes are not treated. Panel (a) shows the spatial distribution in 2002, when 35 pincodes were treated. Panel (b) shows the distribution in 2024, when 258 pincodes were treated. The figure illustrates the broadening spatial footprint of the DAA over time and its concentration in major urban areas.

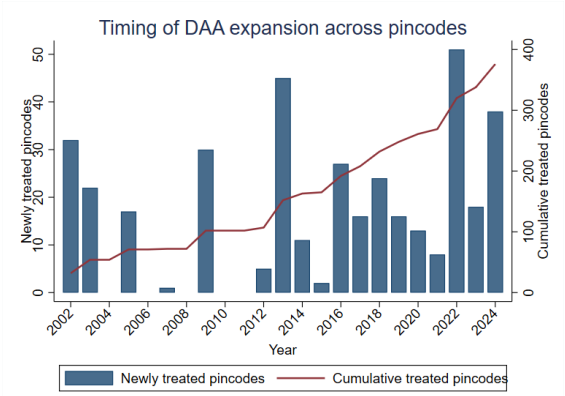
Figure 2: Intensity of DAA Exposure across Gujarat Pincodes



(a) 2002 (Max Intensity = 3) (b) 2024 (Max Intensity = 110)

Notes: The figure maps the intensive margin of DAA exposure across Gujarat at the pincode level. Intensity is measured as the number of distinct disturbed localities mapped into each pincode in a given year. Untreated pincodes are shown in light gray. Treated pincodes are shaded using a common color scale across both panels, so differences across years reflect changes in the local depth of exposure. Panel (a) shows that in 2002 the law had reached only a small number of pincodes and remained shallow within treated areas, with a maximum intensity of 3. Panel (b) shows that by 2024 the law had both spread across more pincodes and deepened substantially within already exposed areas, with some pincodes reaching an intensity of 110.

Figure 3: Staggered Expansion of DAA Exposure across Gujarat Pincodes



Notes: The figure shows the timing of entry into DAA treatment at the pincode level. A pincode is coded as treated in the first year in which it contains at least one locality notified under the Disturbed Areas Act. Bars report the number of newly treated pincodes in each year. The line reports the cumulative number of treated pincodes over time. The figure shows that DAA expansion occurred in waves rather than as a smooth secular trend, generating the staggered cohort variation used in the event-study analysis.

Figure 2 shows the intensive margin: already treated pincodes accumulated more notified localities over time, with some containing more than 100 disturbed localities by 2024. Exposure is strongly urban and concentrated in major cities, especially Ahmedabad and Vadodara. Figure 3 shows that expansion occurred in waves rather than as a smooth secular trend, generating the staggered cohort variation used in the event-study analysis.

3.2 Electoral Rolls and Religious Composition

Electoral rolls provide near-universal adult coverage, individual names, and polling booth locations, making them well suited for measuring neighborhood-level religious composition, although religion must be inferred from names because it is not directly recorded. I use 2014 as the first benchmark year because it is the earliest consistent, research-usable digitized source for Gujarat; these data come from Susewind (2016) and provide polling booth–level religion counts based on name classification. For 2017, I apply the supervised `Pranaam` classifier to individual voter names and polling-station identifiers from the Parsed Indian Electoral Rolls dataset, which reports nearly 98% out-of-sample accuracy for Muslim/Hindu prediction in Hindi and English (Sood et al., 2018; Chintalapati et al., 2024)[3]. For 2025, I construct a new dataset from the full universe of Gujarat electoral-roll PDFs published by the Election Commission of India: I extract individual records, process Gujarati-script names using optical character recognition, transliterate them into standardized Latin script, tokenize given names and surnames, and classify religious identity using a dictionary built from external Gujarati-name sources[4], high-frequency surname fragments, and community-specific naming conventions. Online Appendix A.2 describes this pipeline in detail. I aggregate the resulting classifications to polling booths, map booths to pincodes using geographic identifiers and administrative boundary files, and construct polling booth-year counts of total, Muslim, and Hindu voters. The resulting three-period panel supports the analysis of how legal restrictions

[3] Both 2014 and 2017 electoral rolls data were accessed under restricted-use agreements because they contain personally identifiable information.

[4] This classification is conservative: high-confidence Muslim tokens in voter or household-member names identify Muslim voters, Hindu surname tokens identify Hindu voters when no Muslim token is present, and unresolved cases are classified as Hindu to minimize false positives in Muslim identification.

on property transactions affect residential sorting and segregation. Online Appendix A.3.1 defines the segregation measures used in the analysis. Muslim and Hindu shares capture local group presence, while the log Muslim-to-Hindu ratio captures relative group balance. The mixed-neighborhood index[5] captures how religiously mixed the areas surrounding a polling booth in a pincode are. The isolation contribution[6] measures how much each polling booth contributes to within-pincode Muslim isolation.

3.3 Housing Markets Data

I study housing-market responses using project-level administrative data from the Gujarat Real Estate Regulatory Authority (GujRERA), which records registered real-estate projects from 2005 to 2024 and includes information on project timing, location, scale, prices, costs, bookings, and expected completion. These data allow me to track new formal housing supply in the primary market and the intensity of built-up urban space. I geocode projects using reported locations, map them to pincodes, and aggregate them to the pincode-year level to capture whether legal restrictions on inter-religious property transactions change where development occurs, how intensively land is used, and how the local price and cost structure of housing evolves over time: Total units are the sum of units in projects launched in a pincode-year; price and land cost per square foot are constructed from project sales, land

[5] Motivated by spatial measures of segregation that define local environments using the proximity-weighted composition of surrounding units (Reardon and O'Sullivan, 2004), I construct the mixed-neighborhood index as

$$\text{MixedNeighborhood}_{it} \equiv \sum_{j \neq i} w_{ij} \min\left\{\frac{M_{jt}}{P_{jt}}, \frac{H_{jt}}{P_{jt}}\right\},$$

where w_{ij} are spatial weights over units near polling booth i and P_{jt} is the total electors of nearby unit j. The minimum-share term is largest when Hindu and Muslim shares are balanced and approaches zero as either group becomes absent. The exact minimum-share formulation is constructed for this paper.

[6] Following the standard exposure and isolation measures introduced by Bell (1954) and summarized by Massey and Denton (1988), the contribution of polling booth i to within-pincode Muslim isolation is

$$\text{IsolationContribution}_{it} \equiv \left(\frac{M_{it}}{M_t}\right)\left(\frac{M_{it}}{P_{it}}\right),$$

where M_t is the total Muslim voters in the relevant pincode-year. The first term is the share of the pincode's Muslim voters living in booth i, and the second is the Muslim share within that booth. Summing this expression across all booths in a pincode yields the standard Muslim isolation index (Asher et al., 2026).

costs, and built area and are rescaled by Rs. 10,000; unit density divides total launched units by pincode area in square kilometers; and congestion measures total project built area as a share of pincode land area. Online Appendix A.3.2 defines these measures in detail.

3.4 Public Safety Data

I measure incidence of violence using a georeferenced event dataset: the Global Database of Events, Language, and Tone (GDELT), a broad news-based source covering 1980–2022. I construct pincode-year measures of public safety by reverse-geocoding Gujarat events and using a text model to classify event narratives and actors into three categories – communal violence, civilian-targeting and total fatalities[7]. Online Appendix A.3.3 describes the construction in detail.

3.5 School Inputs and Outcomes Data

I use the Gujarat supplement of the Unified District Information System for Education Plus (UDISE+), India's annual administrative census of recognized schools, to study whether DAA exposure is associated with changes in the local schooling environment from 2005 to 2022. Schools are a relevant margin because local schooling decisions depend on nearby school characteristics, and school quality is central to human-capital formation and persistent inequality (Drèze and Kingdon, 2001; Muralidharan and Sundararaman, 2015; Hanushek and Woessmann, 2012). UDISE+ reports school location, management type, enrollment, staffing, infrastructure, grants, and student outcomes for schools from pre-primary through grade 12. I aggregate school-level records to the pincode-year level and construct measures of school resources, management composition, institutional age, infrastructure, and student performance: grants received and spent per student, pupil–teacher ratios, government- and private-school shares, school age, a year-standardized infrastructure index based on basic

[7] Communal violence, defined as events involving inter-group conflict such as riots, identity-based clashes, and related communal confrontation; civilian targeting, defined as events in which civilians are directly targeted or appear as the targeted side; and fatalities, defined as the sum of reported event-level deaths.

facilities and school amenities, boys' and girls' enrollment shares, and stage-specific exam failure rates. Online Appendix A.3.4 defines these measures in detail.

4 Empirical Strategy

4.1 Unit of Analysis, Sample Construction, and Treatment Definitions

The analysis uses polling booths for residential composition, because they provide fine neighborhood geography, and pincodes for housing, schools, and violence, because those outcomes are reported or more reliably aggregated at that scale. I use the longest feasible panel for each outcome rather than impose a common short sample, and define treatment cohorts by the first year in which a booth or pincode contains any locality notified under the DAA. Treatment is absorbing, since I do not observe systematic denotification. The baseline treatment indicator is an extensive-margin measure: a polling booth or pincode is treated once it contains any locality notified under the Disturbed Areas Act. I also construct pincode-level intensity measures based on the number of distinct disturbed localities in the first year of being notified, which capture how deeply the law penetrates local space and allow me to compare lightly and heavily exposed places. Together, these panels allow me to study how legal restrictions on inter-religious property transactions reshape residential composition, neighborhood amenities, housing markets, and local conflict over time.

4.2 Difference-in-Differences Designs and Identification

I use two difference-in-differences designs, matched to the structure of the outcome data. For residential-segregation outcomes, religious composition is observed only in three benchmark years—2014, 2017, and 2025—so a full dynamic event-study design is infeasible at the polling-booth level. I therefore implement a stacked difference-in-differences design following the

cohort-stacking logic in Cengiz et al. (2019), Wing et al. (2024), and Gardner (2022). The first stack compares polling booths in pincodes first exposed to the Disturbed Areas Act between 2014 and 2017 to booths in pincodes not yet treated over the same period; the second compares booths first exposed between 2017 and 2025 to booths in pincodes that remain untreated over that later window. Each stack is a short two-period difference-in-differences comparison, and the pooled regression is

$$Y_{ipst} = \alpha_{ps} + \lambda_{st} + \beta D_{pst} + \varepsilon_{ipst},$$

where Y_{ipst} is a residential-composition outcome for polling booth i, located in pincode p, in stack s and year t. The term α_{ps} denotes stack-by-pincode fixed effects, and λ_{st} denotes stack-by-year fixed effects. D_{pst} equals one for pincodes assigned to the treated cohort in the post-treatment period of stack s. Because treatment is assigned at the pincode level, all polling booths within a treated pincode receive the same treatment status in a given stack-year. Identification requires a within-stack parallel-trends assumption: absent DAA exposure, booth-level composition in treated and comparison pincodes would have changed similarly over the corresponding electoral-roll interval. The stack-specific fixed effects allow the same calendar year to play different roles across stacks, converting staggered timing into cleaner cohort-specific comparisons.

For housing, school, and violence outcomes, I observe longer annual pincode-year panels and estimate dynamic treatment effects using staggered-adoption event-study designs. A useful benchmark is the conventional event-study regression

$$Y_{it} = \alpha_i + \lambda_t + \sum_{k \neq -1} \beta_k \mathbf{1}\{t - G_i = k\} + \varepsilon_{it},$$

where G_i is the first year in which pincode i is exposed to the DAA and $k = -1$ is the omitted pre-treatment period. Because standard two-way fixed-effects event studies can be misleading when treatment effects are heterogeneous across cohorts or event time (Sun and Abraham, 2021), I use estimators designed for staggered adoption. For housing and violence,

I use the Callaway and Sant'Anna (2021) estimator with not-yet-treated pincodes as controls, comparing units first treated in cohort g to units not yet treated in period t and aggregating group-time average treatment effects $ATT(g,t)$ into event-time profiles. For schooling, I use the Sun and Abraham (2021) interaction-weighted estimator with never-treated pincodes as controls:

$$Y_{it} = \alpha_i + \lambda_t + \sum_g \sum_{k \neq -1} \beta_{gk} \mathbf{1}\{G_i = g\} \mathbf{1}\{t - g = k\} + \varepsilon_{it}.$$

I use never-treated pincodes as the comparison group for schooling because these slower-moving public-good and human-capital outcomes—infrastructure, grants, management composition, enrollment, staffing, and student performance—may already reflect institutional or demographic trajectories associated with eventual DAA designation before treatment. Later-treated pincodes are therefore less suitable controls than for housing or violence, where not-yet-treated units provide useful same-year comparisons. The Sun and Abraham (2021) estimator compares cohort-specific treatment paths to untreated units while avoiding contamination from heterogeneous treatment effects in conventional two-way fixed-effects event studies. In all pincode-level specifications, unit fixed effects absorb persistent location differences, year fixed effects absorb common shocks, and standard errors are clustered at the pincode level.

The identifying assumption is that, absent DAA designation, treated units and their comparison units – not-yet-treated pincodes for housing and violence, and never-treated pincodes for schooling – would have followed similar trends after controlling for unit and year fixed effects. This assumption matters because the DAA is justified as preventive; if designation responds to rising violence or deteriorating neighborhood conditions, treated areas may already be on different trajectories. I therefore place weight on the pre-treatment event-study coefficients, while recognizing that flat pre-trends are diagnostic rather than definitive. I also examine treatment targeting and baseline neighborhood characteristics in the descriptive results and Appendix B.1. Because any cross-pincode spillovers are likely to be highly local, I examine them directly using complementary spatial designs that exclude nearby untreated pincodes from the comparison group and estimate how untreated pincodes change when they

become close to active DAA areas (Butts, 2021). These exercises provide little evidence of substantial spillovers across pincodes, suggesting that they are unlikely to drive the main results.

4.3 Heterogeneity by Baseline DAA Intensity

The Disturbed Areas Act does not enter all places with the same initial depth, and this variation may matter if the law raises transaction frictions within local housing markets. I measure baseline intensity as the number of distinct disturbed localities covered within a pincode when it first enters the DAA regime.[8] I then hold this measure fixed and divide treated pincodes at the median of its baseline distribution into low- and high-intensity groups. Since the median is at one, low-intensity pincodes contain exactly one DAA-notified locality at entry; high-intensity pincodes contain multiple DAA-notified localities at entry. Baseline intensity is fixed at first DAA entry. Using intensity at entry avoids conditioning on subsequent expansions of DAA coverage, which may respond endogenously to changes in the outcomes of interest. For housing and violence outcomes, I estimate separate staggered event-study profiles for each subgroup using the Callaway and Sant'Anna (2021) estimator with not-yet-treated units as controls. For schooling outcomes, I use the Sun and Abraham (2021) estimator with never-treated pincodes as controls, keeping the comparison group consistent with the main schooling specifications. These subgroup estimates test whether DAA effects differ with the local depth of exposure and whether pooled estimates mask offsetting responses across lightly and heavily regulated places.

[8] It may be natural to normalize this measure by pincode area, since larger pincodes could mechanically contain more DAA-notified localities. In the data, however, I find no strong positive relationship between the baseline number of notified localities and pincode size: the Pearson correlation is -0.038. I therefore use the number of notified localities as the baseline intensity measure.

4.4 Spatial Spillovers from DAA Exposure

To study spatial spillovers, I construct a pincode-year exposure panel around active DAA pincodes. The objective is to separate the direct effect of entering the DAA regime from indirect effects operating through proximity to newly treated space, since a place-based restriction on property transactions may redirect households, projects, and local activity beyond formally treated locations and contaminate nearby untreated controls (Butts, 2021). I implement two designs. Design A is a donut-style event study for the direct treatment effect: treated pincodes are compared to untreated pincodes outside an immediate spillover zone, excluding untreated pincodes within 5 kilometers of an active treated pincode. This conservative donut removes the nearby untreated units most likely to be indirectly exposed through redirected transactions, residential adjustment, or local market spillovers. Design B treats proximity to treatment as the object of interest. For untreated pincodes, I define spillover exposure as the first year in which a pincode falls within a distance band of an active treated pincode and estimate separate event studies for mutually exclusive 0–5 kilometer and 5–10 kilometer bands. These specifications test whether untreated pincodes change when they enter the spatial orbit of nearby DAA activity, even if they are not directly designated, and whether spillover responses attenuate with distance. Together, Design A asks whether the direct treatment effects are robust to excluding nearby untreated controls, while Design B asks whether untreated space itself responds dynamically to proximity to active DAA areas.

