# A Framework for Assessing the Distributional Effects of School Closures:

## Evidence from Puerto Rico

Reza Moradi*

September 28, 2026

**Abstract**

I study the distributional consequences of school closures when families differ in their access to substitutes and teachers can be reassigned. I estimate a framework of school choice and educational production using linked student-teacher records, enrollment histories, and population data from Puerto Rico, where the number of operating public schools fell by 44 percent between 2009-10 and 2019-20. The framework separates school contributions to achievement from teacher value added and uses enrollment before and after closures to estimate substitution across schools. Rural-poor households face model-implied losses of access to public school per displaced child that are 84 percent larger than those of urban non-poor households. This disparity persists when distance costs are held common across all groups. In observed students' transfers, changes in the estimated effectiveness of attended schools, including teacher contributions, are about 0.09 test-score standard deviations less favorable for the rural-poor group than for the urban non-poor group. In counterfactual analysis, prioritizing schools with nearby alternatives reduces rural-poor households' total model-implied access loss by 25 percent relative to actual closures holding closure counts fixed. However, it removes larger estimated school contributions relative to surviving alternatives. Moreover, reassigning the observed displaced teachers among compatible local receiving posts produces little improvement in rural-poor students' teacher inputs. These findings show how differences in access to substitutes shape the incidence of consolidation, and why the loss of a school must be separated from the reallocation of its teachers.

*JEL Classification:* I21, I22, I24, I28, H75, R53

*Keywords:* school closures; school consolidation; school effects; teacher value-added; school choice; distance costs; rural schools; incidence; Puerto Rico

---

* Department of Economics, University of Toronto (email: reza.moradi@mail.utoronto.ca). I am deeply grateful to my supervisors, Gustavo Bobonis and Stephan Heblich, and to my committee member Robert McMillan. I also thank Victor Aguirregabiria, Michael Baker, Arthur Blouin, Robert Clark, Eduardo Souza-Rodrigues, and participants in the University of Toronto Applied Micro seminars. The administrative data 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, the School of Graduate Studies, and the Ontario Graduate Scholarship. I have no conflicts of interest to disclose, and all errors are my own.

# 1 Introduction

A school closure removes an option from families' choice sets, but its teachers may continue to teach at schools that remain open. The consequences depend on how readily families can substitute toward those schools and on the educational inputs available there. The contribution specific to a school cannot be assumed to move with its staff. I ask how these differences in substitution opportunities and the mobility of educational inputs shape the distributional consequences of school consolidation. Closing two schools with similar enrollment can have different consequences, even if their pupils and teachers remain in the public-school system.

I examine this question in Puerto Rico, where the number of operating public schools fell by 44 percent between 2009–10 and 2019–20 (NCES, 2021b). I develop and estimate a framework that brings spatial school choice and educational production into a common analysis of this contraction. Linked student–teacher records identify school contributions separately from teacher value added under a conditional assignment restriction. Enrollment before and after closures, school locations, and population data from the Puerto Rico Community Survey inform families' substitution across schools. Separating schools from their teachers permits a distinction between removing an option from the local network and reallocating productive inputs among the remaining schools. The main closure comparisons cover schools whose last operating academic year ended in 2013–17. I also report results through 2016, excluding the destination cohort whose next school year was disrupted by Hurricane María.

The estimates show unequal losses of access to public schooling. Rural-poor households incur 1.84 times the model-implied loss per displaced child of urban non-poor households. This measure expresses the contraction of the public-school menu as an equivalent increase in travel distance; it is neither observed additional travel nor a monetary welfare measure. Because this conversion uses each group's distance coefficient, differences in the coefficients affect the comparison. When distance costs are held common and the model is re-matched to enrollment shares and income composition, the rural-poor to urban non-poor ratio remains between 1.50 and 1.56. Thus the access disparity persists without different distance coefficients, although these scenarios do not isolate a causal effect of geography. The baseline ratio is 1.71 when the 2017 closure wave is excluded.

Changes in educational exposure are also less favorable for the rural-poor group. I measure a school's effectiveness as its estimated school contribution plus the average value added of its teachers. The destination–origin change in this measure is 0.094 test-score standard deviations lower for the rural-poor group than for the urban non-poor group; the difference is 0.138 in the earlier window. These comparisons describe

changes in exposure to estimated educational inputs, rather than treatment effects on children's scores. They use the child's administrative poverty indicator and the origin school's rural catchment share as a weight, whereas the access calculations use residential household types in the choice model. In the main window, the rural-poor group's own effectiveness change is negative but not statistically distinguishable from zero at the five-percent level; in the earlier window it is close to zero. The difference between groups should therefore be distinguished from the level change within one group.

Identification draws on movements of both pupils and teachers and on changes in the available school menu. Student and teacher movements link schools and separate their contributions under the maintained assignment restriction. Forced departure provides useful variation, but does not make destinations random. For school choice, destination ranks after closure, the income composition of public enrollment, and a matched estimate of excess exit from the origin municipality's public schools inform substitution and public-school participation. This mapping requires stable residential composition and school attributes and a link between the excess-exit estimate and the model's removal of schools from the choice set. The model does not reproduce all geographic differences in substitution and exit; the results therefore include restricted distance-cost scenarios and parameter sensitivity. The baseline choice model does not separately identify a preference for schools' achievement contributions. The production estimates instead measure the educational inputs associated with observed and simulated allocations.

