Who Bears the Cost of School Closures? Evidence from Puerto Rico

Reza Moradi*

March 2026

Abstract

Declining child populations have prompted large-scale school consolidation worldwide, yet little is known about the welfare distribution of closures or whether alternative closure rules could improve outcomes. I study Puerto Rico's consolidation program (2010–2018), in which roughly 45 percent of public schools closed—the largest such episode in recent U.S. history. Using a matched event study design, I first document that closures cause persistent enrollment declines concentrated among rural students, with near-zero average test score effects that mask large heterogeneity: students from underperforming schools gain up to 0.4 s.d. in math, while those from overperforming schools lose. I then estimate a logit demand model using novel closure-based instruments that exploit spatial reallocation of displaced students. The preferred IV estimate implies that doubling distance reduces market share by approximately 25 percent, with distance dominating quality in parental choices. Finally, I compare five counterfactual closure rules holding total closures fixed. The actual decisions outperform most alternatives when weighted by affected students, suggesting that the geography of closures matters as much as which schools are selected.

JEL Codes: I21, R12, H75

Keywords: school consolidation, demand estimation, discrete choice, instrumental variables, welfare analysis

Acknowledgements. I am deeply grateful to my supervisors, Gustavo Bobonis and Stephani Heblich, for their guidance and support. I also thank my committee member, Robert McMillan, as well as Victor Aguirregabiria, Eduardo Souza-Rodrigues, and participants in the University of Toronto Applied Micro and IO seminars for valuable comments and

*Department of Economics, University of Toronto. Email: reza.moradi@mail.utoronto.ca

suggestions. The administrative data used in this study were provided by the Puerto Rico Department of Education under a data use agreement. This research was supported by the University of Toronto Department of Economics, School of Graduate Studies, and Ontario Graduate Scholarship (OGS). I have no financial relationships or other potential conflicts of interest to disclose. All errors are my own.

1 Introduction

Around the world, declining fertility and outmigration have left public school systems with far more capacity than they need. Between 2000 and 2020, the school-age population fell by double-digit percentages across much of East Asia, Southern Europe, and the Caribbean. Governments have responded with consolidation programs that shutter underutilized schools and redirect students to the remaining facilities. In the United States alone, thousands of public schools have closed over the past two decades, with major rounds in Chicago, Philadelphia, Detroit, and Washington, D.C. Similar programs have been carried out in China, India, Denmark, Scotland, and South Korea.

The rationale for consolidation is straightforward: maintaining half-empty buildings is expensive, and larger schools may offer richer curricula and better peer environments. But consolidation also imposes costs on the demand side. When a school closes, its students must travel farther to attend an alternative, and the school they end up at may differ in quality, peer composition, and other attributes from the one they left. These costs are not distributed evenly. In settings where schools are spatially dispersed—as they are in rural and semi-rural areas—the distance burden of closure can be severe, and families with fewer resources may be least equipped to absorb it.

Despite the scale and policy importance of consolidation, the empirical literature on its welfare consequences remains limited. Existing work has documented average effects on test scores (typically small or null) and some evidence of differential impacts by student characteristics, but has not connected these reduced-form patterns to a structural framework capable of evaluating alternative closure rules. Without a model of how families choose among schools, it is difficult to quantify the welfare losses from closures or to assess whether a different set of schools could have been closed at lower cost to students.

This paper addresses these gaps in the context of Puerto Rico, which experienced the largest school consolidation in recent U.S. history. Between 2010 and 2018, the Puerto Rico Department of Education (PRDE) closed approximately 688 public schools—roughly 45 percent of the total—in response to a severe demographic decline, a fiscal crisis, and the destruction wrought by Hurricane María in 2017. The scale of this consolidation, combined with detailed administrative data covering the universe of public school students, makes Puerto Rico an unusually powerful setting for studying the distributional consequences of closure policies.

I make three contributions. First, using a matched event study design with an interaction-weighted estimator robust to staggered treatment timing, I document the causal effects of school closures on student enrollment and academic achievement. Displacement re-

duces the probability of remaining in the public system by 2–4 percentage points, with the effect concentrated among rural students and persisting for at least three years. Average test score effects are near zero, but this null masks economically meaningful heterogeneity: students displaced from schools that were underperforming relative to nearby alternatives gain up to 0.4 standard deviations in math, while those from overperforming schools experience losses of similar magnitude. Shorter displacement distances and lower concentrations of displaced students at receiving schools are associated with better outcomes. These patterns establish that the match between displaced students and the characteristics of nearby alternatives—not displacement per se—drives the achievement consequences of closure.

Second, I estimate a discrete choice model of school demand in which families choose among public schools in their municipality-grade-year market, with the option to exit to private school. The key challenge is that distance to school is endogenous: residential sorting on unobserved school quality generates a correlation between distance and the demand unobservable. I address this using novel closure-based instruments that exploit the spatial structure of school closures. When a school closes, students from that school disproportionately enroll at the nearest remaining alternatives, generating a demand shock whose intensity varies with proximity to the closed school. The instrument measures each school's exposure to this displacement, constructed from school-to-school distances and predetermined enrollment at closing schools. This design follows the logic of shift-share instruments (Bartik, 1991; Borusyak et al., 2022), with the additional advantage that school-to-school distances are free of the measurement error inherent in zipcode-centroid-to-school distances, thereby correcting the attenuation bias that afflicts OLS.

The preferred 2SLS estimate of the distance coefficient is $\hat{\lambda} \approx -0.41$ on log distance, roughly four times the OLS estimate with market fixed effects and controls (-0.11). This gap is consistent with both measurement error attenuation and sorting bias. The implied magnitude is substantial: doubling distance reduces a school's market share by approximately 25 percent at the median. By contrast, the coefficient on school value-added is small and imprecisely estimated once distance is instrumented, indicating that distance dominates quality in parental school choices. This finding resonates with a broader literature documenting limited parental responsiveness to school effectiveness measures (Abdulkadiroğlu et al., 2020; Ainsworth et al., 2023).

Third, I use the estimated model to conduct counterfactual welfare analysis. Holding the total number of closures in each year fixed, I compare the consumer surplus implications of five closure rules: the actual PRDE decisions, closing the smallest schools by enrollment, closing the lowest-performing schools by value-added, closing schools with

the lowest utilization rates, and random closure. The actual closures perform well relative to alternatives when outcomes are weighted by the number of affected students. The lowest-performing rule generates the largest welfare losses because low-performing schools tend to be located in areas where remaining alternatives are also far away. The key mechanism is that the geography of closures—where closed schools sit relative to alternatives—matters as much as which schools are selected for closure.

The paper connects to several strands of the industrial organization and education economics literatures. On school consolidation, the most closely related studies are Berry and West (2010), Brummet (2014), and Engberg et al. (2012). Puerto Rico's consolidation is distinguished by its scale (45 percent of schools, compared to 10–33 percent in U.S. cities) and its developing-country context (57 percent child poverty). On demand estimation in education, the paper contributes to a growing literature applying methods from empirical IO to school choice (Bayer et al., 2007; Neilson, 2021; Epple et al., 2018; Agarwal and Somaini, 2020). The closure-based instrument is new to this literature and may be applicable in other settings where school entry or exit generates spatial demand shifters. On welfare evaluation of closure policy, the paper relates to Epple et al. (2018) and Abdulkadiroğlu et al. (2025). The contribution here is to combine structural demand estimation with quasi-experimental variation from a large-scale consolidation to evaluate realized closure decisions against counterfactual alternatives.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and data. Section 3 presents the reduced-form evidence. Section 4 sets up the demand model. Section 5 discusses identification. Section 6 presents estimation results. Section 7 conducts counterfactual welfare analysis. Section 8 concludes.

2 Institutional Background and Data

2.1 Puerto Rico's Education System

Puerto Rico operates a unified public school district under the Puerto Rico Department of Education (PRDE), covering all 78 municipalities on the island. The system is organized into three levels: elementary (grades K–6), junior high (grades 7–9), and high school (grades 10–12). Unlike the mainland United States, where education governance is fragmented across thousands of districts, PRDE functions as a single administrative unit responsible for all public schools. This centralized structure means that consolidation decisions are made at the territorial level rather than by individual districts, and that students are not constrained by district boundaries in their choice of school.

Puerto Rico's economic and demographic context differs markedly from the mainland. Approximately 57 percent of children live below the federal poverty line, and the territory's educational system has historically lagged behind all 50 states in standardized test performance. A protracted fiscal crisis, culminating in the Puerto Rico Oversight, Management, and Economic Stability Act of 2016, put severe pressure on public expenditures, including education. These fiscal constraints coincided with—and were partly driven by—one of the sharpest declines in child population observed anywhere in the Western Hemisphere. Between 2007 and 2017, the population aged 0–5 fell by 30 to 60 percent across most municipalities, driven by a combination of low fertility and sustained outmigration to the mainland. Figure 1 shows the geographic distribution of this decline.