5 Where Is the Disturbed Areas Act Applied?

I begin by examining where the DAA is applied, since its targeting shapes how later treatment effects should be interpreted. The polling-booth results show that DAA designation is more likely where Muslims are locally present and the broader environment is mixed, rather than in already homogeneous space.

Table 1: Targeting of DAA Designation across Polling Booths

	Dependent Variable: Disturbed-Area Designation Indicator			
Booth Muslim Share	0.243***			
	(0.040)			
Booth Hindu Share		-0.028		
		(0.038)		
Mixed-Neighborhood Index			0.197***	
			(0.036)	
Muslim-Majority Booth				-0.004
				(0.056)
Hindu-Majority Pincode (LOO)				0.004
				(0.090)
Muslim-Majority Booth × Hindu-Majority Pincode (LOO)				0.137**
				(0.063)
Constant	0.370***	0.407***	0.367***	0.379***
	(0.019)	(0.038)	(0.019)	(0.090)
District FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dependent Variable Mean	0.384	0.384	0.384	0.386
Adjusted R^2	0.212	0.208	0.212	0.212
Observations	89,240	89,240	89,240	88,634

Notes: The unit of observation is the polling booth-year. The dependent variable is an indicator equal to one if the polling booth falls within a pincode exposed to the Disturbed Areas Act. Each column reports a separate regression. All specifications include district and year fixed effects. Standard errors, reported in parentheses, are clustered at the pincode level. "LOO" denotes a leave-one-out measure of pincode religious composition that excludes the focal booth when constructing the pincode majority indicator. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A 10 percentage-point increase in booth Muslim share is associated with a 2.4 percentage-point higher probability of designation, relative to a mean of 38.4 percent, while the Hindu-share coefficient is negative, small, and statistically insignificant. Mixed local environments are also more likely to be designated: a 10 percentage-point increase in the mixed-neighborhood index raises designation probability by 2.0 percentage points. This is plausible because polling booths are very fine units; many treated booths remain Hindu-majority in levels but have higher Muslim shares than comparison booths and sit in more mixed sur-

roundings. Column 4 sharpens this interpretation using leave-one-out pincode composition[9]: Muslim-majority booths located inside broader Hindu-majority pincode environments are 13.7 percentage points more likely to be designated, relative to a mean of 38.6 percent.

Together, these results suggest that the DAA enters boundary spaces where Muslims are locally present and surrounding areas remain mixed or Hindu-majority. This is consistent with the law entering neighborhoods where cross-community transactions and marginal demographic change are still feasible, so administrative delay and uncertainty can affect who enters, remains, or transacts across boundaries. The same pattern could also reflect a protection rationale if the state intervenes where Muslims are more exposed within mixed or non-Muslim surroundings. I therefore treat this as evidence on where the law enters and rely on later treatment-effect and violence results to distinguish those interpretations. Cohort diagnostics in Online Appendix B.1 reinforce this view, with early-treated booths more Muslim and more mixed, while later-treated booths resemble never-treated and more homogeneous booths before treatment.

6 Hindu–Muslim Residential Segregation

I focus on Hindu–Muslim separation because the residual religious category combines several smaller groups[10] that cannot be reliably measured separately at the polling-booth level, and because the DAA emerged from Gujarat's history of Hindu–Muslim riots and is most relevant in practice for Hindu–Muslim property transfers (Jaffrelot, 2003, 2011; Tejani, 2023; Laliwala et al., 2021). The results should therefore be interpreted as evidence on Hindu–Muslim contact and boundary-hardening, not as a complete measure of all dimensions of religious diversity.

[9] I compute the religious composition of the rest of the pincode after excluding the focal booth, so that the surrounding environment is not mechanically driven by the booth's own composition.

[10] Christians, Sikhs, Buddhists, Jains, other religions and persuasions, and religion not stated (Office of the Registrar General and Census Commissioner, India, 2011b; Press Information Bureau, Government of India, 2015).

6.1 Boundary-Hardening in Treated Neighborhoods

I first estimate treatment effects on booth-level residential composition by comparing treated booths to not-yet-treated areas. The strongest evidence comes from the early-treatment window, which compares booths first exposed between 2014 and 2017 to booths not yet treated over the same period. DAA exposure is associated with a statistically significant 5.3 percentage-point decline in Hindu share, relative to a baseline mean of 81.1 percent, and a 0.893 log-point increase in the Muslim-to-Hindu ratio, implying that the ratio rises by about 144 percent relative to the comparison group. By contrast, the Muslim-share coefficient is small and imprecisely estimated, at -1.6 percentage points relative to a baseline mean of 5.1 percent. The isolation-contribution coefficient is 0.002, about half of the baseline mean of 0.004, indicating that treated booths contribute more to pincode-level Muslim isolation after designation. Thus, the main effect is not broad Muslim inflow but sharper relative Muslim concentration within treated space because Hindu presence declines and treated booths become more central to within-pincode Muslim exposure.

Table 2: Early-Cohort Difference-in-Differences Estimates for Booth-Level Religious Composition, 2014–2017

	(1)	(2)	(3)	(4)
	Muslim share	Hindu share	Log Muslim-to-Hindu ratio	Isolation contribution
DAA Exposure × Post	-0.016	-0.053***	0.893**	0.002***
	(0.016)	(0.017)	(0.381)	(0.001)
Pincode FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dependent Variable Mean	0.051	0.811	-5.463	0.004
Adjusted R^2	0.191	0.639	0.527	0.202
Observations	72,081	72,081	72,081	71,715

Notes: The unit of observation is the polling booth-year. The sample compares booths first exposed to the Disturbed Areas Act between 2014 and 2017 to booths not yet treated over the same period. Treatment is assigned at the pincode level. *Post* is an indicator for the post-designation period. Muslim share and Hindu share do not necessarily sum to one because some voters remain outside the Hindu/Muslim classification in the underlying data. Isolation contribution reports each booth's contribution to within-pincode Muslim isolation. All specifications include pincode and year fixed effects. Standard errors, reported in parentheses, are clustered at the pincode level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The stacked specification, which pools cohort-specific treatment windows while preserving relative event time, delivers the same qualitative conclusion. Because polling booths are much smaller than pincodes, these results indicate that the DAA reshapes very local residential composition rather than simply shifting average shares across broad administrative areas. After designation, treated booths become less mixed relative to the rest of the pincode: Hindu presence declines, the Muslim-to-Hindu ratio rises, and treated booths account for more within-pincode Muslim isolation. A cautious interpretation is that the DAA contributes to early-stage ghettoization. If designation makes a neighborhood politically sensitive, administratively burdensome, or less attractive for cross-community residence, Hindu households may become less willing to remain in or move into those localities. The evidence is therefore consistent with legal designation hardening communal boundaries by shifting mixed or boundary areas toward sharper residential separation.

Table 3: Stacked Difference-in-Differences Estimates for Booth-Level Religious Composition

	(1)	(2)	(3)	(4)
	Muslim share	Hindu share	Log Muslim-to-Hindu ratio	Isolation contribution
DAA Exposure × Post	-0.017	-0.048***	0.849**	0.002***
	(0.015)	(0.016)	(0.359)	(0.001)
Stack × Pincode FE	Yes	Yes	Yes	Yes
Stack × Year FE	Yes	Yes	Yes	Yes
Dependent Variable Mean	0.040	0.866	-6.177	0.004
Adjusted R^2	0.206	0.686	0.510	0.194
Observations	105,826	105,826	105,826	105,181

Notes: The unit of observation is the polling booth-year. The estimates are from a stacked difference-in-differences design that pools cohort-specific treatment windows while preserving relative event time. Treatment is assigned at the pincode level. *Post* is an indicator for the post-designation period within each stack window. Muslim share and Hindu share do not necessarily sum to one because some voters remain outside the Hindu/Muslim classification in the underlying data. Isolation contribution reports each booth's contribution to within-pincode Muslim isolation. All specifications include stack-by-pincode and stack-by-year fixed effects. Standard errors, reported in parentheses, are clustered at the pincode level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

I next test whether composition effects vary by baseline DAA intensity, as defined in Section 4.3. In the stacked specification, the low-intensity effect is captured by *DAA Exposure* × *Post*, while the high-intensity interaction gives the incremental effect in more deeply exposed

places. Muslim and Hindu shares show small and imprecise effect sizes. The log Muslim-to-Hindu ratio rises significantly in low-intensity areas—by 0.948 log points or about 2.6 times as large as untreated areas—while the high-intensity interaction is negative and imprecise. Thus, the clearest effect is a rise in the Muslim-to-Hindu ratio among low-intensity areas, while the high-intensity differentials are small and imprecisely estimated across outcomes; the results therefore support boundary hardening after designation, but not a monotonic relationship with baseline DAA depth.

Table 4: Stacked Difference-in-Differences Estimates for Booth-Level Religious Composition by Baseline DAA Intensity

	(1)	(2)	(3)	(4)
	Muslim share	Hindu share	Log Muslim-to-Hindu ratio	Isolation contribution
DAA Exposure × Post	-0.028	-0.037	0.948*	0.002
	(0.026)	(0.026)	(0.565)	(0.001)
DAA Exposure × Post × High Intensity	0.023	-0.024	-0.205	0.001
	(0.028)	(0.030)	(0.695)	(0.001)
Stack × Pincode FE	Yes	Yes	Yes	Yes
Stack × Year FE	Yes	Yes	Yes	Yes
Dependent Variable Mean	0.040	0.866	-6.177	0.004
Adjusted R^2	0.207	0.686	0.510	0.194
Observations	105,826	105,826	105,826	105,181

Notes: The unit of observation is the polling booth-year. The estimates are from a stacked difference-in-differences design that pools the 2014–2017 and 2017–2025 cohort-specific treatment windows. Treatment is assigned at the pincode level. *DAA Exposure* × *Post* reports the treatment effect for low-intensity treated pincodes. *DAA Exposure* × *Post* × *High Intensity* reports the incremental treatment effect for high-intensity treated pincodes relative to low-intensity treated pincodes. Baseline intensity is the number of distinct notified localities in the pincode when it first enters the DAA regime and is fixed thereafter. High-intensity pincodes are those above the median of this baseline distribution among ever-treated pincodes. Muslim share and Hindu share do not necessarily sum to one because some voters remain outside the Hindu/Muslim classification in the underlying data. Isolation contribution reports each booth's contribution to within-pincode Muslim isolation. All specifications include stack-by-pincode and stack-by-year fixed effects. Standard errors, reported in parentheses, are clustered at the pincode level. The implied treatment effect for high-intensity treated pincodes is the sum of the two reported coefficients. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.2 Spillovers around Treated Neighborhoods

The booth-level spillover results are consistent with the main composition results, but weaker and less uniform. Separating directly treated booths from nearby untreated booths in the 0–5 kilometer ring leaves the direct-treatment estimates close to the baseline: treated booths

show a large decline in Hindu share, a substantial rise in the Muslim-to-Hindu ratio, little change in Muslim share, and a larger contribution to pincode-level Muslim isolation. The Hindu-share effect is -6.6 percentage points, compared with -5.3 percentage points in the baseline early-cohort specification, and the log Muslim-to-Hindu ratio rises by 1.032 log points, implying a ratio about 2.8 times as large as in the comparison group.

Table 5: Early-Cohort Difference-in-Differences Estimates with Direct Treatment and Spillover Exposure, 2014–2017

	(1)	(2)	(3)	(4)
	Muslim share	Hindu share	Log Muslim-to-Hindu ratio	Isolation contribution
DAA Exposure × Post	-0.015	-0.066***	1.032***	0.003***
	(0.016)	(0.017)	(0.385)	(0.001)
Spillover Exposure × Post	0.028***	-0.068***	-0.127	0.003***
	(0.007)	(0.018)	(0.334)	(0.001)
Pincode FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dependent Variable Mean	0.051	0.810	-5.547	0.004
Adjusted R^2	0.203	0.646	0.545	0.225
Observations	53,492	53,492	53,492	53,180

Notes: The unit of observation is the polling booth-year. The sample includes three groups of booths over the 2014–2017 window: directly treated booths in pincodes first exposed to the Disturbed Areas Act between 2014 and 2017, spillover booths in untreated pincodes that newly enter the 0–5 kilometer ring around treated pincodes over the same period, and comparison booths in untreated pincodes that remain outside the 0–5 kilometer ring through 2017. The omitted category is the comparison group: booths in untreated pincodes outside the 0–5 kilometer ring through 2017. *DAA Exposure* is an indicator for booths in directly treated pincodes, *Spillover Exposure* is an indicator for booths in nearby untreated pincodes entering the 0–5 kilometer ring, and *Post* is an indicator for the post-designation period. Muslim share and Hindu share do not necessarily sum to one because some voters remain outside the Hindu/Muslim classification in the underlying data. Isolation contribution reports each booth's contribution to within-pincode Muslim isolation. All specifications include pincode and year fixed effects. Standard errors, reported in parentheses, are clustered at the pincode level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The main treatment pattern is not driven by comparison with untreated units immediately adjacent to treated space. Nearby untreated booths also adjust, but less consistently than directly treated booths. Booths entering the 0–5 kilometer ring show higher Muslim share, lower Hindu share, and a positive contribution to within-pincode Muslim isolation, while the log-ratio response is small and imprecise. Online Appendix B.1 shows the stacked spillover

estimates which display the same pattern: nearby untreated space responds at the margin, but the strongest and most consistent changes remain concentrated inside formally treated areas. A cautious interpretation is that DAA designation redirects some residential adjustment toward nearby untreated locations that remain legally open but are geographically close to treated space. The more limited conclusion is that the DAA is associated with some adjustment in surrounding space, while its clearest booth-level effects remain highly localized within treated areas.

Allowing the direct treatment effect to vary with baseline DAA intensity leaves the main spillover results largely unchanged. In Online Appendix B.1, the direct effect for low-intensity treated pincodes is a 5.1 percentage-point decline in Hindu share, a 1.083 log-point increase in the Muslim-to-Hindu ratio, and a 0.002 increase in isolation contribution. The high-intensity interactions are comparatively small and imprecisely estimated across all four outcomes. Although the point estimates imply somewhat larger declines in Hindu share and increases in isolation in high-intensity treated pincodes, the differences from low-intensity areas are not statistically distinguishable. The spillover coefficients remain similar to the pooled specification: nearby untreated booths experience a 2.4 percentage-point increase in Muslim share, a 6.2 percentage-point decline in Hindu share, little change in the Muslim-to-Hindu ratio, and a positive isolation-contribution effect. Thus, accounting for spillovers continues to support localized boundary hardening in directly treated areas and compositional adjustment in nearby untreated space.

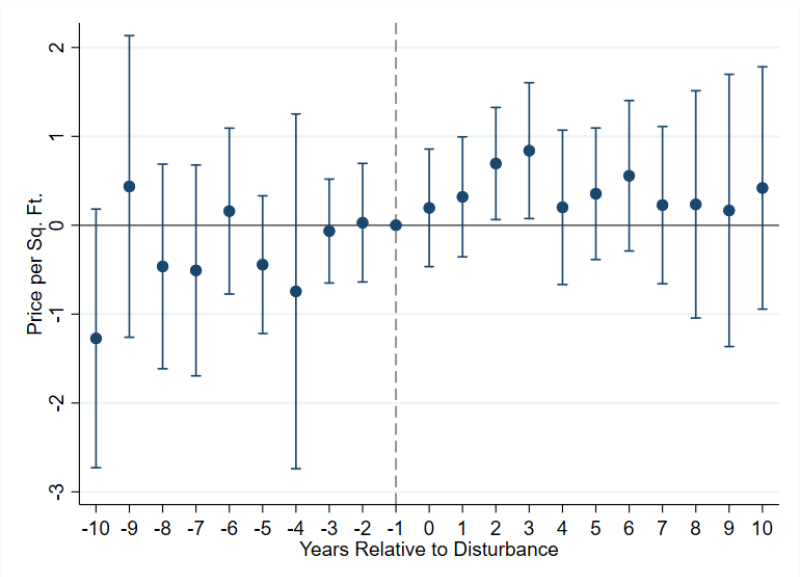
7 Neighborhood Amenities and Urban Inequality

I next examine housing, schooling, and communal violence as downstream margins through which DAA-induced boundary-hardening may reshape neighborhood amenities, public safety, and the geography of urban opportunity.

7.1 Housing Markets

The housing results suggest market reorganization rather than a uniform average effect. Pooled event studies are modest and often imprecise: Figure 4 shows price per square foot remaining close to zero before treatment and rising only gradually afterward. This average pattern is informative but incomplete, because Appendix B.2 shows that it masks substantial heterogeneity across treated places and housing outcomes.