I compare alternative closure rules while holding the number of closures fixed within municipality, wave, and enrollment-size groups. In the preliminary estimates from the previous run, prioritizing schools with nearby surviving alternatives reduces rural-poor households' total model-implied access loss by 24.9 percent relative to actual closures. The same rule increases the index of school contributions removed relative to surviving local schools by 0.054 test-score standard deviations per child displaced by the actual closures. This index compares school contributions across closure sets; it does not predict children's achievement changes. The comparison illustrates how a rule can improve access while removing schools with larger contributions relative to their local alternatives. These results are preliminary pending regeneration with corrected enrollment counts, which can change both the schools selected and the weights used to evaluate them.

The allocation exercises then ask how much educational exposure can change after a given set of closures. Reassigning the observed displaced teachers among compatible local receiving posts yields little improvement in rural-poor students' teacher value added. This is a result for a fixed pool of teachers and posts, not a claim about the general scope for changing teacher allocation. A separate information exercise varies an assumed response to measured school effectiveness, holding the closure list fixed. It examines how educational

exposure would change if families placed additional weight on that measure, under full take-up, unrestricted receiving capacity, and no re-sorting by incumbent families. It does not estimate the effect of an information intervention. These exercises distinguish the educational inputs associated with the schools that close from the scope for changing their allocation afterward.

The paper contributes to research on the consequences of school closures and the allocation of educational inputs. Studies of closures show that outcomes depend on the schools students leave and enter (Brummet, 2014; Steinberg and MacDonald, 2019). The unpublished Puerto Rican study by Bobonis et al. (2022) also examines heterogeneous achievement effects and the alternatives available to displaced students. My contribution is to distinguish access to the remaining school network, school contributions net of teacher composition, and the inputs that can be reallocated through teacher reassignment. This builds on Epple et al. (2018), who evaluate closure decisions allowing parents to re-sort, and on Mansfield (2015), who uses teacher mobility to separate teacher quality from other school inputs. Studies relating demand to school effectiveness (Walters, 2018; Abdulkadiroğlu et al., 2020) also distinguish the schools families prefer from those that contribute more to achievement. Maintaining that distinction here allows access and educational exposure to be evaluated together without treating the demand for a school as a measure of its educational contribution.

The analysis concerns the incidence of a contraction in the school network. Access and educational exposure are separate outcomes, with different units and populations; combining them would require an explicit social valuation. Comparing these outcomes with fiscal savings would additionally require closure-specific savings, receiving-school costs, and transportation costs. Equal numbers of closures do not supply that information, and the analysis does not identify a budget-constrained optimum.

Sections 2 and 3 describe the setting and data; Sections 4 and 5 present the model and its identification. Section 6 reports the estimates and distributional comparisons, Section 7 examines counterfactual allocations, and Section 8 concludes.

## 2 Institutional Background

Puerto Rico's consolidation program reduced a school network serving a declining student population. Closing a school changed the alternatives available to families, but it did not determine where every child enrolled or which teachers taught at the receiving schools. These subsequent choices and assignments are central to the analysis: they connect the geographic change in the network to access and to the educational inputs children received.

## 2.1 Enrollment and the school network

The Puerto Rico Department of Education (PRDE) administered a centralized public school system. It acted as both the state education agency and the sole local education agency, although its internal administration was divided into regions and school districts (PRDE, 2017a). This distinction matters because the municipalities used to define markets in this paper are geographic units, rather than separate school systems with their own budgets.

Public enrollment fell from 473,735 in fall 2010 to 365,181 in fall 2016, before Hurricane María, and to 307,282 in fall 2018 (NCES, 2021a). The number of operating public schools declined from 1,509 in 2009–10 to 847 in 2019–20, a net reduction of approximately 44 percent (NCES, 2021b). These totals describe the scale of the contraction across the public system. A net change in operating schools is not a count of permanent closures, and the totals cover more grades and years than the estimation samples defined below.

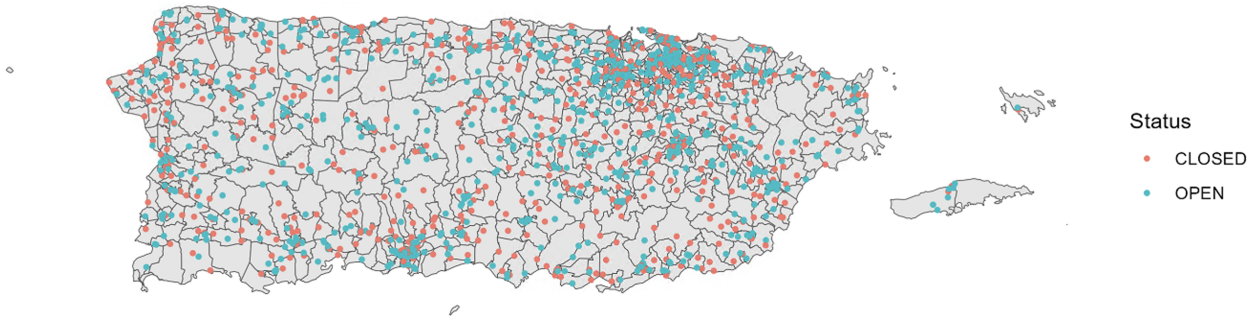
## 2.2 Closure decisions

Enrollment was prominent in the formal procedures, but it was not the sole criterion. The November 2015 circular listed health and safety, enrollment and its persistence, infrastructure, academic indicators, staffing, operating costs, location, and capacity utilization. It also allowed exceptions when transportation was difficult, a school offered a distinctive program, or special education services would be costly to provide elsewhere. Proposed receiving schools were to be assessed for space, staffing, transportation needs, and required improvements (PRDE, 2015).