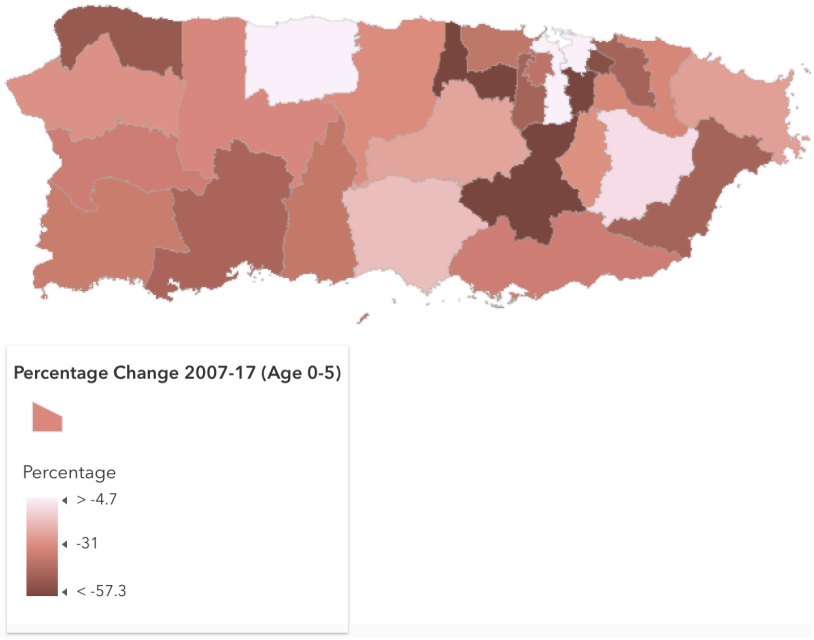


Figure 1: Percentage Change in Population Aged 0–5, 2007–2017

Notes: Darker shading indicates larger declines. Source: Puerto Rico Community Survey.

2.2 The Consolidation Program

The PRDE began closing schools on a small scale around 2010, with a major acceleration in 2014/15 that continued through 2018. By the end of this period, approximately 688 public schools had been shuttered—roughly 45 percent of the pre-consolidation total. This makes Puerto Rico's program the largest school consolidation in recent U.S. history, exceeding Chicago (200 closures, 33 percent), Washington, D.C. (38 closures, 25 percent), and Philadelphia (30 closures, 10 percent) in both absolute and proportional terms. In-

ternationally, the scale is comparable to China's consolidation of rural schools during 2000–2016 and Hong Kong's program in 2003.

The stated criteria for closure included enrollment levels, capacity utilization ratios, per-pupil operating costs, and infrastructure condition. Empirically, the strongest predictors of closure are enrollment and average class size: schools in the bottom quintile of enrollment are dramatically more likely to close than those in higher quintiles, and schools with average class sizes below 15 face elevated closure risk. Importantly, measures of academic performance—such as the share of students scoring proficient on standardized tests—are not significant predictors of closure, conditional on enrollment and class size. Figure 2 illustrates these patterns.

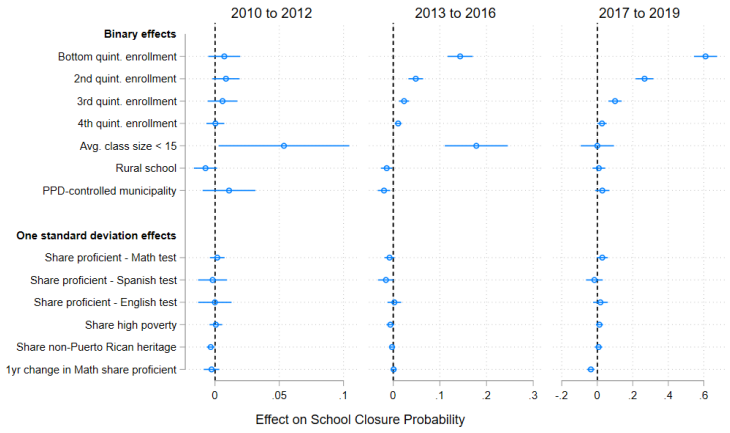


Figure 2: Predictors of School Closure

Notes: Coefficients and 95% confidence intervals from a linear probability model of school closure on school characteristics, estimated separately by closure wave. Regional fixed effects included.

Closed schools are disproportionately rural (approximately two-thirds, compared to 46 percent of non-closed schools) and small (average enrollment of 138–156, compared to 322–386 at non-closed schools). The share of students in poverty is modestly higher at closed schools (about 3 percentage points), and academic performance is somewhat higher, likely reflecting the negative correlation between enrollment and measured proficiency in small-school settings. Fewer than 10 percent of students are directly affected by closure in any given year, since closures target the smallest schools.

2.3 Data Sources

The analysis draws on three primary data sources.

PRDE Administrative Data (2010–2019). The core dataset consists of student-level administrative records from PRDE covering academic years 2010/11 through 2018/19. These records include student demographics (gender, poverty status, special education classification), grade and school enrollment, and end-of-year status (promoted, retained, transferred, or exited). Standardized test scores in math, Spanish, and English are available for students in grades 3–8 and grade 11. I also observe school-level characteristics including total enrollment, number of classrooms, staffing, and geographic coordinates. The universe of records contains approximately 2.8 million student-year observations covering roughly 670,000 unique students.

Puerto Rico Community Survey (PRCS). Block-level data from the PRCS provide demographic and housing characteristics used to construct market-level population counts. I use the school-age population in each municipality-grade-year cell to define market size, which in turn determines the outside option share (private school enrollment and non-enrollment). Zipcode-level centroids from PRCS shapefiles serve as the basis for computing student-to-school distances.

Puerto Rico Council of Education. Aggregate private school enrollment data are obtained from the Puerto Rico Council of Education. These data are used in conjunction with PRCS population counts to construct the outside option share in the demand model.

3 Motivating Evidence

Before developing the structural model, I present reduced-form evidence from a matched event study design. These results serve two purposes: they document the causal effects of displacement on student outcomes, and they establish the empirical patterns—particularly the centrality of distance—that motivate the demand model.

3.1 Event Study Design

I estimate the effect of being displaced by a school closure using a staggered event study framework. The treatment group consists of students who attended a school in its final year of operation and were not enrolled in the terminal grade (and would therefore need to relocate). The control group consists of students who were never displaced and for whom fewer than 5 percent of grade-level peers were displaced in the preceding four

years, excluding students at receiving schools. To ensure comparability, treated and control students are matched using coarsened exact matching (Iacus et al., 2012) on grade, academic year, region, urban/rural status, and a four-year enrollment history.

The estimating equation is:

$$y_{it} = \sum_{k=-4}^{-2} \tau_k D_{ik} + \sum_{k=0}^{3} \tau_k D_{ik} + \alpha_i + \lambda_{tm} + \varepsilon_{it}, \tag{1}$$

where y_{it} is the outcome for student i in year t, $D_{ik} = \mathbf{1}\{k = t - C_i\}$ indicates event time relative to the closure year C_i, α_i are student fixed effects, and λ_{tm} are match-cohort-by-year fixed effects. The period $k = -1$ is omitted as the reference. I apply the interaction-weighted estimator of Sun and Abraham (2021) to account for potential bias from heterogeneous treatment effects across cohorts displaced at different times. Standard errors are clustered at the municipal-district level.

3.2 Results: Displacement Effects

Enrollment. Displacement reduces the probability of continued enrollment in the public system by approximately 2.5 percentage points in the closure year, growing to about 4 percentage points by three years post-closure (Figure 3). Pre-treatment coefficients are close to zero and jointly insignificant, supporting the parallel trends assumption. The enrollment decline is concentrated among rural students, for whom the probability of exit is roughly twice that of urban students (Appendix Figure 12). This gap is consistent with rural students facing longer distances to alternative schools and fewer viable options within commuting range.