Figure 4: Pooled Effect of DAA Exposure on Housing Prices

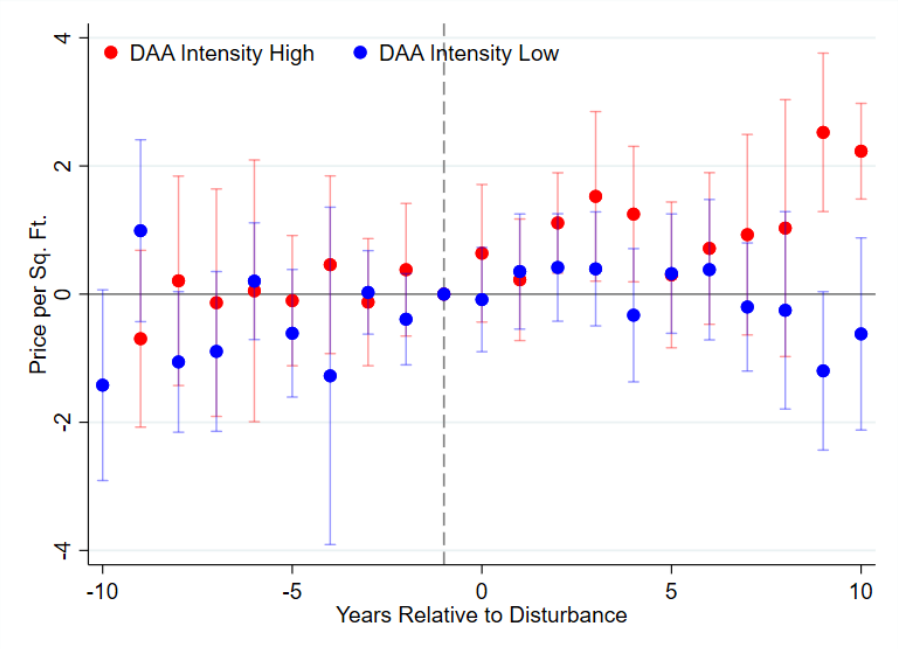


Notes: The figure plots dynamic treatment effects of DAA exposure on price per square foot (in Rs. 10,000) using the Callaway and Sant'Anna estimator with not-yet-treated pincodes as controls. The unit of observation is the pincode-year. Event time is defined relative to the first year in which a pincode contains any locality notified under the Disturbed Areas Act. The omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

To study this heterogeneity, I split treated pincodes by baseline DAA intensity, as defined in Section 4.3. Figure 5 reveals a sharper divergence than the pooled estimates suggest. In high-intensity areas, the estimated effect on price per square foot rises after designation and reaches approximately 2.2–2.5 in the final two event years, equivalent to increases of roughly Rs. 22,000–25,000 per square foot relative to the comparison group. In low-intensity areas,

the corresponding estimates are approximately negative of 1.3 and 0.7, or Rs. 13,000 and Rs. 7,000 lower per square foot, although these estimates are imprecise. By the final event year, the estimated difference between high- and low-intensity areas is approximately Rs. 29,000 per square foot.

Figure 5: Effect of DAA Exposure on Housing Prices by Baseline Treatment Intensity



Notes: The figure plots dynamic treatment effects of DAA exposure on price per square foot (in Rs. 10,000) separately for high- and low-intensity treated pincodes using the Callaway and Sant'Anna estimator with not-yet-treated pincodes as controls. Baseline intensity is the number of distinct notified localities in the pincode when it first enters the DAA regime and is fixed thereafter. High-intensity pincodes are those above the median of this baseline distribution among ever-treated pincodes. The unit of observation is the pincode-year. The omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

Online Appendix B.2 shows similar divergence for other housing outcomes. In high-intensity areas, total units rise gradually to roughly 200–300 additional units, unit density remains mostly stable before increasing by about 100–150 units per square kilometer in the final years, and congestion does not fall, with late coefficients turning positive at roughly 0.001–0.002, or 0.1–0.2 percentage points of pincode land area. In low-intensity areas, housing activity weakens instead: total units fall by about 200 in later post-treatment years, and

unit density declines by roughly 50–100 units per square kilometer. The divergence emerges gradually several years after treatment rather than immediately.

The joint movement of prices and quantities is informative about the nature of this adjustment. In high-intensity areas, prices rise alongside total units and unit density, a pattern more consistent with stronger housing demand accompanied by a developer response than with a pure contraction in housing supply, which would tend to raise prices while reducing quantity. I do not observe housing demand and supply schedules separately, however, so I treat this pattern as suggestive rather than as a structural decomposition.

The residential results in Section 6.1 imply that while DAA designation hardens religious boundaries across treated areas, the incremental high-intensity effects on Muslim share, Hindu share, the Muslim-to-Hindu ratio, and isolation are all imprecisely estimated. The post-treatment housing divergence therefore seems unlikely to reflect a stronger segregation first stage in high-intensity places. Nor does it appear to be a simple continuation of stronger pre-existing housing activity, since pre-treatment prices are similar across the two groups and high-intensity areas do not begin with systematically greater unit density or housing supply.[11] Before treatment, electoral-roll counts suggest that high-intensity pincodes contain a larger absolute residential market – larger number of total and Muslim registered voters – compared to low intensity pincodes.

One interpretation of the housing divergence is a critical-mass hypothesis that becomes relevant for segmented markets, as induced by the DAA. DAA exposure makes some cross-community transactions harder in both intensity groups, but the consequences may depend on the scale of the resulting market. In low-intensity pincodes, which are also smaller residential markets, a single restricted locality may be too small to provide enough alternative transactions, so affected buyers may postpone purchases or search outside the pincode, weakening local housing activity. In high-intensity pincodes, several restricted localities embedded

[11] In the baseline comparisons, Muslim share is 4.1 percent in low-intensity pincodes and 4.3 percent in high-intensity pincodes; Muslim isolation is 0.139 and 0.180, respectively, with the difference statistically imprecise. Pre-treatment price per square foot, unit density and total units are also statistically similar across groups.

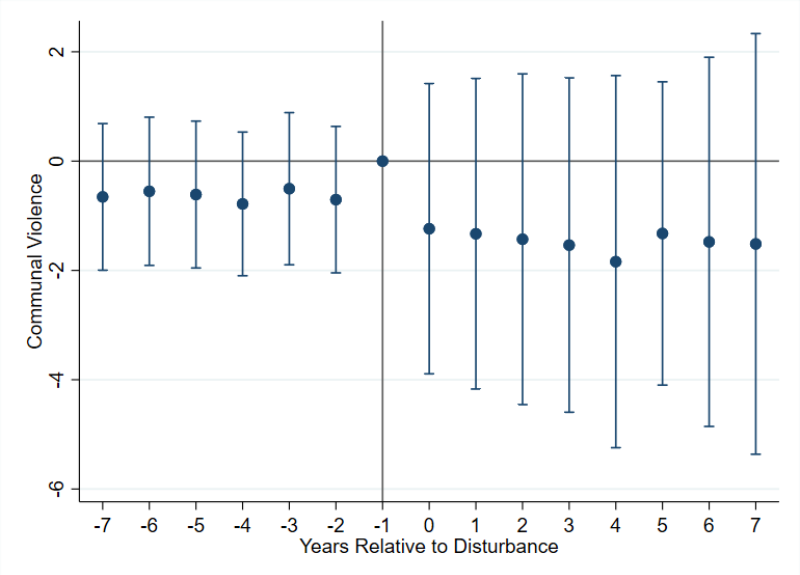
within a larger residential market may instead constitute a sufficiently large segmented submarket for constrained buyers to continue searching locally among feasible sellers, projects, and locations. If this sustains a sufficiently large buyer base, developers have an incentive to add housing supply, generating the observed combination of rising prices and construction. Thus, despite the DAA making exchange harder, a critical mass of DAA localities in a larger residential market may reach the scale required to retain buyers and support a developer response. The larger absolute residential market observed in high-intensity pincodes is consistent with this interpretation, but the channel remains suggestive because I do not observe buyer–seller religion, DAA approval decisions, or separate demand and supply shifts.

The housing results are robust to the spillover-sensitive donut design from Section 4.4, which excludes untreated pincodes within 5 kilometers of active DAA pincodes: pooled estimates remain modest and imprecise, while high-intensity areas continue to show stronger post-treatment movement in price per square foot, land cost per square foot, and total units than low-intensity areas, suggesting that the main patterns are not driven by nearby contaminated controls; see Appendix B.2.1.

7.2 Public Safety

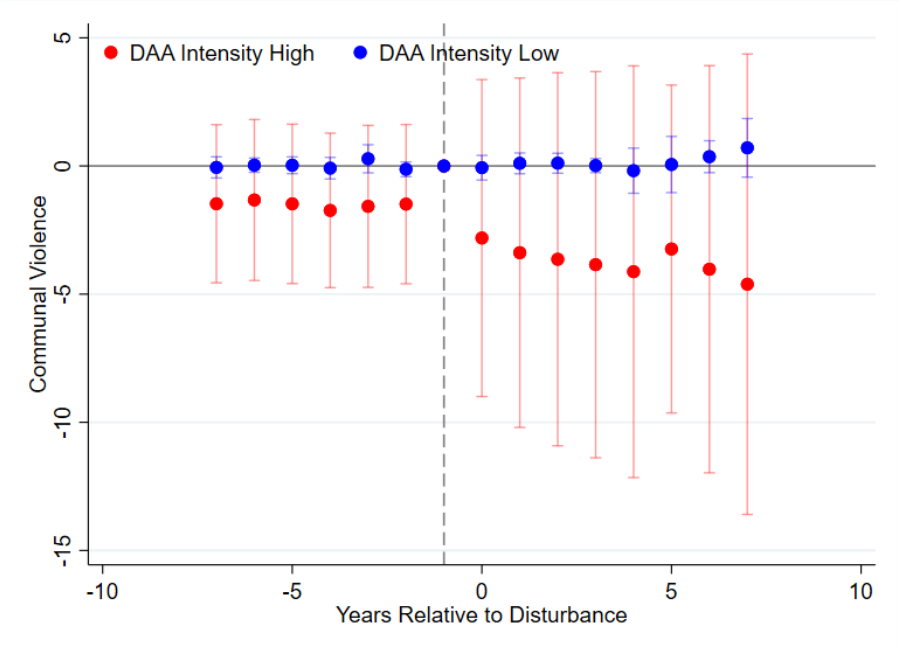
The violence results do not support a simple preventive interpretation of the DAA. If designation were mainly triggered by sharply rising communal risk as expected from the language of the law, then violence should increase before treatment, especially in high-intensity areas. Instead, the pooled GDELT event study for communal violence in Figure 6 and the intensity-split estimates in Figure 7 show pre-treatment coefficients that are broadly flat and centered near zero.

Figure 6: Pooled Effect of DAA Exposure on Communal Violence



Notes: The figure plots dynamic treatment effects of DAA exposure on communal violence using GDELT data and the Callaway and Sant'Anna estimator with not-yet-treated pincodes as controls. The unit of observation is the pincode-year. Event time is defined relative to the first year in which a pincode contains any locality notified under the Disturbed Areas Act. The omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

Figure 7: Effect of DAA Exposure on Communal Violence by Baseline Treatment Intensity



Notes: The figure plots dynamic treatment effects of DAA exposure on communal violence using GDELT data, separately for high- and low-intensity treated pincodes, using the Callaway and Sant'Anna estimator with not-yet-treated pincodes as controls. Baseline intensity is the number of distinct notified localities in the pincode when it first enters the DAA regime and is fixed thereafter. High-intensity pincodes are those above the median of this baseline distribution among ever-treated pincodes. The unit of observation is the pincode-year. The omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

After treatment, the main communal-violence estimates become negative though noisy, particularly in high-intensity areas, while low-intensity coefficients remain closer to zero. This pattern is somewhat preserved under the spillover-sensitive donut design, which excludes untreated units within 5 kilometers of active DAA pincodes, as reported in Figure A.10 in Online Appendix B.3.1: communal-violence coefficients move toward zero and become positive in some later years while remaining statistically imprecise. Figures A.8 and A.9 in Online Appendix B.3 show similarly negative though noisy post-treatment point estimates for civilian targeting and total fatalities in the main specifications. When checking for spillover robustness, as shown in Figures A.11 and A.12 in Online Appendix B.3.1, these

outcomes continue to display imprecise differences in pre-treatment violence incidence while becoming slightly positive although still imprecise post treatment. The evidence therefore does not support a robust causal reduction in public-safety incidents following designation, as was expected from the stated public safety rationale of the law.

These patterns contest a narrow preventive interpretation of the law. The absence of a visible pre-designation escalation in communal violence suggests that DAA designation is not consistently triggered by immediate increases in observed conflict, although public-order concerns may still matter. At the same time, the lack of a statistically significant change in post-treatment violence incidence may reflect stabilization through hardened boundaries and reduced inter-community contact, or displacement of conflict risk into nearby untreated areas or unobserved forms of conflict. However, the distance-based spillover analysis in Section 7.4 finds little evidence of meaningful violence increases in nearby untreated areas. Read alongside the residential-composition results, the evidence is consistent with sharper spatial separation rather than greater social integration, but the public-safety effects themselves remain imprecise, even when controlling for potential spillovers as shown in Online Appendix B.3.1.

7.3 School Inputs and Outcomes

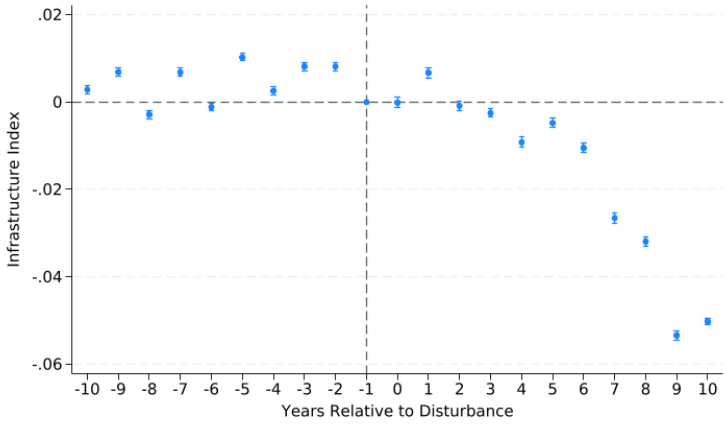
The schooling results, estimated using the Sun–Abraham estimator with never-treated pincodes as controls, suggest that the DAA reshapes local school environments, though interpretation depends on pre-treatment dynamics. I first focus on outcomes with relatively flat pre-trends in both pooled and intensity-split event studies—the infrastructure index, primary failure rates and secondary failure rates. Online Appendix B.4 reports the full results. Among the more plausibly causal outcomes, school infrastructure declines after treatment, with the pooled index falling by about 0.05 to 0.06 standard deviations by the end of the event window. This decline is concentrated in high-intensity areas, where later effects fall by 0.03 to 0.07 standard deviations. Primary failure rates also rise after treatment, especially in high-intensity areas: they increase by about 1–2 percentage points within five to seven

years and by 3–4 percentage points later. Online Appendix B.4 shows similar increases for the secondary failure rates. Together, these results indicate that deeper DAA exposure is associated with weaker school infrastructure and worse student performance, a pattern more consistent with post-treatment deterioration in school environments than with pre-existing sorting into worse schools.