The March 2017 redesign circular replaced these guidelines. It considered transport routes and costs, rural or urban location, building use, and academic offerings, with consultation preceding the committee recommendations and the Secretary's decision (PRDE, 2017c). These documents describe the intended decision process. They do not establish the weight placed on each criterion in practice, and an empirical association between school characteristics and closure cannot recover the Department's deliberations. I therefore treat the realized closure list as a policy to evaluate rather than as a randomly assigned intervention.

Low enrollment can indicate unused capacity while providing little information about the educational contribution of a school or the distance to a suitable replacement. Conversely, a school's observed test scores combine its contribution with the characteristics of its pupils and teachers. The empirical analysis separates these objects before asking what alternative closure rules would have changed. Figure 1 shows the geographic distribution of schools.

**Figure 1.** School Locations




*Notes:* Markers show school locations in the administrative map presented at EARIE 2026. Coral denotes schools recorded as closed and teal denotes schools recorded as open. The map provides geographic context; its complete set of markers is not the estimation sample for a particular closure window or outcome. Section 3 defines those samples. Source: author's EARIE presentation.

## 2.3 Enrollment and teacher reassignment

Families could request schools, subject to administrative approval. The Department's August 2017 enrollment instructions and manual allowed a family to request a school and grade; directors approved new enrollment and could record transfers to other schools. Eligibility for school transportation was determined separately (PRDE, 2017d). Observed enrollment thus reflects family choices under the admission and transportation arrangements they faced. It does not reveal an unconstrained ranking of every school on the island.

Teacher placement followed a different process. Under the January 2017 personnel circular, regions first assessed whether receiving schools could accommodate teachers from closing schools, with those not placed there available for reassignment or transfer. Within position categories, priority depended on the date of permanent status, then experience and qualifications; teachers already at receiving schools were excluded from that seniority comparison (PRDE, 2017b). These rules provide institutional context for the later reassignment exercise, without implying that the same procedure governed every earlier wave or that the records establish how it was implemented in each case.

The separation of enrollment and personnel decisions is the reason to distinguish school effects from teacher value added. A building can close while its teachers continue teaching elsewhere, and pupils and teachers need not move to the same destinations. Linked records make it possible to follow both margins instead of attributing all changes in a receiving school's effectiveness to the school itself.

## 2.4 Timing and hurricane disruption

I date a closure by the school's last operating academic year, labeled by the calendar year in which that academic year ends. A school that last operated in 2016–17 therefore has closure year 2017, with its pupils'

next enrollment observed in 2017–18. The main choice and counterfactual comparisons use closure years 2013–17; the stricter comparison ends in 2016.

Hurricane María made landfall on September 20, 2017, following Hurricane Irma earlier that month (National Hurricane Center, 2017). The recovery assessment documents damage to facilities, interrupted instruction, and reduced schedules after schools reopened (Nelson et al., 2020). Although the schools in the 2017 closure wave had completed their final operating year before María, their pupils attended receiving schools during the hurricane year. This is why I distinguish the timing of a closure from the period in which its consequences are measured. Comparisons ending with 2016 closures avoid that destination cohort, while the reduced-form exit design uses the earlier cohorts to limit this source of confounding. The 2018 closure wave is excluded from the principal choice and counterfactual analyses.

# 3 Data

I combine longitudinal administrative records from PRDE with population data from the Puerto Rico Community Survey (PRCS) and road-network distances. The records link school attendance to the teachers and schools providing instruction, while the population and geographic data describe the alternatives households could reach. The achievement, choice, and counterfactual analyses use different units and samples; Table 1 summarizes these distinctions.

## 3.1 Administrative and population records

The pupil records report school attendance, grade, demographic characteristics, family-poverty status, and standardized test scores. The achievement analysis uses mathematics, Spanish, and English in grades 3–8. The testing system changed during the panel: the *Pruebas Puertorriqueñas de Aprovechamiento Académico* (PPAA) were followed by *Medición y Evaluación para la Transformación Académica de Puerto Rico* (META-PR), first administered in 2015–16 (PRDE, 2016b,a). I normalize scores within subject, grade, and academic year to have mean zero and unit variance. This makes the estimated contributions comparable in student test-score standard deviations; it removes changes in the mean and scale of the tests, but not changes in what they measure.

The student–teacher file identifies the teacher assigned to each pupil in each subject. I link these records by pupil and year and reshape them into pupil–subject–year observations for the production model. School records supply grade offerings, enrollment, operation and closure status, coordinates, and the Department's rural designation. An enrollment total can appear on several rows for the same school and year. I retain

**Table 1.** Data and Samples

| Analysis | Unit | Coverage and restrictions |
|---|---|---|
| Achievement | Pupil–subject–year | Grades 3–8, 2011–19; 3,452,029 scored observations with teacher links, after mobility and estimation restrictions. |
| School choice | Municipality–grade–year | Grades 1–8; origin years 2013–17; 3,023 markets and 272,515 simulated household records. |
| Destination moments | Displaced-pupil record | 22,690 records in the demand sample; 16,327 destinations inside the relevant choice set. Terminal transitions outside the network are excluded. |
| Exit event study | Pupil–event-year | Closure cohorts 2013–16; origin grades 3–7; enrolled and resident in the origin municipality before displacement; 7,138 treated observations at event time zero. |
| Counterfactuals | School, pupil, or household | Closure years 2013–17, with 2013–16 reported separately; the population is specified for each outcome. |

*Notes:* An academic year is labeled by its ending calendar year. Closure year denotes the last operating year, and next-year destinations are observed in year $c+1$ for a closure dated $c$. Records, simulated households, and markets are different units and are not added. Demand counts refer to the accepted estimation sample; they are not counts of all displaced pupils in Puerto Rico. Sources: PRDE administrative records, PRCS, and the estimation samples described in the text.