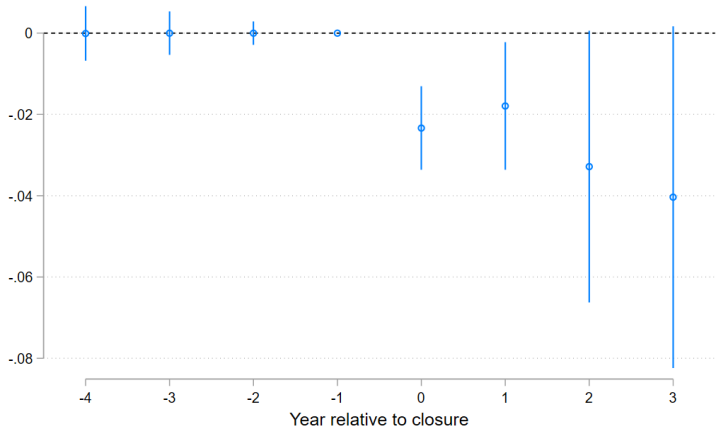


Figure 3: Effects of Displacement on Student Enrollment in the Public System

Notes: Coefficients and 95% confidence intervals from equation (1). Event study estimates adjusted using the Sun and Abraham (2021) interaction-weighted estimator. Standard errors clustered at the municipal-district level.

Test Scores. The average effect of displacement on math test scores is small and statistically insignificant across all post-treatment horizons (Figure 4), consistent with findings in Brummet (2014) and Engberg et al. (2012). However, the average conceals large and offsetting heterogeneous effects. When I stratify by the baseline relative performance of the closed school—defined as the difference in average test scores between the closed school and schools within a 5 km radius during 2010–2012—a striking pattern emerges (Figure 5). Students displaced from underperforming schools experience math score gains of 0.1–0.2 standard deviations that persist and grow over three years post-closure. Students from overperforming schools experience losses of comparable magnitude. Students from average-performing schools show no significant effects.

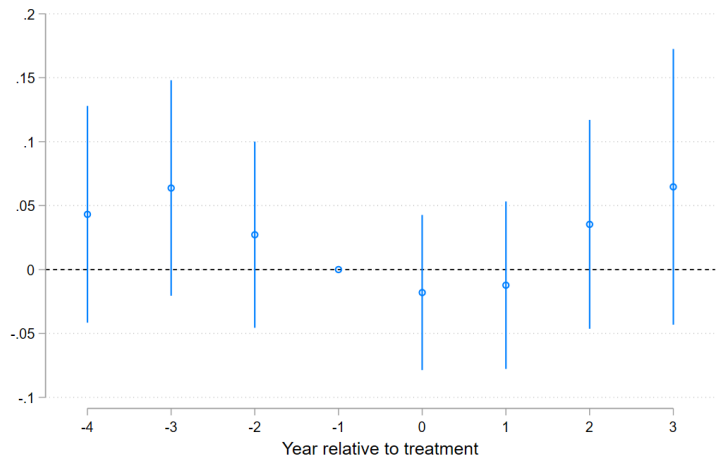


Figure 4: Average Effects of School Closures on Math Test Scores

Notes: Coefficients and 95% confidence intervals from equation (1), with math achievement measured in student-level standard deviation units. Pre-trends test p-value $= 0.71$. Standard errors clustered at the municipal-district level. Sample includes tested students in grades 3–8.

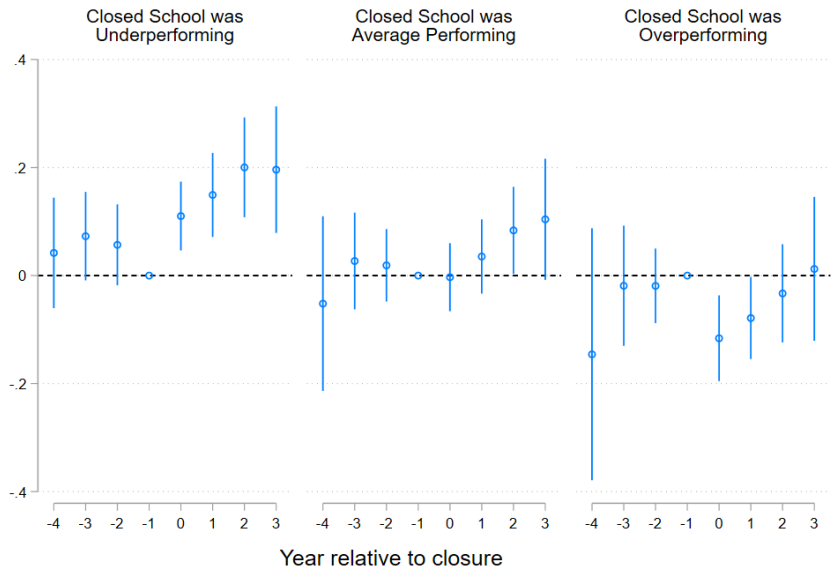


Figure 5: Effects on Math Scores, by Baseline Relative Performance of Closed School

Notes: Coefficients and 95% confidence intervals from the heterogeneous event study specification. Math achievement in student-level standard deviation units. Sample excludes baseline years 2010–2012.

Two additional dimensions of heterogeneity are informative. First, displacement distance matters: students whose median peer traveled less than 1 km to a new school show

the largest gains, while those traveling farther show weaker or null effects (Appendix Figure 14). Second, the concentration of displaced students at receiving schools is associated with worse outcomes (Appendix Figure 15), consistent with congestion or disruption effects.

3.3 Proximity Preference in Substitution Patterns

A further piece of descriptive evidence underscores the centrality of distance. Figure 6 plots the cumulative distribution of displacement distances. Approximately 85 percent of displaced students enroll at a school within 5 km of their closed school, with a median displacement distance of roughly 2 km. This sharp concentration of substitution within a narrow geographic radius is a hallmark of spatially differentiated markets and motivates the demand model's focus on distance as the primary product characteristic.

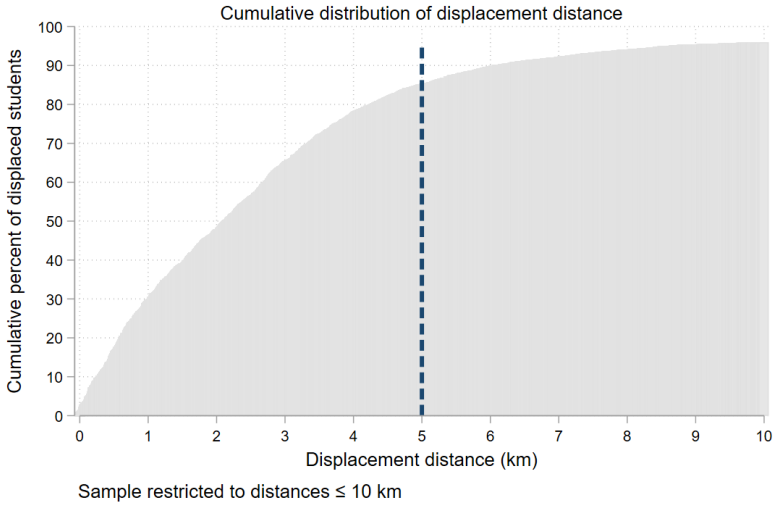


Figure 6: Cumulative Distribution of Displacement Distance

Notes: Sample restricted to displacement distances $\leq$ 10 km. Dashed vertical line marks 5 km.

3.4 Spillover Effects

I also estimate spillover effects on non-displaced students at schools that receive displaced peers, using a specification that interacts the share of displaced peers and displaced teachers with post-closure indicators. A 10 percentage point increase in the share of displaced peers is associated with a small but statistically significant increase in the probability of exiting the public system (0.2–0.6 percentage points) and modest negative effects on math test scores (approximately -0.02 standard deviations).

3.5 Implications for the Structural Model

Three takeaways from the reduced-form evidence motivate the structural approach. First, distance is central to how families respond to closures: the strong proximity preference evident in substitution patterns and the distance gradient in test score outcomes both point to distance as the dominant attribute in school choice. Second, the quality match between displaced students and destination schools drives achievement effects, implying that which schools are closed and what alternatives exist nearby are the key policy levers. Third, the heterogeneity in effects across closure contexts means that a structural model capable of tracing out the choice-set implications of alternative closure rules is needed to evaluate policy.

4 School Choice Model

4.1 School Value-Added

Before specifying the demand model, I estimate school quality using a value-added framework. For student i tested in subject s at school j in year t, the test score equation is:

$$Y_{ist} = q_{j(i,t),t} + \alpha_i + \phi_s + \epsilon_{ist}, \tag{2}$$

where q_{jt} is a school-year value-added parameter, α_i is a student fixed effect capturing permanent ability, and ϕ_s is a subject fixed effect. Identification comes from within-student variation across schools, subjects, and grades, leveraging the network of students who move between schools over the panel. I apply empirical Bayes shrinkage to obtain posterior estimates $q_{jt}^{EB} = \lambda_{jt}\hat{q}_{jt} + (1 - \lambda_{jt})\bar{q}$, which reduces noise in value-added estimates for small schools. The shrunk value-added q_{jt}^{EB} enters the demand model as an observed measure of school quality.