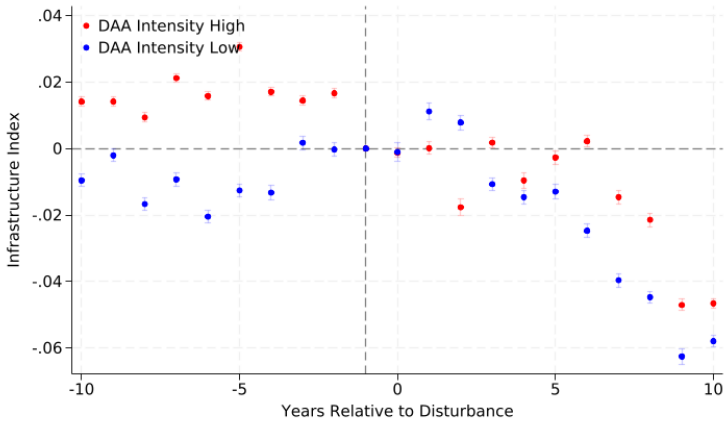
A second set of schooling outcomes points in the same direction in the intensity-split estimates but should be read more cautiously because the pooled event studies show less clean pre-trends, even though the intensity-split specifications display flatter pre-treatment dynamics. Online Appendix B.4 reports these results. School age rises after treatment, especially in high-intensity areas: pooled estimates increase by about 0.5 years within four to five years and 1.0–1.3 years near the end of the window, while high-intensity estimates reach 1.5–1.7 years within three to five years and peak around 2.5–3 years, with low-intensity areas much closer to zero. Grants received and spent per student also fall sharply, especially in high-intensity areas: grants received decline by roughly Rs. 25–45 per student in the pooled estimates and as much as Rs. 140–160 in high-intensity areas, while grants spent fall by Rs. 15–30 pooled and Rs. 110–135 in high-intensity areas. The fraction of private schools also declines more visibly in high-intensity areas, by roughly 3–4.5 percentage points in later post-treatment years, while low-intensity areas are closer to zero or positive. I therefore treat these outcomes as reinforcing rather than standalone causal evidence, but together with the cleaner results on infrastructure and failure rates, they suggest that deeper DAA exposure weakens local school environments: infrastructure deteriorates, failure rates rise, institutional inputs and school composition move adversely. Thus, housing and schooling need not move together. In high-intensity areas, housing prices and construction strengthen even as school amenities and student outcomes worsen. Under the segmented-market interpretation developed in Section 7.1, a sufficiently large constrained housing market can sustain buyer demand and induce new development without any corresponding improvement in local public goods. The housing–school divergence therefore reinforces that rising property values in these areas need not represent broad-based neighborhood improvement. The spillover-robustness checks in Online Appendix B.4.1, which exclude untreated pincodes within 5 kilometers of active

DAA pincodes, leave the main patterns visible, increasing confidence that the core schooling results are not artifacts of nearby contaminated controls.

Figure 8: DAA Exposure and School Infrastructure



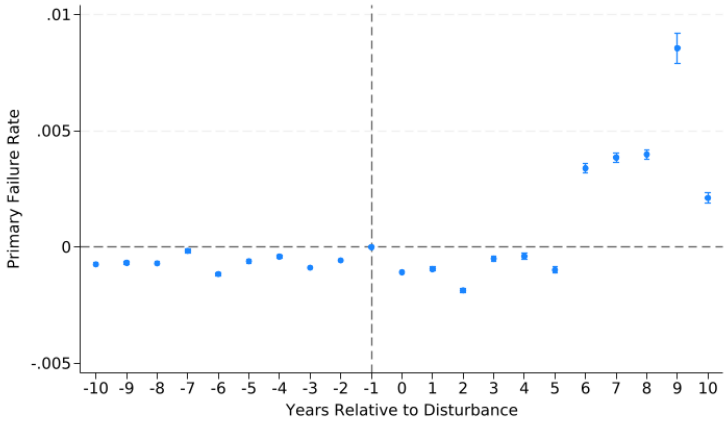
(a) Infrastructure Index (z-score)



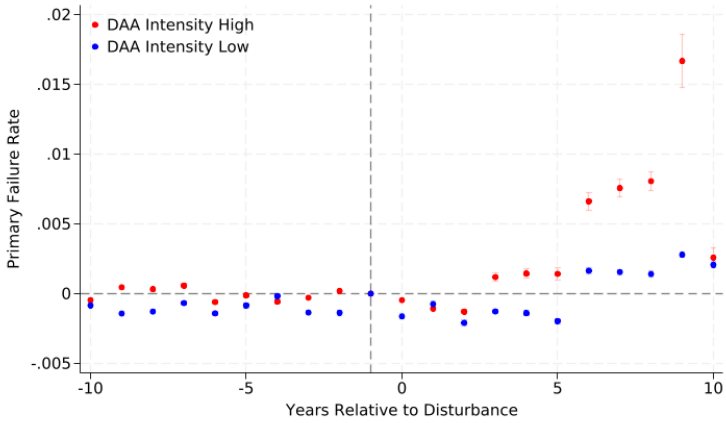
(b) Infrastructure Index by Baseline DAA Intensity

Notes: The figure reports event-study estimates for the school infrastructure index. Panel (a) shows pooled estimates. Panel (b) reports estimates separately for high- and low-intensity treated pincodes, where baseline DAA intensity is defined using the number of disturbed localities mapped to the pincode at the time of first treatment, and fixed thereafter. In both panels, estimates are obtained using the Sun and Abraham estimator with never-treated pincodes as controls. The unit of observation is the pincode-year. Event time is defined relative to first DAA exposure, and the omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

Figure 9: DAA Exposure and Primary-School Failure Rates



(a) Primary-School Failure Rate



(b) Primary-School Failure Rate by Baseline DAA Intensity

Notes: The figure reports event-study estimates for the primary-school failure rate. Panel (a) shows pooled estimates. Panel (b) reports estimates separately for high- and low-intensity treated pincodes, where baseline DAA intensity is defined using the number of disturbed localities mapped to the pincode at the time of first treatment, and fixed thereafter. In both panels, estimates are obtained using the Sun and Abraham estimator with never-treated pincodes as controls. The unit of observation is the pincode-year. Event time is defined relative to first DAA exposure, and the omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

7.4 Spillovers in Nearby Untreated Areas

Design B complements the donut-style robustness exercise in Design A by treating nearby untreated space itself as the object of interest. For untreated pincodes, I define spillover exposure as the first year in which a pincode enters the 0–5 kilometer or 5–10 kilometer ring around any active DAA pincode, and estimate separate event studies for each ring. Online Appendix B.5 reports these estimates for housing, violence, and schooling, using the same estimator as in the main design: Callaway and Sant'Anna for housing and violence, with not-yet-exposed untreated pincodes as controls, and Sun and Abraham for schooling, with never-exposed untreated pincodes as controls. The goal is to test whether DAA expansion shifts housing activity, school development, or conflict into legally open nearby pincodes.

The results suggest limited spillovers. Housing-price responses in untreated pincodes are weak and unstable: price per square foot shows no sustained post-entry break in the 0–5 kilometer ring, while the 5–10 kilometer estimates are more negative on average but noisy and dynamically unstable. Violence spillovers are also limited, with GDELT communal-violence estimates fluctuating around zero in both rings, wide confidence intervals, and no sustained post-entry increase, offering little support for displacement of conflict into nearby untreated areas. Schooling spillovers are harder to interpret causally because infrastructure and secondary-failure-rate series show noticeable pre-treatment movement in several specifications; although some post-entry coefficients remain negative, nearby untreated areas appear negatively selected or already on different trajectories before spillover exposure. Overall, Design B reinforces that the clearest DAA effects are localized within formally treated areas: together with Design A, it shows that excluding nearby untreated controls does not overturn the main direct-treatment results, while directly examining nearby untreated space does not reveal robust causal spillover effects on housing, schooling, or communal violence.

8 Conclusion

This paper shows that segregation can be reproduced not only through private sorting, explicit exclusion, or historical legal barriers, but also through contemporary administrative rules governing ordinary property exchange. The key institution is not a formal ban on minority entry, but a discretionary approval regime that raises the cost and uncertainty of crossing neighborhood boundaries. The DAA disproportionately enters boundary spaces where Muslims are locally present within mixed or Hindu-majority surroundings, and designation subsequently hardens residential boundaries – treated neighborhoods contribute more to local Muslim isolation.

The downstream consequences, however, depend on the initial depth of legal coverage even though the segregation response itself does not vary systematically with intensity. One interpretation is that multiple-locality coverage can generate sufficient scale for a segmented housing submarket to function. A sufficiently large buyer base can then induce developers to add supply. The joint rise in prices and housing quantities in high-intensity areas is consistent with this demand-and-development response, although the data do not separately identify demand and supply shifts or the underlying transaction channel. The divergence between housing values and school environments therefore illustrates a central consequence of segmented housing markets: rising property prices need not signal rising local opportunity when legal rules constrain where households can readily transact. The public-safety results provide little evidence of a robust decline in communal violence, so DAA designation does not generate clearly measurable gains in local safety despite the accompanying increase in residential separation. The broader takeaway is that contemporary legal segregation can operate through procedural discretion, boundary hardening, and market segmentation, with housing-market expansion potentially diverging from local public-good provision. As similar legal frameworks appear beyond Gujarat, including recent developments in Rajasthan (Indian Express, 2026; NDTV, 2026; Hindustan Times, 2026), the relevant policy question is not only whether such laws reduce short-run disorder, but how they alter long-run access to housing, public goods, and urban opportunity.

References

Aaronson, D., Hartley, D., and Mazumder, B. (2021). The Effects of the 1930s HOLC Redlining Maps. *American Economic Journal: Economic Policy*, 13(4):355–392.

Abel, M. (2019). Long-Run Effects of Forced Resettlement: Evidence from Apartheid South Africa. *The Journal of Economic History*, 79(4):915–953.

Alesina, A., Baqir, R., and Easterly, W. (1999). Public Goods and Ethnic Divisions. *Quarterly Journal of Economics*, 114(4):1243–1284.

Almagro, M. and Sood, A. (2025). Choices or Constraints? Disentangling Preferences and Restrictions in Housing Markets. Working paper.

Arora, A. (2024). The Causes and Consequences of Religious Segregation in Rural India. Working paper.

Asher, S., Jha, K., Novosad, P., Adukia, A., and Tan, B. J. (2026). Residential Segregation and Unequal Access to Local Public Services in India: Evidence from 1.5m Neighborhoods. Working Paper 740, Center for Global Development.

Banerjee, A. and Somanathan, R. (2007). The Political Economy of Public Goods: Some Evidence from India. *Journal of Development Economics*, 82(2):287–314.

Bayer, P., McMillan, R., and Rueben, K. (2007). An Equilibrium Model of Sorting in an Urban Housing Market. *Econometrica*, 75(6):1521–1574.

Bazzi, S., Gaduh, A., Rothenberg, A. D., and Wong, M. (2019). Unity in Diversity? How Intergroup Contact Can Foster Nation Building. *American Economic Review*, 109(11):3978–4025.

Bell, W. (1954). A Probability Model for the Measurement of Ecological Segregation. *Social Forces*, 32(4):357–364.

Bharathi, N., Malghan, D., Mishra, S., and Rahman, A. (2021). Fractal Urbanism: City Size and Residential Segregation in India. *World Development*, 141:105397.

Bharathi, N., Malghan, D., Mishra, S., and Rahman, A. (2022). Residential Segregation and Public Services in Urban India. *Urban Studies*, 59(14):2912–2932.

Bharathi, N., Malghan, D., and Rahman, A. (2019). Neighbourhood-Scale Residential Segregation in Indian Metros. *Economic and Political Weekly*, 54(30):64–70.

Black, S. E. (1999). Do Better Schools Matter? Parental Valuation of Elementary Education. *Quarterly Journal of Economics*, 114(2):577–599.

Butts, K. (2021). Difference-in-Differences Estimation with Spatial Spillovers. Working paper.

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

Carrillo, B., Charris, C., and Iglesias, W. (2023). Moved to Poverty? A Legacy of the Apartheid Experiment in South Africa. *American Economic Journal: Economic Policy*, 15(4):183–221.

Cellini, S. R., Ferreira, F., and Rothstein, J. (2010). The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design. *Quarterly Journal of Economics*, 125(1):215–261.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The Effect of Minimum Wages on Low-Wage Jobs. *Quarterly Journal of Economics*, 134(3):1405–1454.

Chen, B., Lu, M., and Zhong, N. (2015). How Urban Segregation Distorts Chinese Migrants' Consumption? *World Development*, 70:133–146.

Chetty, R., Friedman, J. N., Hendren, N., Jones, M. R., and Porter, S. R. (2026). The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility. *American Economic Review*, 116(1):1–51.

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

Chetty, R. and Hendren, N. (2018b). The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates. *Quarterly Journal of Economics*, 133(3):1163–1228.

Chetty, R., Hendren, N., and Katz, L. F. (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.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States. *Quarterly Journal of Economics*, 129(4):1553–1623.

Chintalapati, R., Dar, A., and Sood, G. (2024). Pranaam: Predict religion and caste based on name. Software.

Christensen, P. and Timmins, C. (2022). Sorting or Steering: The Effects of Housing Discrimination on Neighborhood Choice. *Journal of Political Economy*, 130(8):2110–2163.

Cutler, D. M., Glaeser, E. L., and Vigdor, J. L. (1999). The Rise and Decline of the American Ghetto. *Journal of Political Economy*, 107(3):455–506.

Department of Posts, Government of India (2026). All India Pincode Directory. Open Government Data Platform India.

Drèze, J. and Kingdon, G. G. (2001). School Participation in Rural India. *Review of Development Economics*, 5(1):1–24.

Election Commission of India (2024a). Data of Number of Electors, 2024. Final publication of electoral roll with January 1, 2024 as the qualifying date.

Election Commission of India (2024b). Theme 7: Polling Station. Election Commission of India training material.

Fishback, P. V., Rose, J. D., Snowden, K. A., and Storrs, T. (2024). New Evidence on Redlining by Federal Housing Programs in the 1930s. *Journal of Urban Economics*, 141:103708.

García-López, M.-À. and Moreno-Monroy, A. I. (2018). Income Segregation in Monocentric and Polycentric Cities: Does Urban Form Really Matter? *Regional Science and Urban Economics*, 71:62–79.

Gardner, J. (2022). Two-Stage Differences in Differences. Working Paper.

Gonzalez-Navarro, M. and Quintana-Domeque, C. (2016). Paving Streets for the Poor: Experimental Analysis of Infrastructure Effects. *Review of Economics and Statistics*, 98(2):254–267.

Government of Gujarat (1991). The Gujarat Prohibition of Transfer of Immovable Property and Provision for Protection of Tenants from Eviction from Premises in Disturbed Areas Act, 1991. Gujarat Act No. 12 of 1991.

Government of Gujarat (2020). The Gujarat Prohibition of Transfer of Immovable Property and Provision for Protection of Tenants from Eviction from Premises in Disturbed Areas (Amendment) Act, 2020. Gujarat Government Gazette Extraordinary, 15 October 2020.

Hanushek, E. A. and Woessmann, L. (2012). Schooling, Educational Achievement, and the Latin American Growth Puzzle. *Journal of Development Economics*, 99(2):497–512.

Hindustan Times (2026). Rajasthan Assembly Passes Disturbed Areas Property Bill amid Opposition Criticism. Published March 6, 2026.

India Justice Report (2025a). Gujarat State Analysis: Police. India Justice Report 2025.

India Justice Report (2025b). Population per Police Station (Rural). India Justice Report 2025. Indicator calculated as rural population divided by rural police stations, using National Commission on Population and BPR&D Data on Police Organizations.

India Justice Report (2025c). Population per Police Station (Urban). India Justice Report 2025. Indicator calculated as urban population divided by urban police stations, using National Commission on Population and BPR&D Data on Police Organizations.

Indian Express (2026). Rajasthan Clears Bill on Property Transfers in 'Disturbed' Areas. Published March 7, 2026.

Jaffrelot, C. (2003). Communal Riots in Gujarat: The State at Risk? *Heidelberg Papers in South Asian and Comparative Politics*, (17).

Jaffrelot, C. (2011). *Religion, Caste, and Politics in India*. C. Hurst & Co., London.

Kok, C. C. and Lim, G. J. (2024). Ethnic Proximity and Politics: Evidence from Colonial Resettlement in Malaysia. SoDa Laboratories Working Paper Series 2024-06, Monash University.

Kuminoff, N. V., Smith, V. K., and Timmins, C. (2013). The New Economics of Equilibrium Sorting and Policy Evaluation Using Housing Markets. *Journal of Economic Literature*, 51(4):1007–1062.

Laliwala, S., Jaffrelot, C., Thakkar, P., and Desai, A. (2021). Paradoxes of Ghettoization: Juhapura in Ahmedabad. Technical report, Centre for Equity Studies.

Lambiri, D. and Vargas, M. (2011). Residential Segregation and Public Housing Policy: The Case of Chile. Working Paper 29, Facultad de Economía y Empresa, Universidad Diego Portales.

Larrañaga Jiménez, O. and Sanhueza Riveros, C. (2007). Residential Segregation Effects on Poor's Opportunities in Chile. Working Paper 259, University of Chile, Department of Economics.

Liao, Y. and Zhang, J. (2021). Hukou Status, Housing Tenure Choice and Wealth Accumulation in Urban China. *China Economic Review*, 68:101638.

Lowe, M. (2021). Types of Contact: A Field Experiment on Collaborative and Adversarial Caste Integration. *American Economic Review*, 111(6):1807–1844.

Marx, B., Stoker, T. M., and Suri, T. (2019). There Is No Free House: Ethnic Patronage in a Kenyan Slum. *American Economic Journal: Applied Economics*, 11(4):36–70.

Massey, D. S. and Denton, N. A. (1988). The Dimensions of Residential Segregation. *Social Forces*, 67(2):281–315.