**Table 2.** School Participation by Income

| | Income-poor | Non-poor |
|---|---|---|
| Outside public schools (percent) | 8.7 | 42.6 |
| Not enrolled (percent) | 1.8 | 1.8 |

*Notes:* Weighted proportions among children aged 8–14 in the 2013–17 PRCS Public Use Microdata Sample. Income poverty is defined using the Census poverty threshold and differs from the administrative poverty indicator used in the pupil records. Outside public schools includes private enrollment and non-enrollment. These are descriptive population statistics, not estimates from the school-choice model.

a single school total when describing school size; a sum across those rows would confound size with the number of grades and geographic cells represented in the file.

PRCS block-group data supply the residential school-age population and geographic weights. I supplement them with the 2013–17 Public Use Microdata Sample to measure public-school participation by income poverty. These population data include children outside the public system, which is necessary because administrative public-school enrollment alone cannot determine the size or composition of the outside option. Among children aged 8–14, 8.7 percent of those in income-poor families are outside public schools, compared with 42.6 percent of the non-poor; non-enrollment is 1.8 percent in both groups, so private enrollment accounts for the difference. Table 2 reports these observed proportions. They describe reliance on the public system, without establishing the affordability of a private school for any particular family.

### 3.2 Geography and household groups

I calculate road distances between schools and from block groups to schools using the OpenStreetMap network and the OSRM routing engine. Counterfactual comparisons also use geometric approximations where a common measure across candidate pairs is needed; Appendix D describes these constructions. School-to-school distance describes relocation from an origin to a destination, or the distance from a closing school to its nearest surviving alternative. Block-group-to-school distance enters household utility. Neither is an observed change in an individual child's commute.

The pupil records do not identify residential block groups. Instead of assigning every pupil to a ZIP centroid, the demand model integrates over possible residential block groups using the population distribution in each municipality–grade–year market. Initial attendance also carries information about residence: conditional on attendance at a particular origin, the residential weights depend on the model's probability of attending that school. Section 5.2 explains how this conditioning enters the destination moments.

I distinguish three geographic classifications because they serve different purposes. Simulated households are rural or urban according to a block-group density classification calibrated to the Department's school designation. For observed destination moments, I use the rural share of the closing school's catchment, constructed from nearby population with distance weights declining at a bandwidth of 2 kilometers. A pupil contributes fractionally to the rural and urban moments using this share and its complement. The student-level school-effect comparisons also use these fractional catchment weights, combined with each pupil's poverty indicator. A pupil can therefore contribute to both geographic groups, and the numbers of contributing records cannot be added across those groups. The exit event study instead uses the Department's designation of the origin school. These are related geographic measures, not interchangeable observations of a child's residence.

The administrative family-poverty indicator defines the low-income group in the student-level comparisons and the poverty interaction in distance costs. Where the indicator is missing, I use the pupil's modal recorded status. PRCS income poverty is a separate measure used to describe population composition and participation in the public system. I use "rural-poor" as shorthand for the relevant rural classification combined with the administrative poverty indicator, and specify whether a result concerns observed pupils or simulated households. The household groups in the current analysis are not defined by splitting schools at their median poverty rate. Appendix A gives the construction of these measures.

### 3.3 Samples and observation periods

The production model uses the scored, linked panel over 2011–19, with estimation restrictions determined by the pupil and teacher mobility needed to separate the components. The reported estimation sample contains 3,452,029 pupil–subject–year observations; this is not a count of distinct pupils. The lagged-score ANCOVA robustness specification requires a prior test and therefore begins in grade 4, whereas the main production sample includes grade 3. The two specifications answer related questions on different samples, as described in Appendix B.4.

The demand frame covers grades 1–8 and origin years 2013–17. Its principal estimation sample contains 3,023 municipality–grade–year markets and 272,515 simulated household records. The displaced-pupil file used with this sample contains 22,690 records, of which 16,327 have destinations inside the relevant choice set. Terminal-grade transitions outside the modeled network and attendance at a public school in another municipality do not supply within-market destination ranks. The model treats cross-municipality public enrollment as part of the outside option, alongside private enrollment and non-enrollment.

The principal counterfactuals consider schools last operating in 2013–17, with 2013–16 as a separate comparison. Their populations depend on the object being measured. Geographic access can be calculated for children represented in the choice frame, whereas changes in school-effect exposure require measured effects for the relevant origins and destinations. Comparing pupil-level school effects also requires observed next-year public enrollment and a usable poverty measure. I keep these restrictions explicit rather than treating the samples as a single balanced panel. Access and achievement measures retain their own populations and units throughout the paper.

### 3.4 Enrollment after closure

For the outside option, it is necessary to distinguish leaving the public system from leaving the origin municipality's public schools. I use three mutually consistent outcomes: a pupil is not observed in Puerto Rico's public schools in the following year; the pupil is observed at a public school in another municipality; or either of these occurs. The last outcome corresponds to the outside option of the municipality-level choice model. Absence from the public panel can reflect private enrollment, migration, non-enrollment, or record-linkage problems, so it cannot be interpreted as dropout alone.