4.2 Demand Specification

I define a market as a municipality $\times$ grade $\times$ year cell. Within market m, the choice set $\mathcal{J}_m$ consists of all public schools operating in that municipality and serving the relevant grade level. The outside option ($j = 0$) captures private school attendance and non-enrollment; the market size N_m is the total school-age population in the municipality-grade-year cell.

Student i in market m choosing school j obtains utility:

$$U_{ijm} = \delta_{jm} + \varepsilon_{ijm}, \qquad \varepsilon_{ijm} \sim \text{Type I Extreme Value}, \tag{3}$$

where δ_{jm} is the mean utility of school j in market m and ε_{ijm} is an idiosyncratic taste shock. The mean utility is specified as:

$$\delta_{jm} = \bar{\beta} \cdot q_{jm} + \bar{\lambda} \cdot d_{jm} + X'_{jm}\theta + \xi_{jm}, \tag{4}$$

where q_{jm} is school value-added from (2), d_{jm} is log distance from the zipcode centroid to school j (averaged across students in the market), X_{jm} is a vector of school characteristics (log capacity, student-teacher ratio, average class size, urban/rural indicator), and ξ_{jm} is unobserved school quality.

Under the logit assumption, the market share of school j in market m is:

$$s_{jm} = \frac{\exp(\delta_{jm})}{1 + \sum_{k \in \mathcal{J}_m} \exp(\delta_{km})}, \tag{5}$$

and the outside option share is $s_{0m} = 1/(1 + \sum_{k \in \mathcal{J}_m} \exp(\delta_{km}))$. Taking the log-ratio yields the standard linear estimating equation:

$$\ln s_{jm} - \ln s_{0m} = \bar{\beta} \cdot q_{jm} + \bar{\lambda} \cdot d_{jm} + X'_{jm}\theta + \xi_{jm}. \tag{6}$$

The parameter of primary interest is $\bar{\lambda}$, the marginal (dis)utility of log distance. Its magnitude governs both the substitution patterns following closures and the welfare cost of removing schools from the choice set.

4.3 Discussion

The logit specification offers tractability and transparency at the cost of restrictive substitution patterns via the independence of irrelevant alternatives (IIA) property. I view this as a reasonable starting point given the policy question: the counterfactual exercise compares closure rules at the level of aggregate consumer surplus, where the logit's closed-form welfare expression is a significant advantage. An extension incorporating random coefficients on distance and demographic interactions is in progress, following the BLP framework (Berry et al., 1995), to allow for richer heterogeneity in distance sensitivity across student types.

5 Identification

5.1 Endogeneity of Distance

The key econometric challenge in estimating (6) is that d_{jm} is correlated with the unobserved demand shock ξ_{jm}. This correlation arises from residential sorting: families with strong preferences for school quality may locate near high-quality schools, inducing a negative correlation between distance and ξ_{jm}. In addition, the distance variable is measured using zipcode centroid-to-school distances averaged across enrolled students, which introduces classical measurement error that attenuates the OLS estimate of $\bar{\lambda}$ toward zero.

To fix ideas, decompose the unobservable as $\xi_{jm} = \lambda_m + \omega_{jm}$, where λ_m is a market-level component absorbed by market fixed effects and ω_{jm} is a school-specific residual. Consistent estimation of $\bar{\lambda}$ requires an instrument that shifts the distance landscape for school j without correlating with ω_{jm}.

5.2 Closure-Based Instruments

I construct an instrument that measures each school's exposure to demand shocks generated by nearby closures. The intuition is illustrated in Figure 7. When a school k closes, its students must enroll elsewhere. Because substitution is predominantly local (85 percent within 5 km), schools that are geographically close to the closed school absorb a disproportionate share of the displaced students. This generates variation in enrollment—and, through the market share equation, in the effective distance faced by each school—that is driven by the locations of closing schools rather than by the quality of receiving schools.

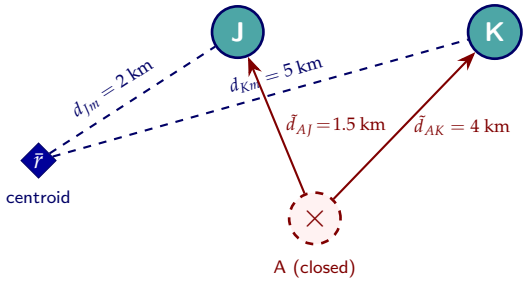


Figure 7: Instrument Visual Intuition

Notes: - - - d_{jm}: centroid-to-school distance (enters utility, endogenous). $\longrightarrow$ $\tilde{d}_{jk}$: school-to-school distance (builds Z^c, exogenous). School J, near the closure, faces high Z^c_J and absorbs displaced demand. School K, far from the closure, faces low Z^c_K. Comparing Δs_J vs. Δs_K identifies $\bar{\lambda}$.

Formally, the instrument for school j in market m is:

$$Z_{jm}^{c} = \ln\left(1 + \sum_{\substack{k \neq j \\ k \in \mathcal{C}_m}} n_{k,m_j'(t-1)} \cdot \kappa(\tilde{d}_{jk})\right), \tag{7}$$

where $\mathcal{C}_m$ is the set of schools that closed prior to period m, $n_{k,m_j'(t-1)}$ is the enrollment at closing school k in the previous period, $\tilde{d}_{jk}$ is the distance between schools j and k, and $\kappa(\cdot)$ is a distance-decay kernel. The instrument is high for schools located near large recently-closed schools and low for schools far from any closure.

5.3 Identifying Assumptions

The instrument relies on two key assumptions.

Relevance. Proximity to closed schools predicts the effective distance variable because the instrument captures the spatial structure of displaced demand. Schools near closures experience a positive demand shock that shifts their position in the enrollment-distance gradient. The first-stage relationship is strong: the Kleibergen-Paap F-statistic exceeds 20 in the preferred specification with market fixed effects and controls.

Exclusion. The instrument affects enrollment shares only through the distance channel, not through direct effects on unobserved school quality. This is supported by the empirical finding that closures are driven by enrollment and class size thresholds rather than by school quality measures. Conditional on market fixed effects and observable school characteristics, the proximity of school j to recently closed schools is plausibly unrelated to ω_{jm}.

An additional advantage of the instrument is that it is constructed from school-to-school distances ($\tilde{d}_{jk}$), which are fixed geographic features free of the aggregation error present in the centroid-to-school distance regressor. The instrument therefore isolates variation in true spatial access while purging attenuation bias from centroid-based measurement error, in the spirit of split-sample IV approaches to errors-in-variables problems.

5.4 First Stage

Table 1 reports first-stage regressions of log average distance on the closure exposure instrument. Without market fixed effects (columns 1–2), the instrument has a significant

negative coefficient, consistent with the prediction that schools near closures absorb displaced students who were, on average, farther from the centroid. Adding market fixed effects (column 3) strengthens the first stage substantially, with an F-statistic of 19.4. The preferred specification (column 4) adds control variables and yields an F-statistic of 23.0, well above conventional weak instrument thresholds.

Table 1: First Stage: Closure Exposure Instrument

	(1) Log Avg. Dist.	(2) Log Avg. Dist.	(3) Log Avg. Dist.	(4) Log Avg. Dist.
Potential Displacement	-0.145^{**}	-0.144^{**}	-0.304^{***}	-0.129^{***}
	(0.063)	(0.062)	(0.050)	(0.028)
Lag of School VA		0.061	0.104^{**}	0.036
		(0.062)	(0.043)	(0.033)
Market FE	No	No	Yes	Yes
Control Variables	No	No	No	Yes
Observations	33,488	33,488	33,387	33,368
Clusters	78	78	77	77
F-statistic	5.28	3.17	19.36	22.97

Notes: Standard errors clustered at the municipality level in parentheses. Control variables are lags of: log capacity, student-teacher ratio, average class size, share in poverty, share female, and share in special education. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6 Estimation Results

Table 2 reports estimates of the demand equation (6) under five specifications. Column (1) is OLS without fixed effects. Columns (2) and (3) add market fixed effects and then control variables. Columns (4) and (5) instrument distance using the closure exposure variable, with and without controls.