Menon, S. and Joji, J. (2023). Juhapura: The Making of a Muslim Ghetto in Ahmedabad. *Journal of Social Inclusion Studies*, 9(2):169–183.

Miguel, E. and Gugerty, M. K. (2005). Ethnic Diversity, Social Sanctions, and Public Goods in Kenya. *Journal of Public Economics*, 89(11–12):2325–2368.

Monarrez, T. E. (2023). School Attendance Boundaries and the Segregation of Public Schools in the United States. *American Economic Journal: Applied Economics*, 15(3):210–237.

Muralidharan, K. and Sundararaman, V. (2015). The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India. *Quarterly Journal of Economics*, 130(3):1011–1066.

NDTV (2026). Rajasthan Passes Bill Restricting Property Transfers in 'Disturbed' Areas. Published March 6, 2026.

Office of the Registrar General and Census Commissioner, India (2011a). HH-01: Normal Households by Household Size, Gujarat, Census 2011. Census of India 2011. Reports 12,181,718 households in Gujarat.

Office of the Registrar General and Census Commissioner, India (2011b). Primary Census Abstract by Religion, India, Census 2011. Census of India 2011. PCA REL: Primary Census Abstract by Religion. Table categories include Hindu, Muslim, Christian, Sikh, Buddhist, Jain, Other Religions and Persuasions, and Religion Not Stated.

Press Information Bureau, Government of India (2015). RGI Releases Census 2011 Data on Population by Religious Communities. Release based on Census of India 2011 religion data.

PRS Legislative Research (2022). Annual Review of State Laws 2021.

Rao, G. (2019). Familiarity Does Not Breed Contempt: Generosity, Discrimination, and Diversity in Delhi Schools. *American Economic Review*, 109(3):774–809.

Reardon, S. F. and O'Sullivan, D. (2004). Measures of Spatial Segregation. *Sociological Methodology*, 34(1):121–162.

Rosen, S. (1974). Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, 82(1):34–55.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2):143–186.

Selod, H. and Zenou, Y. (2001). Location and Education in South African Cities under and after Apartheid. *Journal of Urban Economics*, 49(1):168–198.

Shrivastava, D. (2022). Socio-spatial Consequences of Disturbed Areas Act 1991 on Urbanizing Spaces in Gujarat. *Socio-Legal Review*, 18(2):176–199.

Sood, A. and Ehrman-Solberg, K. (2023). The Long Shadow of Housing Discrimination: Evidence from Racial Covenants. Working Paper, Revised and Resubmitted.

Sood, G., Laohaprapanon, S., Sanjeevi, M., and Tran, K. (2018). Parsed Indian Electoral Rolls.

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

Susewind, R. (2016). Data on Religion and Politics in India.

Tejani, S. (2023). Saffron Geographies of Exclusion: The Disturbed Areas Act of Gujarat. *Urban Studies*, 60(4):597–619.

Troesken, W. and Walsh, R. (2019). Collective Action, White Flight, and the Origins of Formal Segregation Laws. *The Journal of Law, Economics, and Organization*, 35(2):289–318.

Wing, C., Freedman, S. M., and Hollingsworth, A. (2024). Stacked Difference-in-Differences. Technical Report 32054, National Bureau of Economic Research.

Wong, M. (2013). Estimating Ethnic Preferences Using Ethnic Housing Quotas in Singapore. *The Review of Economic Studies*, 80(3):1178–1214.

Wong, M. (2014). Estimating the Distortionary Effects of Ethnic Quotas in Singapore Using Housing Transactions. *Journal of Public Economics*, 115:131–145.

SUPPLEMENTAL ONLINE APPENDIX

A Data Appendix

A.1 Construction of the Disturbed Areas Dataset

Table A.1: Construction of the Disturbed Areas Treatment Panel

Component	Construction and role in analysis
Legal unit of treatment	The DAA is applied to notified disturbed areas within police-station jurisdictions, not to broad administrative units such as districts or cities. Treatment is therefore built from locality-level legal designations and then aggregated to polling booths and pincodes.
Official notifications	Gujarat Revenue Department notifications declaring new disturbed areas form the core legal record. Because no consolidated public archive exists, these records were assembled through field visits and legal contacts. Notifications vary from clean legal notices to scanned documents and semi-structured tables; each is standardized into a locality-year file containing notified locality, police-station jurisdiction, city, and year of first notification.
Newspaper supplementation	Newspaper articles on disturbed-area notifications, property-transfer disputes, and related legal developments supplement incomplete official records. A text model first identifies whether an article mentions a disturbed area and extracts the location in detail, including neighborhood, city, and state when available. A second prompt uses the article text and extracted location to recover month-year of designation when available, or year otherwise.
Record consolidation	Repeated locality references across articles are reconciled and merged with the official notification database, producing a unified list of disturbed localities and years of designation. Newspaper evidence recovers missing locality names and dates and provides an independent cross-check on timing and geography.
Geocoding and pincode recovery	Each disturbed locality is submitted to the Google Maps Geocoding API to recover latitude and longitude, including the main point and, where available, additional coordinates from returned bounds or viewport. These coordinates are reverse-geocoded to recover candidate pincodes, creating mapped locality-to-pincode links.
Pincode treatment measures	A pincode is coded as treated in the first year in which any disturbed locality mapped to it appears in the notification or newspaper record. I also construct a time-varying intensity measure equal to the number of distinct disturbed localities associated with each pincode-year. These measures are used for housing, school, and violence outcomes.
Polling-booth treatment	Locality-level designations are also aggregated to polling booths, the fine electoral geography used for residential-composition analysis. Polling booths inherit the treatment status and treatment year of their pincode. A booth is therefore treated beginning in the first year in which its pincode contains any DAA-notified locality.

Notes: Official notifications remain the core legal record. Newspaper evidence supplements missing locations and dates and cross-checks the timing and reach of the law. Geocoding converts locality-level legal designations into the empirical units used in the analysis. Polling booths approximate fine residential catchments; pincodes are broader postal units used for housing, school, and violence outcomes.

A.2 2025 Electoral Rolls and Religious Classification

Table A.2: Construction of 2025 Electoral-Roll Religious Composition Measures

Step	Construction and role in analysis
Raw source	Public Gujarat electoral-roll PDFs released by the Election Commission of India. The rolls contain full polling-booth voter lists, including voter names, household relationships, and location information, and form the universe of 2025 voter records used to construct neighborhood religious composition.
OCR, cleaning, and parsing	Scanned Gujarati-script PDFs are converted into text, cleaned to remove formatting artifacts, normalize spacing and character encoding, handle invisible characters, and correct common OCR errors. Each voter entry is then parsed to isolate the voter name and, when available, the guardian or household-head name.
Transliteration and tokenization	Gujarati-script names are transliterated into a standardized Latin representation and decomposed into meaningful tokens, typically given names and surnames. Honorifics and non-informative suffixes are removed; final name tokens receive particular attention because they are more likely to encode caste or community identity.
Dictionary construction	The classification dictionary combines external Gujarati name and surname sources, public lists of Muslim names, and high-frequency surname tokens from the electoral data that are manually reviewed. Systematic spelling variants, suffix modifications, and transliteration variants are incorporated, producing token-level Muslim and Hindu classifications with reliability metadata.
Ambiguity handling	Ambiguous tokens and tokens appearing across multiple communities are flagged and treated conservatively. Short, ambiguous, or cross-community tokens are excluded or prevented from driving classification, reducing false classification from weak name evidence.
Hierarchical classification	Names are scanned first for high-confidence Muslim-associated tokens in the voter name or associated household-member name. If no Muslim token is found, final name components are searched for Hindu surname tokens. If neither source provides evidence, the individual is conservatively classified as Hindu.
Aggregation	Individual classifications are aggregated to polling booths and mapped to pincodes using geographic identifiers and administrative boundary files, yielding polling booth-year and pincode-year counts of total, Muslim, and Hindu voters for the segregation analysis.

Notes: The classification is intentionally conservative: high-confidence Muslim tokens identify Muslim voters; Hindu surname tokens identify Hindu voters only when no Muslim token is present; unresolved cases are classified as Hindu. This rule prioritizes minimizing false positives in Muslim identification.

A.3 Additional Data Construction

A.3.1 Residential Segregation Measures

Let i index a local unit, t index year, and M_{it}, H_{it}, and O_{it} denote Muslim, Hindu, and other voters. Let $P_{it} \equiv M_{it} + H_{it} + O_{it}$, and let $M_t \equiv \sum_i M_{it}$.

Table A.3: Residential Composition and Segregation Measures

Measure	Definition	Interpretation
Muslim share	M_{it}/P_{it}	Local Muslim presence in the unit.
Hindu share	H_{it}/P_{it}	Local Hindu presence. This is not mechanically the complement of Muslim share because other religious groups are present.
Log Muslim-to-Hindu ratio	$\log\{(M_{it}+c)/(H_{it}+c)\}$, with $c > 0$	Relative Muslim-Hindu balance. Positive values indicate Muslim concentration relative to Hindus; negative values indicate Hindu concentration.
Mixed-neighborhood index	$\sum_{j \neq i} w_{ij} \min\{M_{jt}/P_{jt}, H_{jt}/P_{jt}\}$, with $\sum_{j \neq i} w_{ij} = 1$	Exposure-weighted measure of how mixed nearby units are. Higher values indicate more interspersed Hindu-Muslim surroundings.
Isolation contribution	$(M_{it}/M_t)(M_{it}/P_{it})$	Booth-level contribution to pincode-year Muslim isolation. It is larger when a booth contains both a larger share of local Muslims in the pincode and a higher Muslim share internally.

Notes: For the mixed-neighborhood index, w_{ij} denotes spatial weights over nearby units. The isolation contribution sums across booths within a pincode-year to yield the pincode-year Muslim isolation index.

A.3.2 Housing Market Measures

Table A.4: Housing Market Measures

Measure	Definition	Interpretation / construction note
Total units launched	$\sum_{p \in \mathcal{P}_{it}} \text{Units}_{pit}$	Total number of housing units in projects launched in pincode i and year t.
Price per square foot	$\text{Sales}_{pit}/\text{BuiltArea}_{pit}$	Project sales divided by built area. The measure is rescaled by Rs. 10,000 for estimation and presentation.
Land cost per square foot	$\text{LandCost}_{pit}/\text{BuiltArea}_{pit}$	Land cost divided by built area, also rescaled by Rs. 10,000.
Unit density	$\text{TotalUnits}_{it}/\text{Area}_{i}$	Housing units launched per square kilometer of pincode area.
Built-space congestion	$\sum_{p \in \mathcal{P}_{it}} \text{BuiltArea}_{pit}/\text{PincodeArea}_{i}$	Total project built area as a share of pincode land area in square feet.

Notes: p indexes projects, i pincodes, t years, and $\mathcal{P}_{it}$ the set of projects in pincode i beginning in year t. Housing price and cost outcomes use both pincode-year means across projects and aggregate pincode-year ratios formed from sums of numerators and denominators.

A.3.3 Public Safety Measures

Table A.5: GDELT Public Safety Measures

Measure	Definition	Construction note
GDELT: Communal Conflict	$\sum_{g \in \mathcal{G}_{it}} \mathbf{1}\{\text{event } g \text{ is communal conflict}\}$	Events involving confrontation between religious, ethnic, or other identity-based groups. Classification uses actor identities together with event narratives indicating inter-group clashes, riots, mob violence, or related communal confrontation.
GDELT: Civilian Targeting	$\sum_{g \in \mathcal{G}_{it}} \mathbf{1}\{\text{event } g \text{ targets civilians}\}$	Events in which civilians are identified as the direct target of violence, threats, assault, arson, or other coercive action. Classification combines civilian actor labels with the direction of the reported interaction and information contained in the event narrative.
GDELT: Total Fatalities	$\sum_{g \in \mathcal{G}_{it}} Fatalities_g$	Sum of reported event-level fatalities in a pincode-year.

Notes: $\mathcal{G}_{it}$ denotes the set of GDELT events assigned to pincode i in year t. Gujarat events are reverse-geocoded to pincodes. Event categories are constructed using classified actor labels, the direction of interaction between actors, and information contained in event narratives. Actor classifications distinguish Indian state forces, foreign state forces, Indian civilians, foreign civilians, and other non-state actors. The three categories need not be mutually exclusive.

A.3.4 School Inputs and Outcomes

Table A.6: School Input and Outcome Measures

Measure	Definition	Interpretation / construction note
Grants received per student	$\text{GrantsReceived}_{sit}/N_{sit}$	School grants received divided by total enrollment.
Grants spent per student	$\text{GrantsSpent}_{sit}/N_{sit}$	School grants spent divided by total enrollment.
Government-school share	$\text{GovtSchools}_{it}/\text{Schools}_{it}$	Share of schools in a pincode-year that are government schools.
Private-school share	$\text{PrivateSchools}_{it}/\text{Schools}_{it}$	Share of schools in a pincode-year that are private schools.
School age	$t - \text{EstYear}_s$	Years since school establishment.
Infrastructure index	$K^{-1}\sum_{k=1}^{K} Z_{sit}^{(k)}$, where $Z_{sit}^{(k)}$ is year-standardized	Average of year-standardized infrastructure components, including electricity, drinking water, toilets, classroom quality, computers, ICT lab, library, playground, medical check-ups, and ramp access.
Stage-specific failure rate	$\#\ \text{Fails}_{sit}/\#\ \text{Enrolled}_{sit}$	Exam failures divided by enrollment at the primary, secondary, or higher-secondary stage.

Notes: s indexes schools, i pincodes, t years, and N_{sit} total enrollment. The infrastructure index averages year-standardized components. School-year outcomes are aggregated to the pincode-year level using sums for counts and ratios or means as appropriate for each outcome.

B Results

B.1 Residential Segregation

Table A.7: Polling-Booth Cohort Comparisons

	Muslim share	Hindu share	Mixed-neighborhood index	Isolation contribution
Late-treated relative to early-treated	-0.090***	0.329***	-0.091***	-0.004***
	(0.019)	(0.015)	(0.012)	(0.001)
Adjusted R^2	0.134	0.577	0.147	0.024
Observations	8,463	8,463	8,463	8,463

Notes: The unit of observation is the polling booth. The table compares baseline neighborhood characteristics of early-treated booths, first exposed between 2014 and 2017, to late-treated booths, first exposed between 2017 and 2025. The omitted group is early-treated booths, so positive coefficients indicate that late-treated booths have higher baseline values of the corresponding characteristic. Isolation contribution reports each booth's contribution to within-pincode Muslim isolation. Standard errors, reported in parentheses, are clustered at the pincode level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.8: Baseline Characteristics of Late-Treated and Never-Treated Polling Booths

	Mean	SD	Median	N
Late-treated booths (first exposed 2017–2025)				
Muslim share	0.018	0.055	0.002	5,939
Hindu share	0.982	0.055	0.998	5,939
Mixed-neighborhood index	0.033	0.091	0.004	5,939
Never treated by 2025				
Muslim share	0.013	0.043	0.000	26,803
Hindu share	0.987	0.043	1.000	26,803
Mixed-neighborhood index	0.025	0.079	0.000	26,803

Notes: The unit of observation is the polling booth. The table reports baseline characteristics for booths first exposed to the Disturbed Areas Act between 2017 and 2025 and for booths never treated by 2025.