The matched event-study design uses closure cohorts 2013–16 and pupils in nonterminal grades 3–7 who were enrolled and resident in the origin municipality before displacement. I estimate cohort-specific event-time effects and aggregate them using the interaction-weighted approach of Sun and Abraham (2021),

**Table 3.** Enrollment Responses to Closure

| | Effect | Standard error |
|---|---|---|
| *Panel A. Pooled effects* | | |
| Not observed in public schools | 3.20 | (0.89) |
| Public school in another municipality | 5.32 | (1.21) |
| Either outcome | 8.52 | (1.47) |
| *Panel B. Either outcome, by origin* | | |
| Rural school | 12.0 | (2.5) |
| Urban school | 3.2 | (1.2) |

*Notes:* Effects in percentage points in the first postclosure year, relative to the last pre-displacement year. The two component outcomes are mutually exclusive and add to the third. The matched interaction-weighted event study uses closure cohorts 2013–16 and origin grades 3–7, requiring enrollment and residence in the origin municipality before displacement. Rurality is the Department's origin-school designation. Standard errors in parentheses are clustered by administrative district. The estimation sample includes 7,138 treated observations at event time zero. This table reports reduced-form estimates; the pooled effect on either outcome supplies the demand model's outside-option target.

with coarsened-exact-matching weights and standard errors clustered by administrative district. Appendix A.6 gives the specification and the sample comparisons. A causal interpretation requires parallel counterfactual trends within the comparison design and no anticipatory response in the reference period; it does not follow from a claim that schools were randomly closed.

Closure increases the probability of leaving the origin municipality's public schools in the first postclosure year by 8.5 percentage points (standard error 1.5). The estimate comprises 3.2 points not observed in the public system and 5.3 points attending another municipality's public school. The corresponding effect is 12.0 points for pupils from schools designated rural and 3.2 points for those from urban schools. Table 3 reports the estimates, which enter the paper both as evidence on enrollment responses and, for the pooled effect, as the target used to discipline substitution to the outside option. The geographic differences are retained as a validation comparison rather than fitted separately.

These records connect the two parts of the model. Observed destinations describe where pupils enrolled after an origin disappeared, and the linked achievement records describe the school and teacher inputs they encountered there. The model uses that information to evaluate changes in the set of schools and in the allocation of pupils and teachers, while keeping predicted access losses distinct from changes in students' own test scores.

# 4 Model

The framework has two components. Families choose among the schools available to them, and achievement depends on the school attended and the teachers who teach the child. A closure changes the available

set, while teacher reassignment changes the inputs provided at schools that remain open. I use the choice component to measure access and the production component to measure differences in educational inputs. Keeping these outcomes separate allows the policy comparisons to show where a change improves access, where it changes educational exposure, and where the two considerations point in different directions.

The choice model recovers school attractiveness from enrollment and substitution patterns. It does not separately identify families' preferences for the estimated school and teacher contributions to achievement. Thus, the framework does not attach a utility or dollar value to a test-score standard deviation. Teacher assignments are observed inputs to production; the model does not describe teachers' choices of employer or schools' staffing decisions.

### 4.1 Educational Production

Let $i$ index students, $s$ subjects, $t$ academic years, $k$ teachers, and $j$ schools. The subjects are mathematics, English, and Spanish. Write $k(i,s,t)$ for the teacher assigned to student $i$ in subject $s$ and year $t$, and $j(i,t)$ for the school attended. The production equation is

$$Y_{ist} = a_i + \mu_{k(i,s,t),s} + \phi_{j(i,t)} + e_{ist}. \tag{1}$$

Here $Y_{ist}$ is the test score standardized within subject, grade, and year, $a_i$ is a student fixed effect, $\mu_{ks}$ is a teacher–subject effect, and $\phi_j$ is a school effect. The residual $e_{ist}$ contains the remaining achievement shocks. Teacher and school effects are measured in student test-score standard deviations.

The distinction between $\mu_{ks}$ and $\phi_j$ matters for consolidation. The teacher contribution moves with the teacher under the portability restriction in (1). The school effect is attached to the school and can reflect its organization, peer environment, and other persistent inputs left after controlling for student and teacher effects. The equation does not separate these mechanisms. Nor does it imply that closing a school lowers each child's score by that school's effect: the child attends another school, may receive different teachers, and may experience changes not represented by the fixed effects.

Student and school effects are pooled across subjects, whereas teacher effects are subject-specific. Any subject intercept can therefore be absorbed into the teacher–subject effects. Both teacher and school effects are held constant over the estimation period. School-year and classroom-year shocks enter the residual, and the estimation below allows these shocks to correlate observations within the same school and classroom.

For school $j$ in year $t$, let $n_{jt,ks}$ count scored student–subject observations taught by teacher $k$ in subject $s$,

and let $n_{jt} = \sum_{k,s} n_{jt,ks}$. The school's teacher input is

$$q_{jt} = \sum_{k,s} \frac{n_{jt,ks}}{n_{jt}} \mu_{ks}, \qquad \Psi_{jt} = \phi_j + q_{jt}. \tag{2}$$

I refer to $q_{jt}$ as teachers' average value-added and to $\Psi_{jt}$ as school effectiveness. The pooled teacher mean is $\bar{q}_j = \sum_t (n_{jt}/n_j) q_{jt}$, where $n_j = \sum_t n_{jt}$. These school-level means summarize teacher composition. An individual student's teacher input in (1) still depends on the teacher assigned to that student. The measures concern performance on the tested subjects, which need not capture schools' effects on other outcomes (Jackson, 2018).

## 4.2 School Choice

A market $m$ is a municipality of residence, grade, and year. Its inside choice set $\mathscr{J}_m$ contains public schools in that municipality serving the grade. Alternative 0 combines private schooling, nonenrollment, and enrollment at a public school in another municipality. These alternatives lie outside the modeled local public-school menu; choosing 0 need not mean leaving public education.