Table 2: Plain Logit Demand Estimation

	OLS			2SLS	
	(1)	(2)	(3)	(4)	(5)
Avg. Distance (log)	-0.611^{***}	-0.285^{***}	-0.110^{***}	-0.439^{***}	-0.409^{***}
	(0.196)	(0.021)	(0.017)	(0.077)	(0.121)
Lag School VA	0.512^{**}	-0.094^{***}	-0.005	-0.079^{***}	0.007
	(0.230)	(0.026)	(0.015)	(0.027)	(0.020)
Market FE	No	Yes	Yes	Yes	Yes
Control Variables	No	No	Yes	No	Yes
Observations	30,195	30,162	30,148	30,162	30,148
Clusters	75	75	75	75	75
KP F-stat				32.80	21.18

Notes: The dependent variable is $\ln s_{jm} - \ln s_{0m}$. Standard errors clustered at the municipality level in parentheses. Control variables are lags of: log capacity, student-teacher ratio, average class size, share in poverty, share female, and share in special education. KP F-stat is the Kleibergen-Paap rank Wald F-statistic. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.1 Interpreting the Distance Coefficient

The progression across specifications is informative. The OLS estimate without fixed effects (column 1) is -0.61, but this reflects a mix of true distance preferences, market-level heterogeneity, and sorting. Adding market fixed effects (column 2) reduces the estimate to -0.29, absorbing the component of distance variation correlated with market-level unobservables. Adding controls (column 3) further reduces it to -0.11, as school characteristics absorb some of the within-market variation in distance that correlates with ξ_{jm}.

The 2SLS estimates tell a different story. The IV estimate without controls (column 4) is -0.44, and the preferred specification with controls (column 5) yields $\hat{\lambda} = -0.41$. The fact that IV exceeds OLS with controls by a factor of nearly four is consistent with two reinforcing mechanisms: (i) measurement error from zipcode-centroid aggregation, which attenuates OLS toward zero, and (ii) sorting bias, which in this setting appears to work in the same direction as attenuation. The closure-based instrument, by construction, corrects both.

The preferred estimate of $\hat{\lambda} = -0.41$ implies that a doubling of distance (a one-unit

increase in log distance) reduces the relative log market share by 0.41. For a school with a 10 percent market share, this translates to an approximately 25 percent decline in enrollment. This is an economically large effect that underscores the importance of spatial access in school choice.

6.2 The Role of School Quality

Across all IV specifications, the coefficient on lagged school value-added ($\bar{\beta}$) is small and statistically insignificant once controls are included (column 5: 0.007, $p > 0.70$). This is not necessarily an indictment of school quality as irrelevant to families—the measurement error in value-added, the limited variation in quality across schools within a municipality, and the weak instrument problem for quality (closure instruments are spatial shifters, not quality shifters) all contribute to imprecision. Nevertheless, the finding that distance dominates quality in the reduced-form demand equation is consistent with the motivating evidence from the event study, where proximity to the new school is a stronger predictor of post-closure outcomes than the quality of the new school.

7 Counterfactual Welfare Analysis

7.1 Welfare Metric

Under the logit model, expected consumer surplus per student in market m is given by the familiar log-sum formula:

$$E[CS_m] = \frac{1}{|\bar{\lambda}|} \ln \left(1 + \sum_{j \in \mathcal{J}_m} \exp(\delta_{jm}) \right) + \text{const.} \tag{8}$$

The constant cancels in welfare comparisons across scenarios. The normalizing coefficient $|\bar{\lambda}|$ converts the logit inclusive value from utils to distance-equivalent units. For a closure scenario c that removes a subset of schools from $\mathcal{J}_m$, the per-student welfare change is:

$$\Delta CS_m^c = \frac{1}{|\bar{\lambda}|} \left[\ln \left(1 + \sum_{j \in \mathcal{J}_m^c} \exp(\delta_{jm}) \right) - \ln \left(1 + \sum_{j \in \mathcal{J}_m} \exp(\delta_{jm}) \right) \right], \tag{9}$$

where $\mathcal{J}_m^c \subset \mathcal{J}_m$ is the post-closure choice set. Aggregate welfare is $\Delta CS^c = \sum_m N_m \cdot \Delta CS_m^c$, weighted by market size N_m. Weighting is essential: unweighted averages can reverse sign relative to population-weighted averages when small markets with extreme

welfare changes have disproportionate influence, a form of Simpson's paradox.

7.2 Closure Scenarios

I compare five closure rules, each calibrated to close the same number of schools as the actual PRDE decisions in each year:

1. **Actual closures**: the schools that PRDE chose to close.
2. **Close smallest**: schools with the lowest enrollment.
3. **Close lowest-performing**: schools with the lowest value-added.
4. **Close lowest utilization**: schools with the lowest enrollment-to-capacity ratio.
5. **Random closures**: schools selected uniformly at random, averaged over multiple draws.

Each scenario is computed non-cumulatively: I analyze each closure wave independently using the contemporaneous choice set, so that a school closed in 2015 is absent from the 2015 counterfactual but the 2016 counterfactual uses the choice set as it actually stood in 2016. This approach avoids the path-dependence issues that would arise from sequentially removing schools.

7.3 Results

The counterfactual results are presented in Figures 8 and 9. Several patterns emerge.

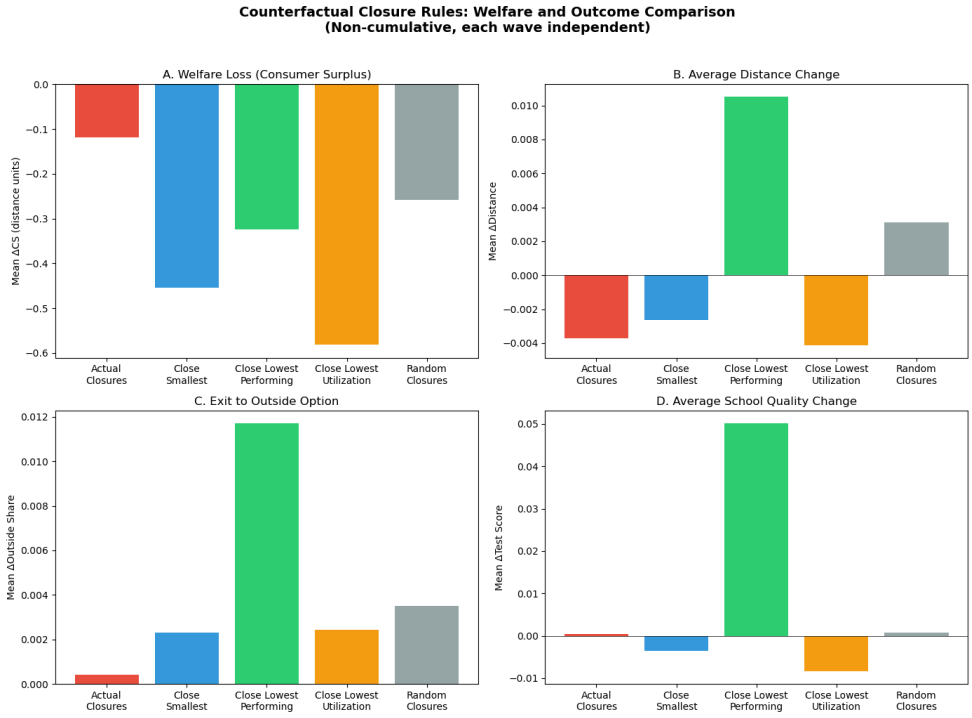


Figure 8: Counterfactual Closure Rules: Welfare and Outcome Comparison

Notes: Each panel reports mean changes across markets for each closure rule, holding total closures fixed within each wave. Panel A: consumer surplus change (distance units). Panel B: average distance change. Panel C: change in outside option share. Panel D: average school quality change. Non-cumulative, each wave analyzed independently.

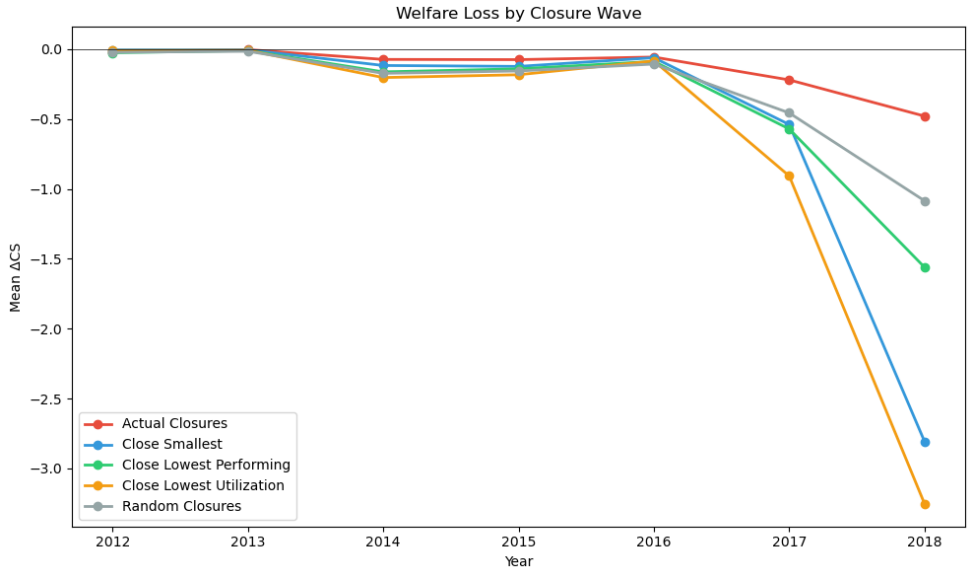


Figure 9: Welfare Loss by Closure Wave

Notes: Mean ΔCS (distance units) by closure year for each rule. The divergence across rules widens in 2017–2018, when a large wave of closures hit rural and semi-urban areas.