Table A.9: Stacked Difference-in-Differences Estimates with Direct Treatment and Spillover Exposure

	(1) Muslim share	(2) Hindu share	(3) Log Muslim-to-Hindu ratio	(4) Isolation contribution
DAA Exposure × Post	-0.015	-0.062***	0.987***	0.003***
	(0.015)	(0.016)	(0.364)	(0.001)
Spillover Exposure × Post	0.024***	-0.062***	-0.027	0.002***
	(0.007)	(0.016)	(0.312)	(0.001)
Stack × Pincode FE	Yes	Yes	Yes	Yes
Stack × Year FE	Yes	Yes	Yes	Yes
Dependent Variable Mean	0.040	0.864	-6.196	0.004
Adjusted R^2	0.218	0.691	0.527	0.215
Observations	77,318	77,318	77,318	76,824

Notes: The unit of observation is the polling booth-year. The estimates are from a stacked difference-in-differences design that pools cohort-specific treatment windows while preserving relative event time. Treatment and spillover exposure are assigned at the pincode level. The spillover group consists of untreated pincodes entering the 0–5 kilometer spillover ring within each stack window. *Post* is an indicator for the post-designation period within each stack window. Muslim share and Hindu share do not necessarily sum to one because some voters remain outside the Hindu/Muslim classification in the underlying data. Isolation contribution reports each booth's contribution to within-pincode Muslim isolation. All specifications include stack-by-pincode and stack-by-year fixed effects. Standard errors, reported in parentheses, are clustered at the pincode level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

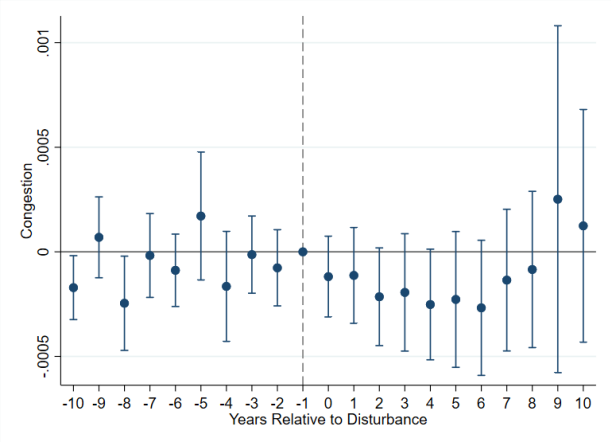
Table A.10: Stacked Difference-in-Differences Estimates for Booth-Level Religious Composition with Baseline DAA Intensity and Spillover Exposure

	(1)	(2)	(3)	(4)
	Muslim share	Hindu share	Log Muslim-to-Hindu ratio	Isolation contribution
DAA Exposure × Post	-0.026	-0.051*	1.083*	0.002**
	(0.027)	(0.026)	(0.568)	(0.001)
DAA Exposure × Post × High Intensity	0.023	-0.023	-0.197	0.001
	(0.028)	(0.030)	(0.695)	(0.001)
Spillover Exposure × Post	0.024***	-0.062***	-0.026	0.002***
	(0.007)	(0.016)	(0.312)	(0.001)
Stack × Pincode FE	Yes	Yes	Yes	Yes
Stack × Year FE	Yes	Yes	Yes	Yes
Dependent Variable Mean	0.040	0.864	-6.196	0.004
Adjusted R^2	0.218	0.691	0.527	0.215
Observations	77,318	77,318	77,318	76,824

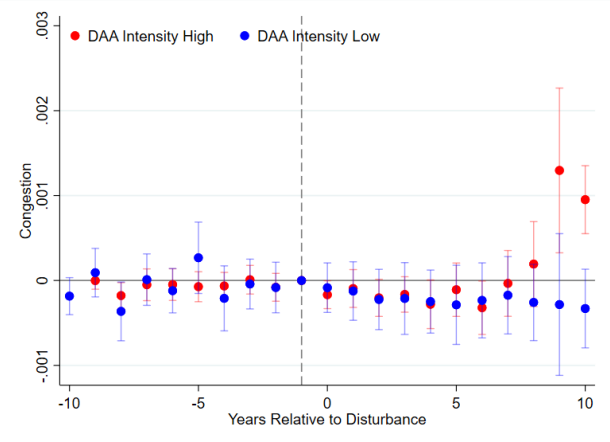
Notes: The unit of observation is the polling booth-year. The estimates are from a stacked treatment-and-spillover difference-in-differences design that pools the 2014–2017 and 2017–2025 cohort-specific windows. Treatment and spillover exposure are assigned at the pincode level. *DAA Exposure × Post* reports the treatment effect for low-intensity treated pincodes. *DAA Exposure × Post × High Intensity* reports the incremental treatment effect for high-intensity treated pincodes relative to low-intensity treated pincodes. *Spillover Exposure × Post* reports the spillover effect for nearby untreated booths. Baseline intensity is the number of distinct notified localities in the pincode when it first enters the DAA regime and is fixed thereafter. High-intensity pincodes are those above the median of this baseline distribution among ever-treated pincodes. Muslim share and Hindu share do not necessarily sum to one because some voters remain outside the Hindu/Muslim classification in the underlying data. Isolation contribution reports each booth's contribution to within-pincode Muslim isolation. All specifications include stack-by-pincode and stack-by-year fixed effects. Standard errors, reported in parentheses, are clustered at the pincode level. The implied treatment effect for high-intensity treated pincodes is the sum of the two direct-effect coefficients. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B.2 Housing Markets

Figure A.1: DAA Exposure and Housing Congestion

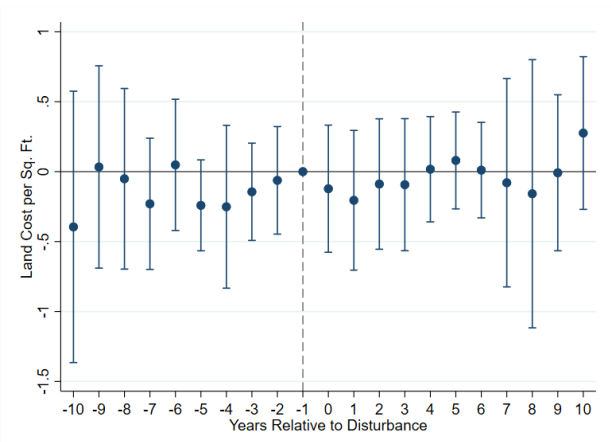


(a) Congestion

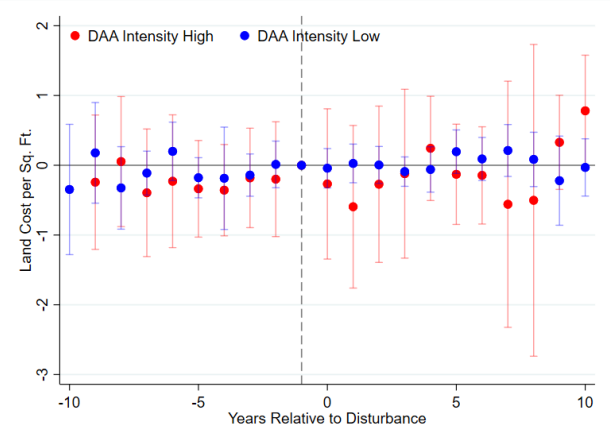


(b) Congestion by Baseline DAA Intensity

Figure A.2: DAA Exposure and Land Cost per Sq. Ft. (in Rs. 10,000)

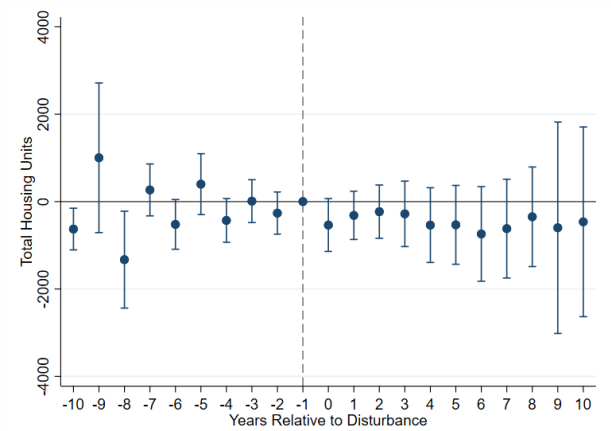


(a) Land Cost per Sq. Ft.

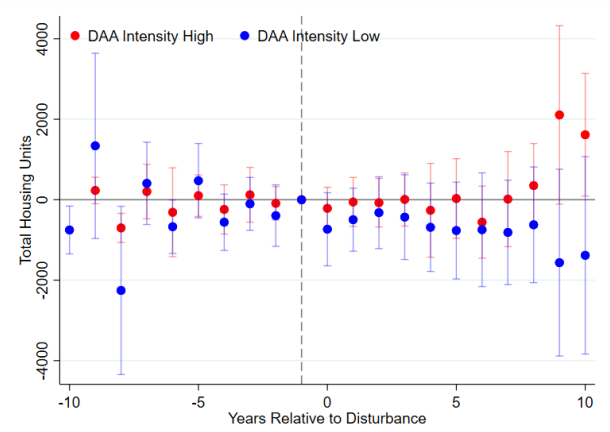


(b) Land Cost per Sq. Ft. by Baseline DAA Intensity

Figure A.3: DAA Exposure and Total Housing Units

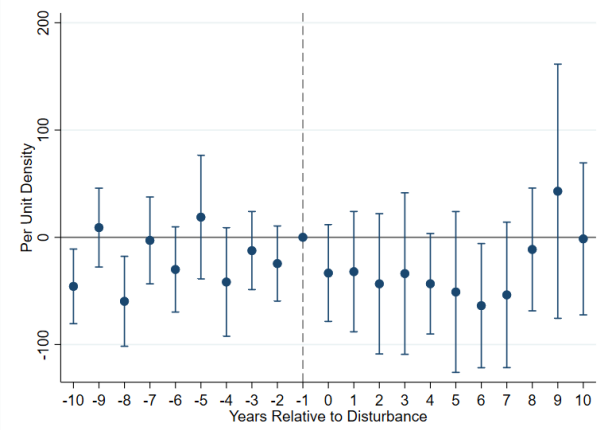


(a) Total Housing Units

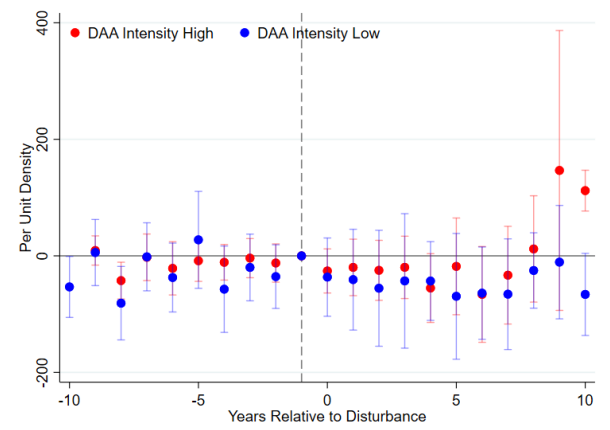


(b) Total Housing Units by Baseline DAA Intensity

Figure A.4: DAA Exposure and Unit Density



(a) Unit Density

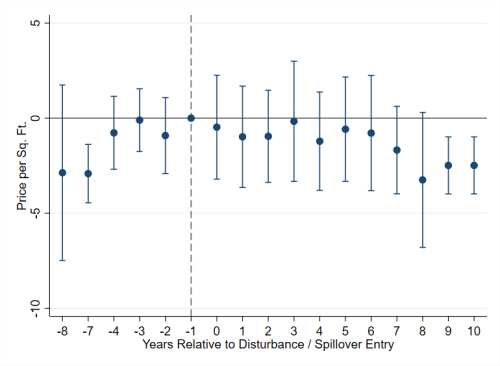


(b) Unit Density by Baseline DAA Intensity

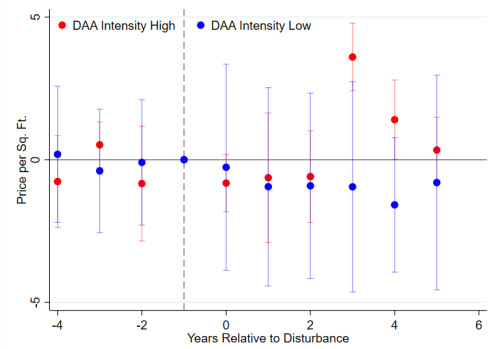
Notes: The figures above report pooled and intensity-split event-study estimates for each housing market attribute. Panel (a) uses not-yet-treated pincodes as controls. Panel (b) splits treated pincodes by baseline DAA intensity. The unit of observation is the pincode-year. In both panels, estimates are obtained using the Callaway and Sant'Anna estimator with not-yet-treated pincodes as controls. The omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

B.2.1 Spillover Robustness

Figure A.5: DAA Exposure and Price per Sq. Ft. (in Rs. 10,000)

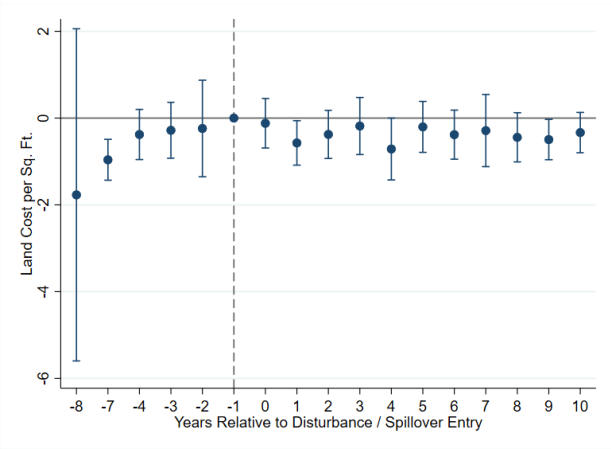


(a) Price per Sq. Ft.

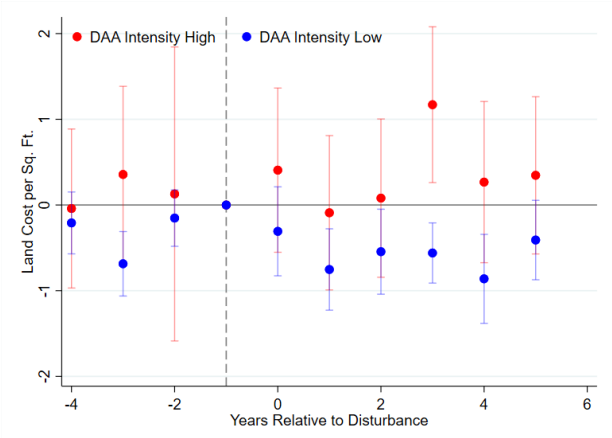


(b) Price per Sq. Ft. by Baseline DAA Intensity

Figure A.6: DAA Exposure and Land Cost per Sq. Ft. (in Rs. 10,000)

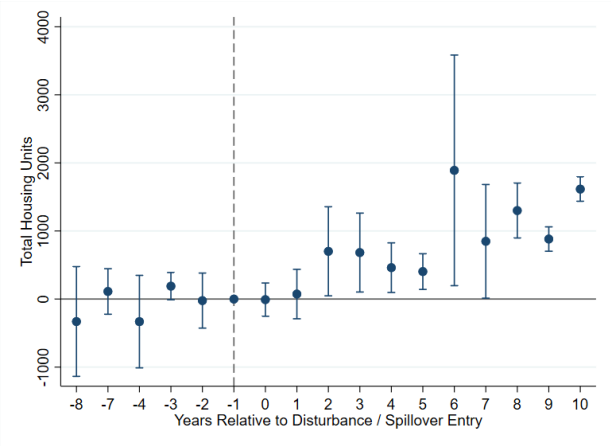


(a) Land Cost per Sq. Ft.