Household residence is represented by a census block group $b_i$. Let $R_i$ indicate a rural block group, $P_i$ administrative family poverty, and $I_i$ Census income poverty, as defined in Section 3. Road distance from the block-group centroid to school $j$ is $d_{ij}$, measured in kilometers. The utility of school $j \in \mathscr{J}_m$ is

$$\begin{aligned} u_{imj} &= \delta_{mj} + (1-\rho)\widetilde{\lambda}_i d_{ij}^{p} + \pi I_i + \varepsilon_{imj}, \quad u_{im0} = \varepsilon_{im0}, \\ \widetilde{\lambda}_i &= \widetilde{\lambda}_0 + \widetilde{\lambda}_R R_i + \widetilde{\lambda}_P P_i, \qquad\qquad p = 0.8. \end{aligned} \tag{3}$$

School–market mean utility $\delta_{mj}$ enters identically for all households in the market. It is recovered from enrollment shares and can include school effectiveness and other observed or unobserved attributes. Income poverty shifts the utility of the local public-school nest relative to the outside option. Administrative poverty and rurality shift the distance coefficient; there is no rural–poverty interaction.

The disturbances have the nested-logit distribution with one nest for local public schools and a singleton outside option, with $0 \le \rho < 1$ (Cardell, 1997; Train, 2009). The coefficients $\widetilde{\lambda}$ are expressed in within-nest utility units. Their counterparts in the utility index are $\lambda = (1-\rho)\widetilde{\lambda}$. This normalization makes distance responses comparable when the nesting parameter changes. The exponent $p$ is fixed, while $\rho$ is estimated using the exit response described below.

Suppressing the market index, define

$$
\begin{aligned}
D_i &= \sum_{j \in \mathcal{J}_m} \exp\left(\widetilde{\delta}_{mj} + \widetilde{\lambda}_i d_{ij}^p\right), \quad \widetilde{\delta}_{mj} = \frac{\delta_{mj}}{1-\rho}, \\
r_{ij} &= \frac{\exp(\widetilde{\delta}_{mj} + \widetilde{\lambda}_i d_{ij}^p)}{D_i}, \qquad T_i = \exp(\pi I_i) D_i^{1-\rho}, \\
s_{i0} &= \frac{1}{1+T_i}, \qquad s_{ij} = (1-s_{i0}) r_{ij}.
\end{aligned} \tag{4}
$$

Thus $r_{ij}$ is the probability of choosing school $j$ conditional on choosing a local public school, whereas $s_{ij}$ is its unconditional choice probability. I integrate these probabilities over block-group and poverty cells using PRCS population weights.

The distance coefficients describe effective responses to proximity under the information, transportation, and placement arrangements families face. The data do not separate travel preferences from these other influences. The baseline specification does not include separately estimated interactions between household characteristics and teacher value-added. Common school characteristics remain in $\delta_{mj}$, so this restriction does not say that families disregard educational quality.

### 4.3 Educational Exposure

Let $o(i)$ and $d(i)$ denote a displaced student's observed origin and destination. Let $a(j)$ denote school $j$'s municipality, $c_i$ the closure wave (the last operating year), and $g_i$ the grade used for the local comparison. Write $\bar{\phi}_i(C)$ for the corresponding surviving-school benchmark. The change in school-effect exposure is

$$
\Delta\hat{\phi}_i = \hat{\phi}_{d(i)} - \hat{\phi}_{o(i)} = \left(\bar{\phi}_i(C) - \hat{\phi}_{o(i)}\right) + \left(\hat{\phi}_{d(i)} - \bar{\phi}_i(C)\right). \tag{5}
$$

For municipality $a$, wave $c$, and grade $\ell$, let $\mathscr{S}_{a,c,\ell}(C_c)$ contain the surviving schools serving that grade with a measured school effect. The benchmark is

$$
\bar{\phi}_{a,c,\ell}(C_c) = \frac{\sum_{j \in \mathscr{S}_{a,c,\ell}(C_c)} E_{jc\ell} \hat{\phi}_j}{\sum_{j \in \mathscr{S}_{a,c,\ell}(C_c)} E_{jc\ell}}, \qquad \bar{\phi}_i(C) = \bar{\phi}_{a(o(i)),c_i,g_i}(C_{c_i}), \tag{6}
$$

where $E_{jc\ell}$ is preclosure enrollment at school $j$ in grade $\ell$. The first term in (5) compares the closed school with this local benchmark; the second compares the observed destination with the benchmark. Grade eligibility does not establish that a school had an available seat or was feasible for every family. Numerical comparisons using the benchmark also require the enrollment reconciliation discussed in the policy section; the direct origin–destination difference does not depend on those benchmark weights.

Replacing $\phi$ with $\Psi$ also counts teachers' average value-added. The observed comparison is $\hat{\Psi}_{d(i),c_i+1} - \hat{\Psi}_{o(i),c_i}$, so it includes the destination roster in the following school year. A comparison of schools available at the closure decision instead holds teacher composition at that decision date. I keep these comparisons distinct. Neither measures the treatment effect of closure on the student's own achievement.

### 4.4 Public-School Access

Consider removing a set $C_m \subset \mathscr{J}_m$ while holding surviving schools' mean utilities, residence, population weights, and demand parameters fixed. Write $r_{iC} = \sum_{j \in C_m} r_{ij}$. The decline in the within-public-school inclusive value is

$$L_i^{\text{pub}}(C_m) = \log D_i - \log D_i^{(-C_m)} = -\log(1 - r_{iC}), \qquad \ell_i(C_m) = \frac{-\log(1 - r_{iC})}{r_{iC}}. \tag{7}$$

The second expression reports the loss per modeled displaced child. Under the maintained within-nest choice law with fixed rankings, removing schools leaves the best available public option unchanged for families whose first choice survives. Dividing the aggregate within-nest loss by the probability of attending a closing school therefore expresses the loss on the displaced margin.