First, all closure rules generate welfare losses, as removing schools from the choice set can only reduce (or leave unchanged) the inclusive value. The magnitudes vary substantially, however. The actual PRDE closures generate moderate welfare losses that are smaller than most alternatives when weighted by the number of affected students.

Second, the smallest-schools rule yields welfare losses that are comparable to or slightly better than the actual closures in unweighted comparisons, reflecting the fact that small schools contribute less to the inclusive value. However, because small schools serve fewer students, the weighted welfare impact is modest, and the rule misses opportunities to close larger but underutilized schools in dense areas where alternatives abound.

Third, the lowest-performing rule generates the *largest* welfare losses among the deterministic rules. This counterintuitive result arises because low-performing schools tend to be located in peripheral areas where the remaining schools are also far away. Closing them forces students into long-distance substitution, and the distance penalty dominates any quality gain from attending a (weakly) better school. This finding reinforces the structural result that distance dominates quality in school choice and has direct policy implications: targeting low-performing schools for closure, without considering the spatial distribution of alternatives, can make families worse off.

Fourth, the lowest-utilization rule performs comparably to the smallest-schools rule. Both are proxies for excess capacity, and both generate similar patterns of spatial reallo-

cation.

Fifth, the time profile of welfare losses is revealing (Figure 9). Through 2016, all rules generate roughly similar (and modest) losses, because early closures involved small numbers of schools in relatively dense areas. The divergence across rules widens dramatically in 2017–2018, when a massive wave of closures hit rural and semi-urban areas. In this wave, the actual closures substantially outperform the alternatives, suggesting that PRDE's selection criteria—while not explicitly optimizing welfare—effectively avoided closing schools whose removal would strand students in low-access areas.

The key takeaway is that the geography of closures matters as much as which schools are selected. A closure rule that appears sensible on school-level criteria (e.g., close the lowest performers) can produce large welfare costs if it happens to remove schools in areas where alternatives are sparse. Conversely, a rule that targets small schools in dense areas minimizes welfare losses even if the closed schools were not the "worst" by any academic metric.

8 Conclusion

This paper studies the demand-side consequences of school consolidation in Puerto Rico, combining reduced-form evidence on the heterogeneous effects of closures with a structural discrete choice model that quantifies welfare costs and evaluates alternative closure rules. Three findings stand out.

First, distance is the dominant factor in school choice. The IV estimate of the distance coefficient, identified using closure-based instruments that exploit the spatial reallocation of displaced students, is roughly four times the OLS estimate and implies that doubling distance reduces a school's market share by about 25 percent. School quality, as measured by value-added, plays a secondary role that is difficult to identify cleanly with spatial instruments.

Second, the welfare effects of closures depend critically on the spatial configuration of remaining alternatives, not just on the characteristics of the schools that close. The reduced-form evidence shows this clearly: test score effects are driven by the match between displaced students and destination quality, while enrollment effects are driven by distance to the nearest alternative. The structural model formalizes this insight and shows that closure rules targeting low-performing schools can generate the largest welfare losses because they tend to remove schools from areas with poor spatial access.

Third, Puerto Rico's actual closure decisions, while not derived from an explicit welfare optimization, perform reasonably well relative to simple alternatives. The actual clo-

sures appear to have avoided closing schools in areas where the remaining alternatives are sparse, implying a de facto sensitivity to the geographic distribution of the school network. Whether this outcome reflects deliberate policy design or fortuitous correlation with enrollment-based criteria is an open question.

Several extensions are in progress. The logit model imposes restrictive substitution patterns via IIA; I am extending the analysis to a BLP random coefficients specification that allows distance sensitivity to vary with student demographics and unobserved heterogeneity, using post-closure substitution patterns as micro-moments to identify the random coefficient on distance. I am also investigating the housing market consequences of closures using block-level PRCS data, testing whether closures induce residential reallocation that compounds the demand-side welfare effects estimated here. Finally, a supply-side analysis of how closures affect school quality provision—through changes in class size, teacher assignment, and resource allocation at receiving schools—would complete the welfare picture by accounting for equilibrium responses that the current partial-equilibrium demand model holds fixed.

From a policy perspective, the results suggest that consolidation planners should think of school closure not primarily as a question of which schools are worst, but as a question of spatial network design. The welfare cost of removing a school depends less on that school's quality than on the distance its students must travel to the next-best alternative. Tools from industrial organization—demand estimation, consumer surplus analysis, and market structure evaluation—are well suited to this problem.

References

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

Abdulkadiroğlu, A., P. A. Pathak, and C. R. Walters (2025). Who gets what may not matter: Understanding school match effects. Working Paper.

Agarwal, N. and P. Somaini (2020). Revealed preference analysis of school choice models. *Annual Review of Economics*, 12(1):471–501.

Ainsworth, R., R. Dehejia, C. Pop-Eleches, and M. Urquiola (2023). Why do households leave school value added on the table? *American Economic Review*, 113(2):456–490.

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute.

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

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Berry, C. R. and M. R. West (2010). Growing pains: The school consolidation movement and student outcomes. *The Journal of Law, Economics, & Organization*, 26(1):1–29.

Black, S. E. (1999). Do better schools matter? Parental valuation of elementary education. *The Quarterly Journal of Economics*, 114(2):577–599.

Borusyak, K., P. Hull, and X. Jaravel (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213.

Brummet, Q. (2014). The effect of school closings on student achievement. *Journal of Public Economics*, 119:108–124.

Chetty, R., J. N. Friedman, and J. E. Rockoff (2014). Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates. *American Economic Review*, 104(9):2593–2632.

Edwards, D. S. (2023). Another one rides the bus: The impact of school transportation on student outcomes in Michigan. *Education Finance and Policy*, 19(1):1–31.

Engberg, J., B. Gill, G. Zamarro, and R. Zimmer (2012). Closing schools in a shrinking district: Do student outcomes depend on which schools are closed? *Journal of Urban Economics*, 71(2):189–203.

Epple, D., A. Jha, and H. Sieg (2018). The superintendent's dilemma: Managing school district capacity as parents vote with their feet. *Quantitative Economics*, 9(1):483–520.

Gibbons, S. and S. Machin (2003). Valuing English primary schools. *Journal of Urban Economics*, 53(2):197–219.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.

Iacus, S. M., G. King, and G. Porro (2012). Causal inference without balance checking: Coarsened exact matching. *Political Analysis*, 20(1):1–24.

Neilson, C. (2021). Targeted vouchers, competition among schools, and the academic achievement of poor students. *American Economic Review*, 111(10):3443–3491.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Appendix A: Summary Statistics

Table 3: School-Level Summary Statistics

	AY 12/13 to 15/16			AY 16/17 to 18/19		
	Pre-closure	No closure	Diff.	Pre-closure	No closure	Diff.
School Enrollments						
End-of-year enrollment	143.2	344.2	−201.0	155.6	321.9	−166.3
Avg. class size	18.6	22.6	−4.0	19.7	20.8	−1.1
Student Body						
Share in poverty	0.85	0.81	0.04	0.84	0.80	0.04
Share special education	0.31	0.29	0.02	0.34	0.34	0.00
Academic Performance						
Math share proficient	0.41	0.34	0.07	0.56	0.41	0.15
Locale						
Urban	0.34	0.54	−0.20	0.35	0.54	−0.19
Rural	0.66	0.46	0.20	0.65	0.46	0.19
Observations	183	2,637		425	1,988	

Notes: "Pre-closure" refers to schools observed one year before closure. Standard deviations omitted for brevity.

Table 4: Student-Level Summary Statistics

	Overall	Treated	Control	Spillovers
Student-year observations	2,821,017	252,603	1,704,051	2,403,686
Unique students	669,274	37,227	329,331	590,233
Characteristics				
Female	0.49	0.48	0.49	0.49
Special education	0.23	0.35	0.23	0.22
Attended rural school	0.50	0.57	0.50	0.48
Student age	11.69	10.51	13.11	12.07
Outcomes				
Enrolls in PRDE next year	0.75	0.83	0.79	0.74
Normalized math score	0.00	−0.00	0.01	0.01
Class size	22.77	21.19	23.20	22.94
Closure Conditions				
Distance to new school (km)		3.21		
Share at modal receiving school		0.40		

Notes: Standard deviations omitted for brevity. Treated students are those displaced by a school closure. Controls are never-displaced students with $< 5\%$ displaced peer exposure.