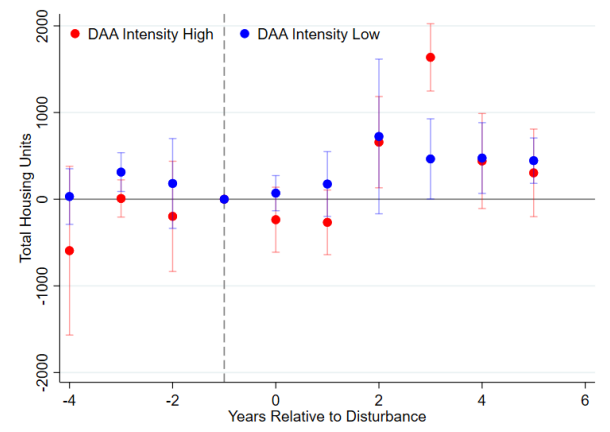


(b) Land Cost per Sq. Ft. by Baseline DAA Intensity

Figure A.7: DAA Exposure and Total Housing Units



(a) Total Housing Units

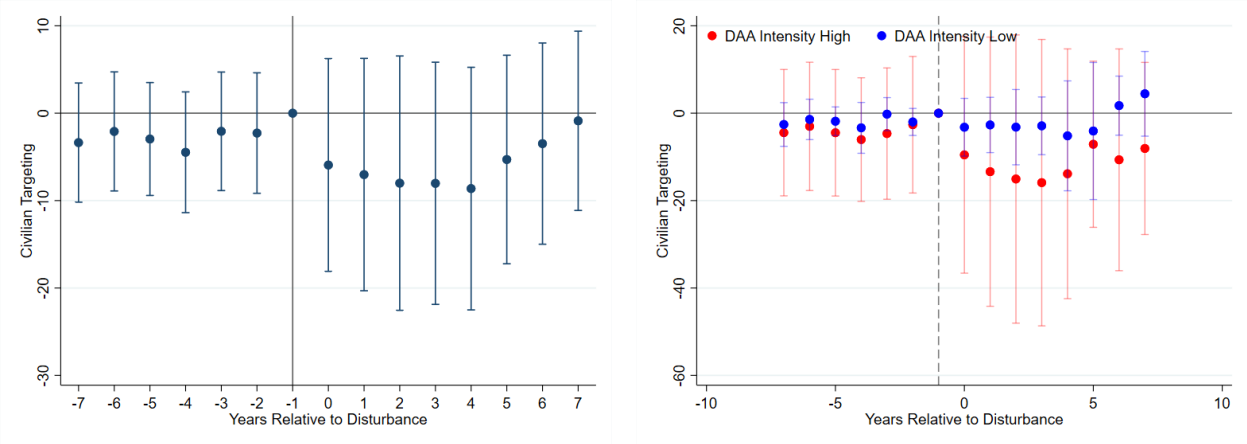


(b) Total Housing Units by Baseline DAA Intensity

Notes: The figures above report pooled and intensity-split event-study estimates for different housing market attributes under the donut-sample design. Panel (a) uses untreated pincodes outside the 0–5 kilometer ring around active DAA pincodes as controls. Panel (b) splits treated pincodes by baseline DAA intensity. In both panels, estimates are obtained using the Callaway and Sant'Anna estimator under the donut-sample design, with not-yet-treated pincodes outside the 0–5 kilometer ring around active DAA pincodes serving as controls. The unit of observation is the pincode-year. The omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

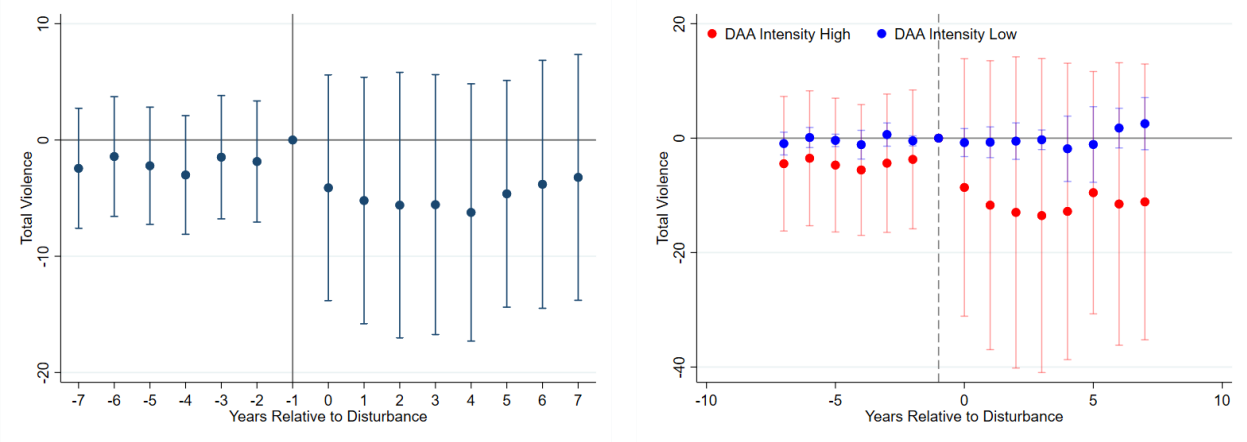
B.3 Public Safety

Figure A.8: DAA Exposure and Civilian Targeting: GDELT



(a) Civilian Targeting (b) Civilian Targeting by Baseline DAA Intensity

Figure A.9: DAA Exposure and Total Fatalities: GDELT

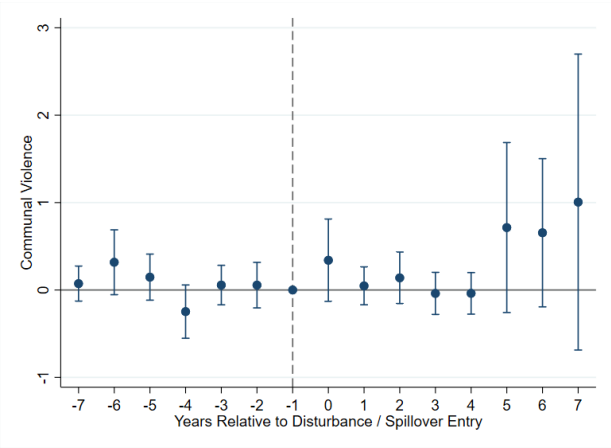


(a) Total Fatalities (b) Total Fatalities by Baseline DAA Intensity

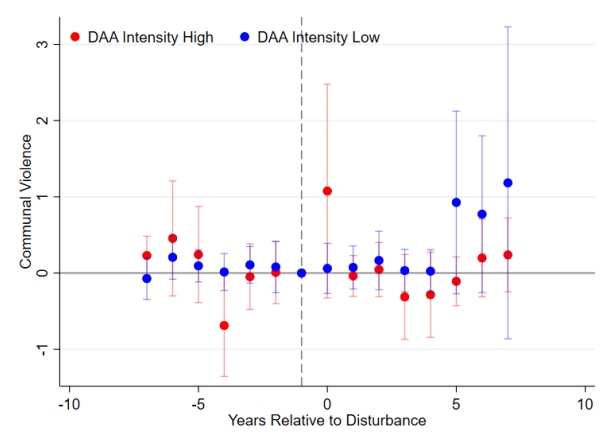
Notes: The figure reports pooled and intensity-split event-study estimates for public safety incidents using GDELT data. Panel (a) reports pooled event-study estimates. Panel (b) reports the corresponding estimates separately for high- and low-intensity treated pincodes. In both panels, estimates are obtained using the Callaway and Sant'Anna estimator with not-yet-treated pincodes as controls. The unit of observation is the pincode-year. The omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

B.3.1 Spillover Robustness

Figure A.10: DAA Exposure and Communal Violence Incidence: GDELT

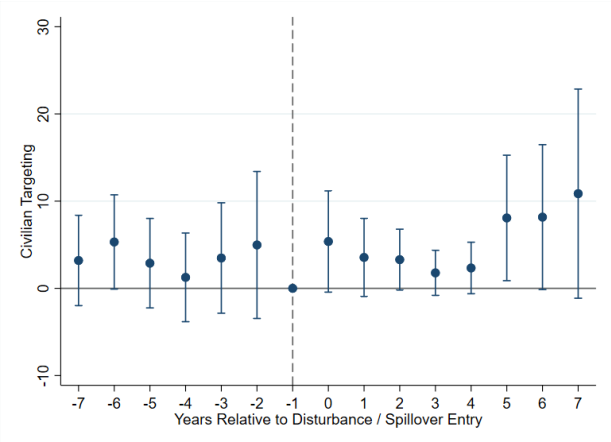


(a) Communal Violence Incidence

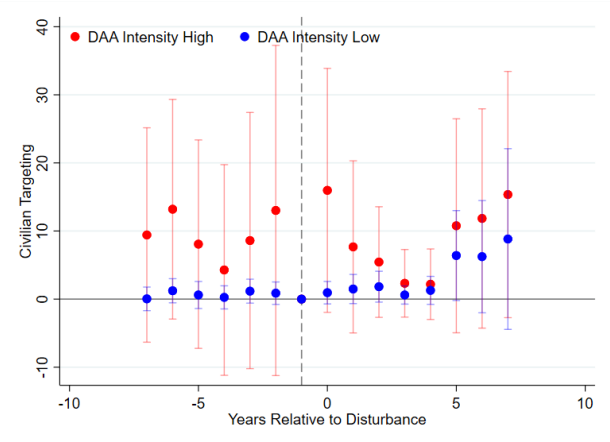


(b) Communal Violence Incidence by Baseline DAA Intensity

Figure A.11: DAA Exposure and Civilian Targeting: GDELT

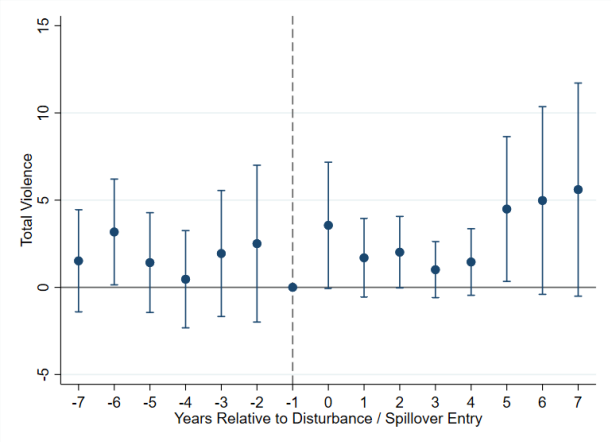


(a) Civilian Targeting

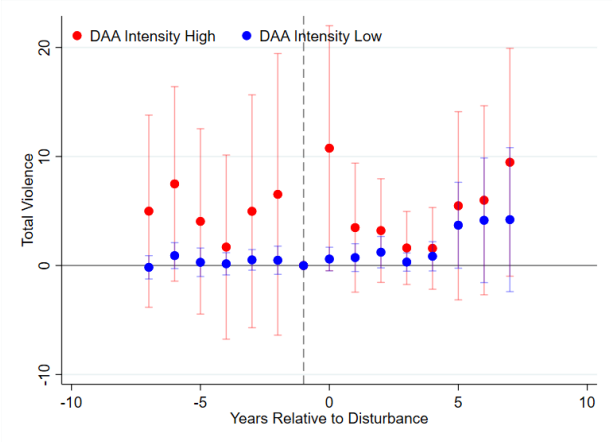


(b) Civilian Targeting by Baseline DAA Intensity

Figure A.12: DAA Exposure and Total Fatalities: GDELT



(a) Total Fatalities

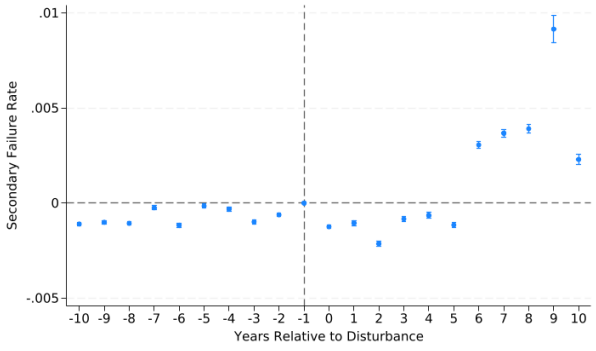


(b) Total Fatalities by Baseline DAA Intensity

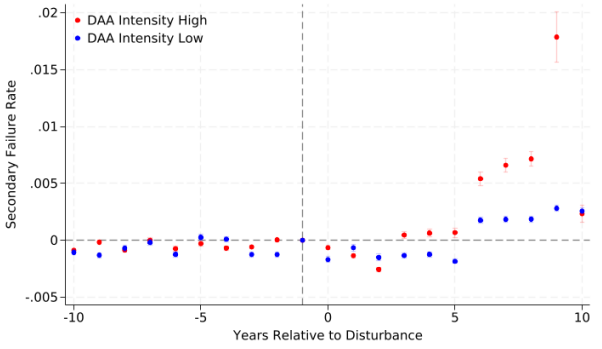
Notes: The figure reports pooled and intensity-split event-study estimates for public safety incidents using GDELT data under the donut-sample design. Panel (a) uses untreated pincodes outside the 0–5 kilometer ring around active DAA pincodes as controls. Panel (b) reports the corresponding estimates separately for high- and low-intensity treated pincodes. In both panels, estimates are obtained using the Callaway and Sant'Anna estimator under the donut-sample design, with not-yet-treated pincodes outside the 0–5 kilometer ring around active DAA pincodes serving as controls. The unit of observation is the pincode-year. The omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