To express this index in familiar units, I find the increase $x_i(C_m)$ in distance from a reference distance $d_* = 3.9$ kilometers satisfying

$$|\widetilde{\lambda}_i| \left[ (d_* + x_i(C_m))^p - d_*^p \right] = \ell_i(C_m). \tag{8}$$

This is a kilometer equivalent, not a predicted increase in a child's journey. The nonlinear distance function requires this conversion; dividing utility loss by a constant distance coefficient would produce a different measure. Markets with no surviving inside alternative do not have a finite within-public-school index and are excluded from comparisons requiring it. The index therefore does not describe access losses for those markets.

For population aggregation, let $w_i$ be the integration weight, summing to one within market, and $M_m$ the market population. The modeled displaced mass represented by agent $i$ is

$$h_i = M_m w_i (1 - s_{i0}) r_{iC}. \tag{9}$$

For group $g$, average loss is $\sum_{i \in g} h_i x_i / \sum_{i \in g} h_i$. The group's share of displaced mass and share of total kilometer-equivalent loss use $h_i$ and $h_i x_i$, respectively. These are model-implied population weights, which need not reproduce the distribution of observed displaced records.

This measure deliberately concerns access within the local public-school menu. The full nested-model logsum would be $W_i = \log(1 + T_i)$ and would also value adjustment to the outside option. It is not the index in (7). Moreover, interpreting choice-based utility as fully informed welfare requires stronger assumptions about information (Kapor et al., 2020). I therefore report public-school access, school-effect exposure, and fiscal accounts in their own units, without adding them into a single welfare total.

# 5 Identification and Estimation

The production and choice components use different variation from the same consolidation episode. Changes in students' and teachers' assignments link schools in the production equation. Enrollment and displaced students' destinations reveal substitution in the choice equation. An independently estimated change in enrollment outside the origin municipality's public schools supplies the extensive-margin target. This section states what is needed to interpret that variation and explains where the data limit the resulting policy analysis.

## 5.1 Educational Production

I estimate (1) by least squares after absorbing student effects. A student observed with different teachers contributes comparisons of teacher inputs holding permanent student heterogeneity fixed. Students who change schools and teachers who teach the same subject at different schools link the school and teacher–subject effects across institutions. This uses the mobility-based fixed-effect logic of Abowd et al. (1999); Mansfield (2015) provides a related educational application separating teacher and school-course components. Neither mobility nor a large sample alone establishes identification: the normalized joint design must have sufficient rank, and movements must satisfy the assignment restriction below.

**Assumption 1** (Assignment and achievement shocks)**.** *Let $\mathscr{A}$ collect the student, teacher–subject, and school assignment histories and the observation pattern in the estimation sample. Conditional on the fixed effects, the remaining achievement disturbance has mean zero:*

$$\mathbb{E}[e_{ist} \mid \mathscr{A}, a, \mu, \phi] = 0 \quad \textit{for each observed } (i,s,t). \tag{10}$$

The restriction permits sorting on the permanent differences represented by the fixed effects. It excludes assignment responding to residual achievement shocks, including persistent responses to earlier shocks not absorbed by $a_i$. Such responses are a central concern in value-added estimation (Rothstein, 2010). A closure

removes the option to remain at the origin, but it does not randomize which school or teacher the child receives. Ordinary transfers also provide links and are subject to the same restriction.

The additive specification further requires teacher contributions to be portable and school effects to be sufficiently stable over the estimation period. I examine residual teacher–school match variation, school-effect stability, and the ability of estimated effectiveness changes to forecast achievement changes. These exercises address particular implications of the specification. A successful forecast of the sum $\phi + q$ would not, by itself, validate its division into school and teacher components.

Estimation error inflates the dispersion of fixed-effect estimates. I therefore adapt the leave-out variance-component estimator of Kline et al. (2020), allowing for classroom-year and school-year shocks. The common-shock variances are calibrated inputs, and I report sensitivity to them. I also compare effects estimated on disjoint school-year halves. Splitting at the school-year level prevents the same school-year shock from entering both halves; splitting classrooms alone would not. Cross-half covariance removes measurement-error variance only when estimation errors are uncorrelated across halves. Appendix B gives the calculation and its limitations.

As a robustness check, I replace student fixed effects with grade-interacted cubic controls for lagged mathematics, English, and Spanish scores and demographic controls, retaining teacher–subject and school effects. This ANCOVA specification conditions on prior achievement and therefore imposes a different assignment restriction (Chetty et al., 2014). It uses grades 4–8 with the required lagged scores. Its estimates need not equal the student-fixed-effect estimates, and agreement would not establish either identifying assumption. Appendix B.4 reports the available comparison.

### 5.2 School Choice

Aggregate enrollment shares alone cannot distinguish distance responses from school attractiveness. For any candidate parameter vector, I recover mean utilities that match those shares. The additional information comes from the schools that displaced families attend after their origin closes and from the income composition of public enrollment. This uses the information in substitution choices emphasized by Berry et al. (2004), with an additional stability restriction because the destinations here are observed after an actual closure rather than elicited in a hypothetical choice exercise.

Let $z_i = (b_i, P_i, I_i)$ collect residence and the two poverty indicators, and let $F_m$ be their distribution in the market population. For each candidate $\theta = (\vartheta', \rho)'$, where $\vartheta = (\widetilde{\lambda}_0, \widetilde{\lambda}_R, \widetilde{\lambda}_P, \pi)'$, share inversion solves

$$S_{mj}^{\text{obs}} = \int s_{mj}(z; \mathscr{J}_m, \theta, \delta_m(\theta))\, dF_m(z) \quad \text{for every } j \in \mathscr{J}_m. \tag{11}$$

Here $S_{mj}^{\text{obs}}$ is the enrollment share and $\delta_m(\theta)$ is the vector of school–market mean utilities (Berry, 1994). Consequently, unobserved attractiveness common to households may be correlated with educational effectiveness. It cannot absorb an unrestricted pattern of school attributes valued differently by household type.