Appendix B: Additional Context

Table 5: School Consolidation in U.S. Jurisdictions

Jurisdiction	Population (m)	Student Enrollment	Closures	% Closed
Philadelphia	1.58	202,767	30 (2013)	10%
Washington, D.C.	0.70	49,056	38 (2008–14)	25%
Chicago	2.71	361,314	200 (2002–18)	33%
Puerto Rico	**3.20**	**300,000**	**688 (2010–18)**	**45%**

Notes: Population figures refer to 2018. Student enrollment refers to the 2018–19 academic year.

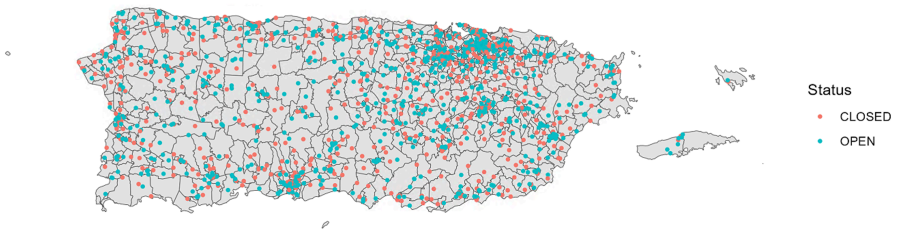


Figure 10: Open and Closed Schools in Puerto Rico

Notes: Each dot represents a school. Closed schools in red; open schools in blue.

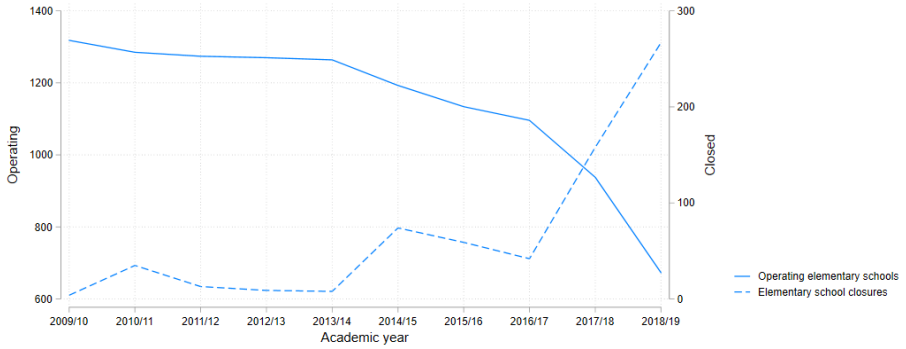
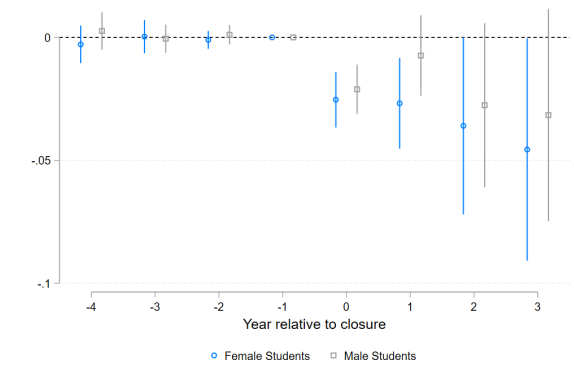


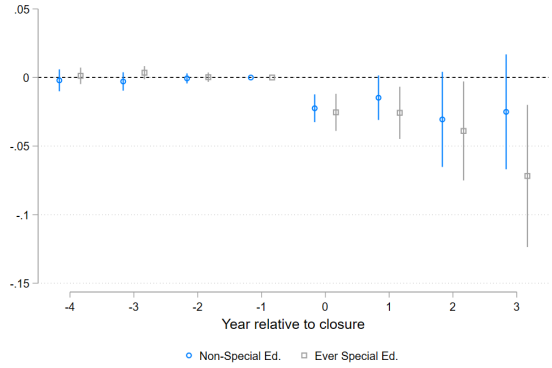
Figure 11: Operating Elementary Schools and Closures Over Time

Notes: Solid line: operating elementary schools (left axis). Dashed line: cumulative closures (right axis).

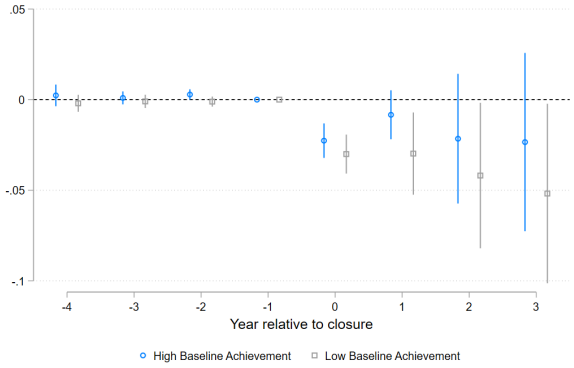
Appendix C: Additional Event Study Results



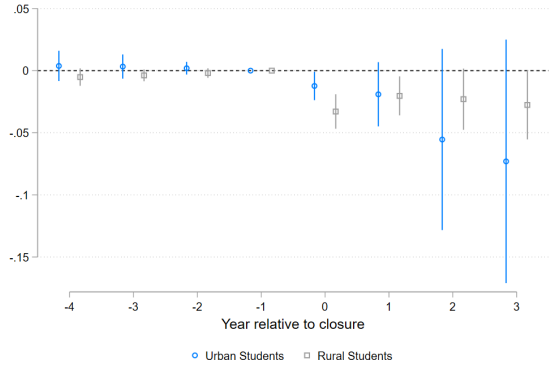
(a) Gender



(b) Special Education Status



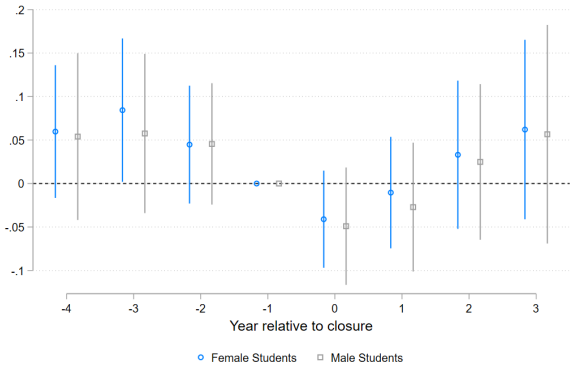
(c) Baseline Achievement



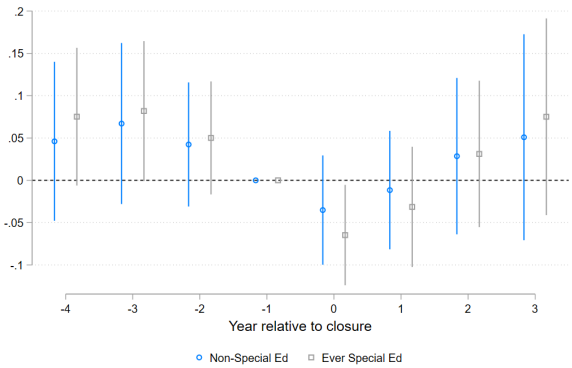
(d) Urban or Rural

Figure 12: Enrollment Heterogeneity: Effects of Displacement on System Enrollment

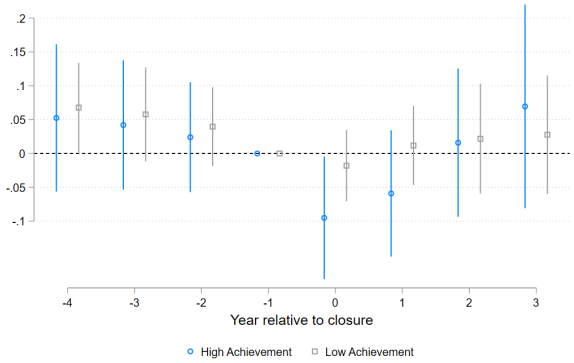
Notes: Coefficients and 95% confidence intervals from equation (1) interacted with binary student attributes. Standard errors clustered at the municipal-district level.



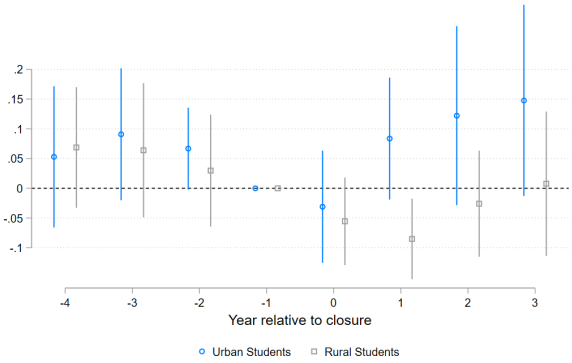
(a) Gender



(b) Special Education Status



(c) Baseline Achievement



(d) Urban or Rural

Figure 13: Test Score Heterogeneity: Effects of Displacement on Math Scores

Notes: Math achievement in student-level standard deviation units. Standard errors clustered at the municipal-district level.