B.4 School Inputs and Outcomes

Figure A.13: DAA Exposure and Secondary-School Failure Rates

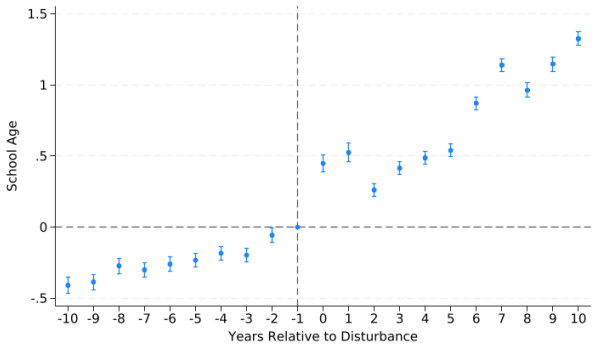


(a) Secondary-School Failure Rate

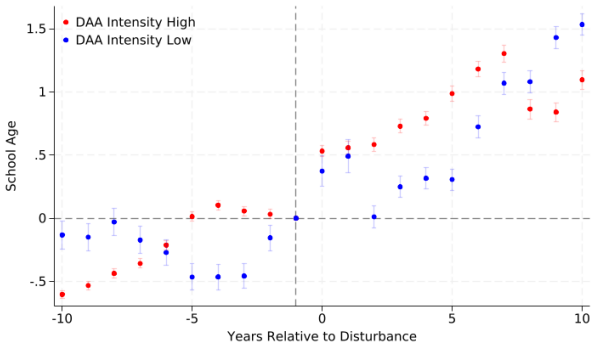


(b) Secondary-School Failure Rate by Baseline DAA Intensity

Figure A.14: DAA Exposure and School Age

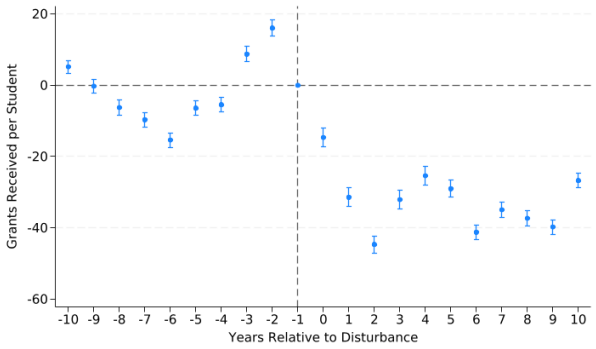


(a) School Age

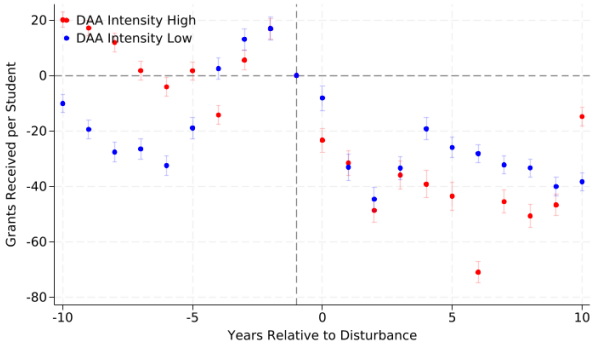


(b) School Age by Baseline DAA Intensity

Figure A.15: DAA Exposure and Grants Received per Student

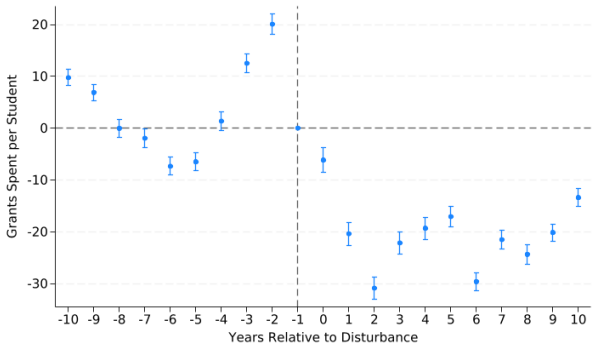


(a) Grants Received per Student

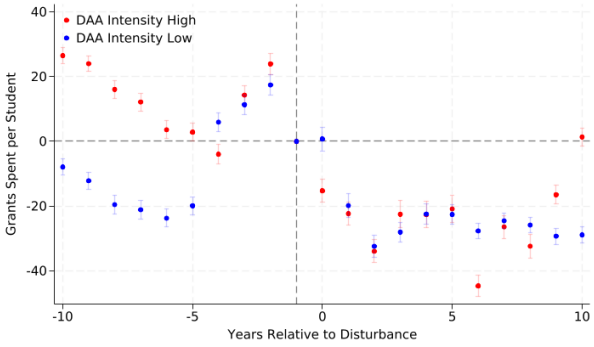


(b) Grants Received per Student by Baseline DAA Intensity

Figure A.16: DAA Exposure and Grants Spent per Student

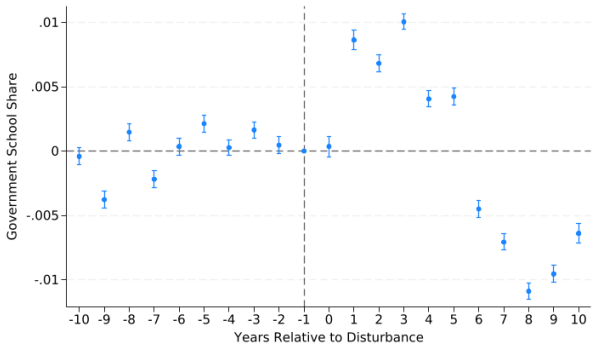


(a) Grants Spent per Student

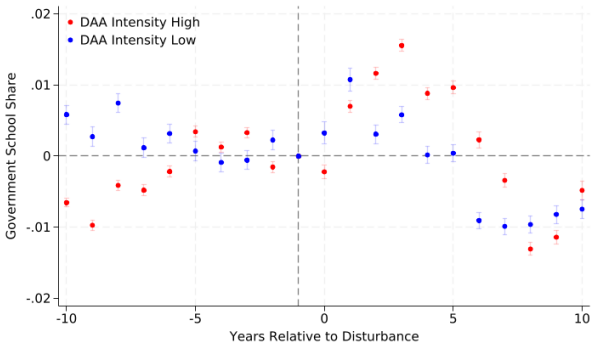


(b) Grants Spent per Student by Baseline DAA Intensity

Figure A.17: DAA Exposure and the Government-School Share

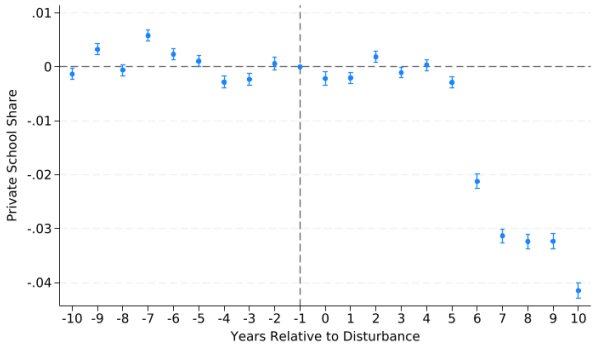


(a) Government-School Share

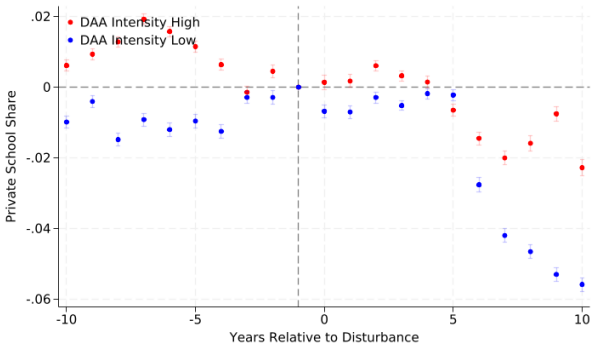


(b) Government-School Share by Baseline DAA Intensity

Figure A.18: DAA Exposure and the Private-School Share



(a) Private-School Share

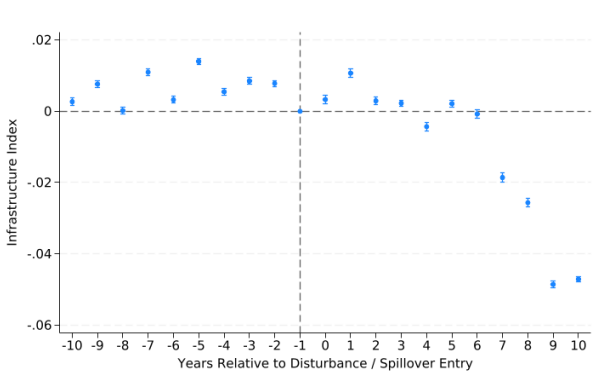


(b) Private-School Share by Baseline DAA Intensity

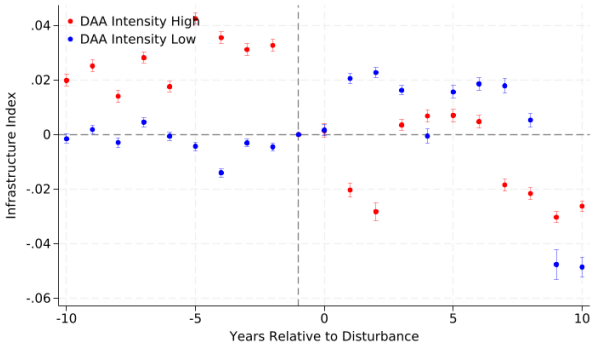
Notes: The figures above report event-study estimates for the different school-related attributes. Panel (a) shows pooled estimates. Panel (b) reports estimates separately for high- and low-intensity treated pincodes. In both panels, estimates are obtained using the Sun and Abraham estimator with never-treated pincodes as controls. The unit of observation is the pincode-year. Event time is defined relative to first DAA exposure, and the omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

B.4.1 Spillover Robustness

Figure A.19: Design A Spillover Robustness: Infrastructure

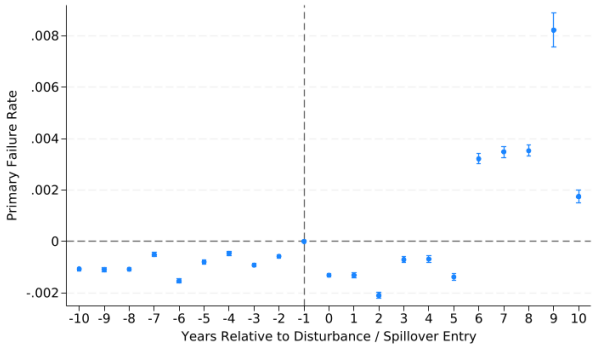


(a) Infrastructure Index

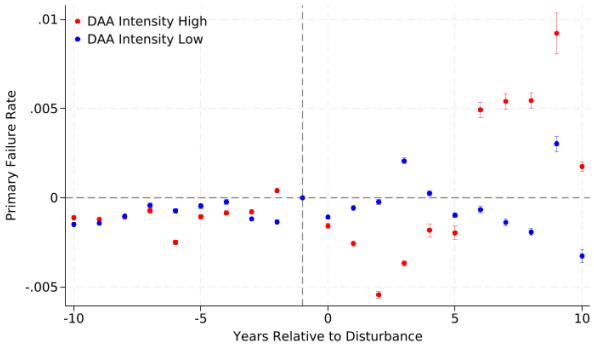


(b) Infrastructure Index by Baseline DAA Intensity

Figure A.20: Design A Spillover Robustness: Primary School Exam Failure Rates

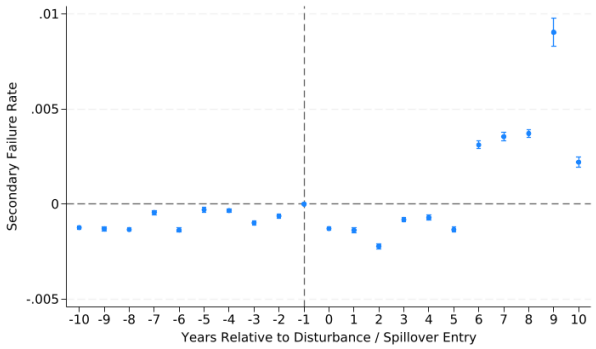


(a) Primary School Exam Failure Rate

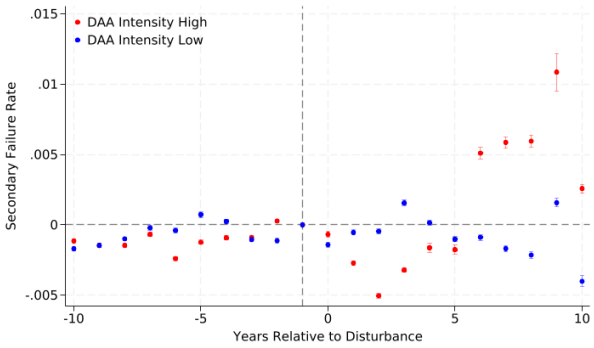


(b) Primary School Exam Failure Rate by Baseline DAA Intensity

Figure A.21: Design A Spillover Robustness: Secondary School Exam Failure Rates



(a) Secondary School Exam Failure Rate

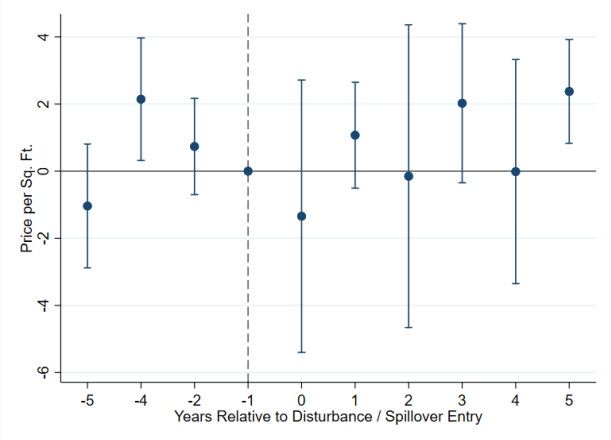


(b) Secondary School Exam Failure Rate by Baseline DAA Intensity

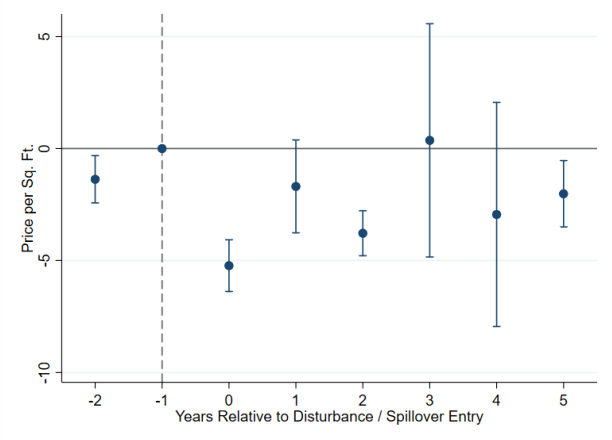
Notes: The figures above report pooled and intensity-split event-study estimates for the different school-related attributes under the donut-sample design. Panel (a) uses untreated pincodes outside the 0–5 kilometer ring around active DAA pincodes as controls. Panel (b) reports the corresponding estimates separately for high- and low-intensity treated pincodes. In both panels, estimates are obtained using the Sun and Abraham estimator under the donut-sample design, with never-treated pincodes outside the 0–5 kilometer ring around active DAA pincodes serving as controls. The unit of observation is the pincode-year. The omitted period is the year immediately before treatment. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.

B.5 Distance-Based Spillover Effects in Untreated Areas

Figure A.22: Distance-Based Spillover Effects on Price per Sq. Ft. (in Rs. 10,000) in Untreated Areas

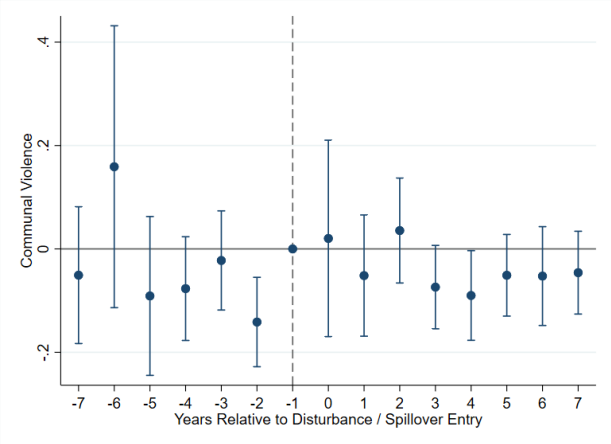


(a) Price per Sq. Ft. (0–5 km ring)

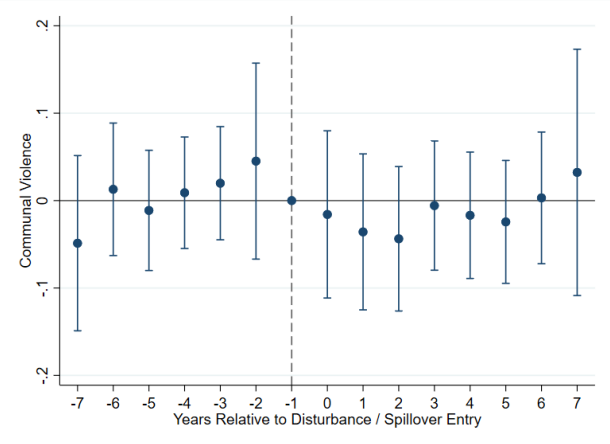


(b) Price per Sq. Ft. (5–10 km ring)

Figure A.23: Distance-Based Spillover Effects on Communal Violence in Untreated Areas: GDELT

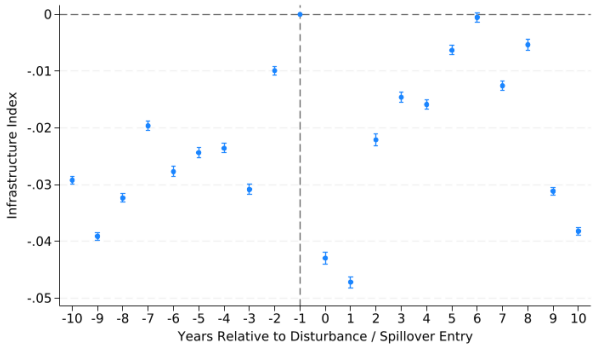


(a) Communal Violence (0–5 km ring)

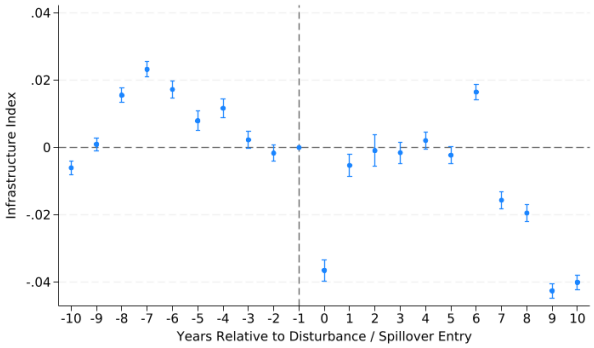


(b) Communal Violence (5–10 km ring)

Figure A.24: Distance-Based Spillover Effects on School Infrastructure Index in Untreated Areas

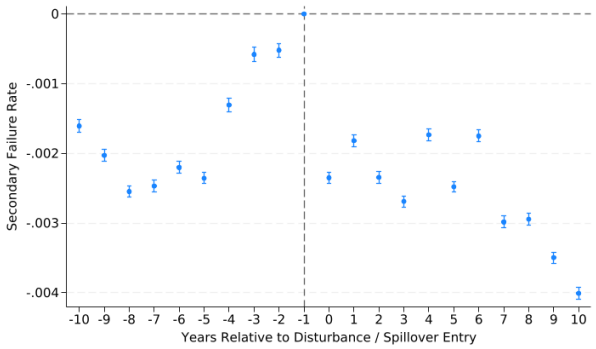


(a) Infrastructure Index (0–5 km ring)

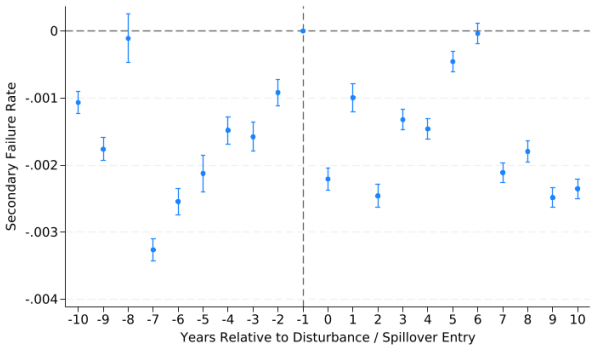


(b) Infrastructure Index (5–10 km ring)

Figure A.25: Distance-Based Spillover Effects on Secondary School Exam Failure Rate in Untreated Areas



(a) Secondary School Exam Failure Rate (0–5 km ring)



(b) Secondary School Exam Failure Rate (5–10 km ring)

Notes: The figures above report event-study estimates of spillover exposure on different measures of neighborhood amenities in untreated pincodes. Panel (a) uses untreated pincodes that first enter the 0–5 km ring around active DAA pincodes as the spillover cohort. Panel (b) uses untreated pincodes that first enter the 5–10 km ring. In both panels, estimates are obtained using the Callaway and Sant'Anna estimator for housing/public safety outcomes or Sun and Abraham estimator for school-related outcomes, with untreated pincodes that not yet (housing/public safety) or never (schooling) enter the relevant exposure ring during the sample period serving as controls. The unit of observation is the pincode-year. Event time is defined relative to first entry into the relevant spillover ring, and the omitted period is the year immediately before spillover entry. Vertical bars denote 95 percent confidence intervals, with standard errors clustered at the pincode level.