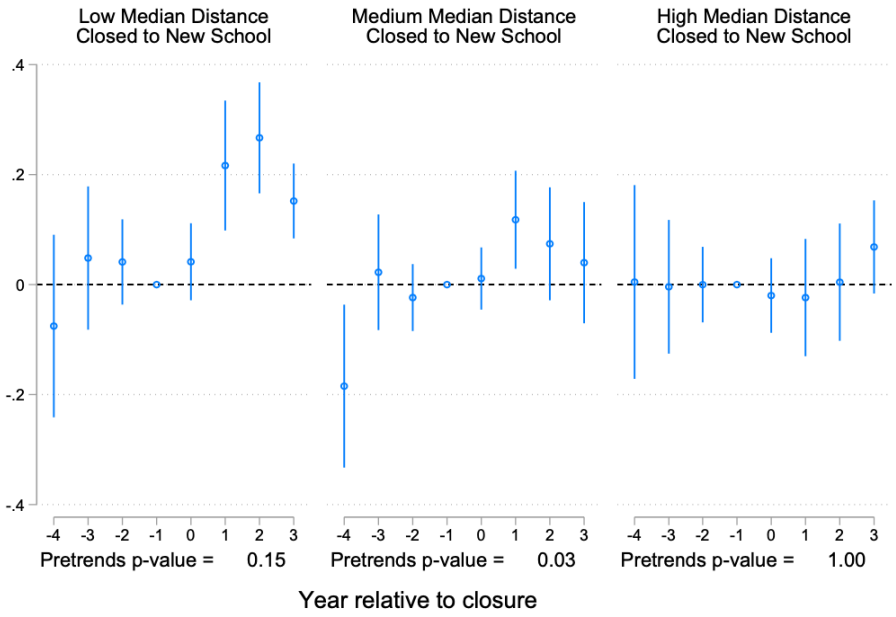


Figure 14: Effects on Math Scores, by Median Distance to New School

Notes: Low median distance: median student traveled < 1 km. High: ≥ 2.75 km. Standard errors clustered at the municipal-district level.

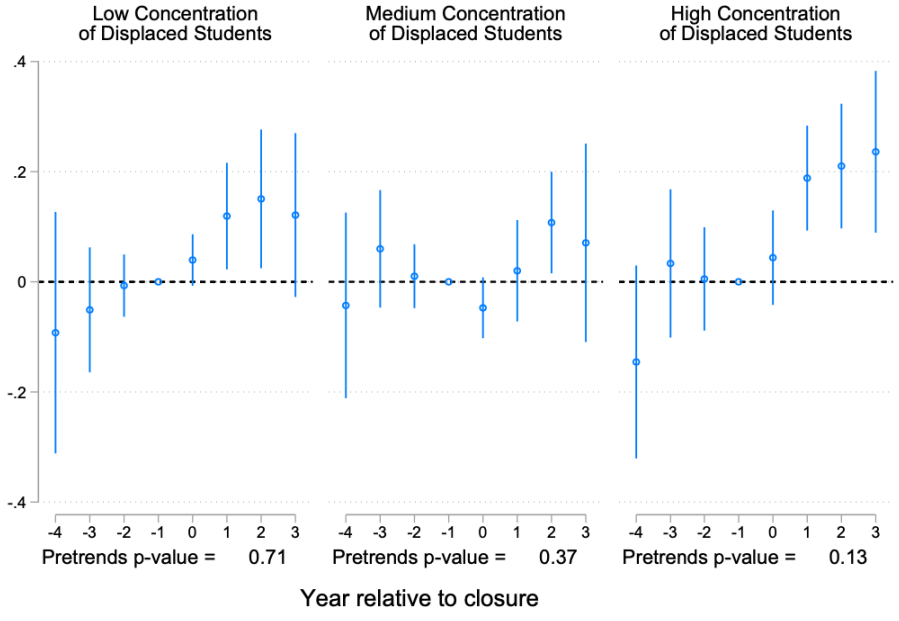
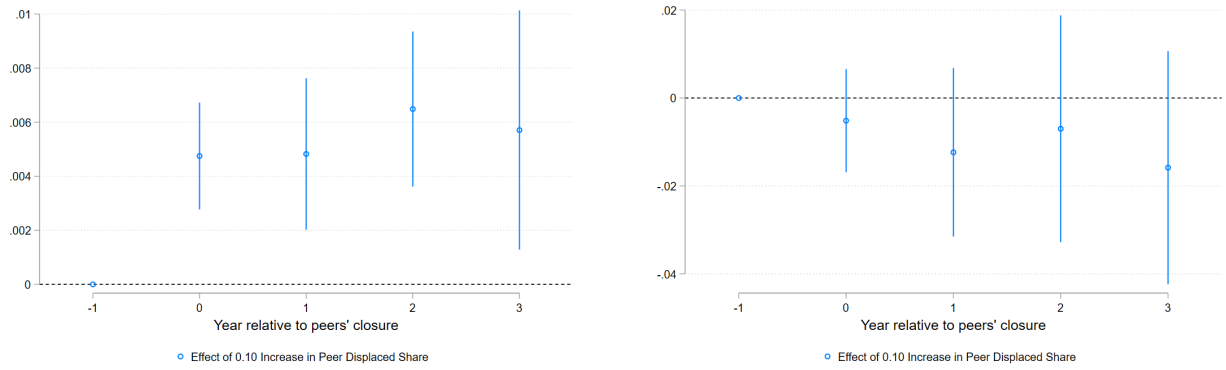


Figure 15: Effects on Math Scores, by Concentration of Displaced Students

Notes: Concentration defined as the share of displaced students from the closed school attending the modal receiving school. Standard errors clustered at the municipal-district level.



(a) Exiting Outside System (b) Exiting Within System

Figure 16: Spillover Effects on Attrition

Notes: Effect of a 0.10 increase in the share of displaced peers. Standard errors clustered at the municipal-district level.

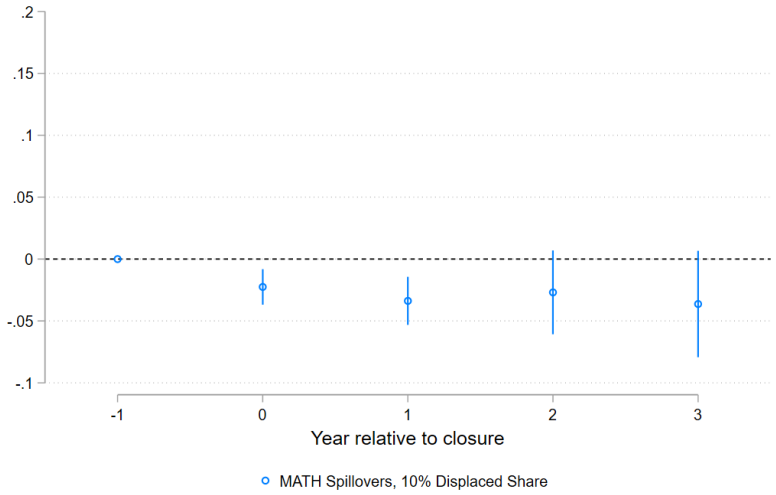


Figure 17: Spillover Effects on Math Test Scores

Notes: Effect of a 0.10 increase in the share of displaced peers. Math achievement in student-level s.d. units.

Appendix D: BLP Extension

The logit model is being extended to incorporate random coefficients and demographic interactions. The full utility specification is:

$$U_{ijm} = \delta_{jm} + \underbrace{(\Pi_\beta D_i + \sigma_\beta \nu_i^\beta)}_{\text{quality heterogeneity}} \cdot q_{jm} + \underbrace{(\Pi_\lambda D_i + \sigma_\lambda \nu_i^\lambda)}_{\text{distance heterogeneity}} \cdot d_{ijm} + \varepsilon_{ijm}, \tag{10}$$

where D_i includes gender, poverty status, and special education indicators; d_{ijm} is student-specific distance (zipcode to school); and ν_i are Halton draws for unobserved taste heterogeneity. A lognormal specification for the distance coefficient ensures that all agents have a negative distance preference. Post-closure substitution patterns provide micro-moments that identify the random coefficient on distance, since second-choice data reveal the distribution of distance preferences in a way that first-choice data alone cannot.