Initial attendance is itself a choice. A household with characteristics $z$ attends origin $o$ with probability $s_{mo}(z)$, so its distribution among that origin's students is proportional to $s_{mo}(z)\,dF_m(z)$. The integration must account for this selection; using residential population weights alone would assign the same residential composition to every origin.

**Assumption 2** (Closure-induced substitution)**.** *Conditional on initial attendance, the attributes entering the market model, and the specified distribution of household characteristics, selection into the closure episode does not introduce additional selection on latent utility rankings. For the menu-removal experiment, survivors' utilities and household residence remain fixed, and the destination of a household initially attending a closing school is*

$$J_i(C_m) = \underset{j \in (\mathscr{J}_m \setminus C_m) \cup \{0\}}{\arg\max} \; u_{imj}, \tag{12}$$

*using the same utility realization that represents initial attendance.*

The assumption allows the closure list to depend on enrollment and school attributes. It requires those variables and the choice model to account for relevant selection in the displaced population. It also rules out closure-induced changes in information, transport, placement, capacity, or school attributes that alter rankings beyond the modeled removal. These are substantive restrictions. A school's closure is not a randomized assignment of families to their remaining alternatives.

For any survivor $j$, this assumption gives the group-removal identity

$$\Pr(O_i \in C_m, J_i(C_m) = j \mid z) = s_{mj}(z; \mathscr{J}_m \setminus C_m) - s_{mj}(z; \mathscr{J}_m), \tag{13}$$

including $j = 0$. Removing schools can change the selected alternative only for households whose initial choice belongs to the removed set. The probabilities hold mean utilities at the values recovered from the origin-year shares. They are not obtained by fitting the postclosure shares with new mean utilities.

The rank moments compare observed and predicted shares attending the second, third, fourth–fifth, and sixth-or-more distant surviving school, for four poverty-by-rural-exposure cells. The nearest-school category is omitted because the categories add to one. Observed distance ranks are measured from the origin school. Model utilities use distances from each possible residential block group; the model then predicts the same

origin-based destination ranks. The comparison therefore does not assume that the origin school is each child's home.

Administrative poverty is observed for students, but residential rurality is not. On the observed side, an origin contributes to the rural cell with its catchment rural-exposure weight $A^R(o)$ and to the urban cell with $1 - A^R(o)$. In the choice model, rurality is the block group's indicator $R_i$. The rural coefficient thus relies on the residential distribution and differences across catchments as well as substitution within markets. Catchment exposure is constructed using a two-kilometer kernel bandwidth. It is not a directly observed rural indicator for each displaced child.

For the rank moments, the implemented estimator deletes one origin at a time and excludes co-closing schools from eligible destinations. For the exit target and policy calculations, it deletes all schools closing in the market together. These operations need not give the same integrated destination distribution. Appendix C.1 states the implemented rank moments explicitly; their application to markets with simultaneous closures maintains that the single-origin approximation adequately represents the relevant conditional rank shares.

### 5.3 The Exit Target and Estimation

The extensive-margin outcome is departure from the origin municipality's public-school system, combining departure from public schooling and transfer to a public school in another municipality. I estimate its closure response using a matched event-study design with the interaction-weighted estimator of Sun and Abraham (2021). Matching uses pre-event characteristics; the regressions include student-by-match and match-year-by-calendar-year effects, with district-clustered inference. The outcome and sample definitions are given in Section 3.

The causal interpretation requires parallel untreated outcome trends after matching, no anticipation relative to the reference period, and no differential changes in observation unrelated to the modeled closure response. Pre-event coefficients provide evidence about trends, but cannot establish these restrictions. The target is the first-year effect for resident students in nonterminal tested grades in closure cohorts 2013–2016. Transferring it to the choice model also requires comparability with the model's broader grade coverage.

Let $A(\vartheta, \rho)$ denote the model's pooled probability of outside diversion conditional on initially attending a closing school, weighted to the exit-target population. I use sequential moment estimation (Newey, 1984):

$$\begin{aligned} \hat{\vartheta}(\rho) &= \arg\min_{\vartheta} g(\vartheta, \rho)' W g(\vartheta, \rho), \\ A(\hat{\vartheta}(\hat{\rho}), \hat{\rho}) &= \hat{\tau}, \end{aligned} \tag{14}$$

where $g$ contains the sixteen rank discrepancies and the discrepancy in the income-poor share of public enrollment, and $\hat{\tau}$ is the matched exit effect. The composition target informs $\pi$, which shifts all inside options relative to the outside. The rank moments jointly inform the distance coefficients after share inversion. They do not assign one independent moment to each coefficient.

The mapping in the second line is an identifying restriction: the event-study excess exit is represented by the model's diversion from removing the origin schools, with other utilities fixed. The event study adjusts for exit that would have occurred without closure, whereas the structural experiment keeps the original choice environment fixed. Equating them requires these two descriptions of the counterfactual to be sufficiently comparable. The single exit target exactly identifies the nesting parameter conditional on the other restrictions; matching it is not an independent test of the model's exit predictions.

The distance exponent is fixed at 0.8 before this estimation, with 0.74 and 1 examined as sensitivity specifications. The weighting matrix uses municipality-bootstrap variances of the rank targets and places additional weight on the precisely measured composition target. Standard errors stack the first-stage estimating equations with the exit restriction and use their full derivative, retaining the role of nonzero fitted moment discrepancies (Hall and Inoue, 2003). The reported calculation assumes zero covariance between the rank targets and the exit estimate and conditions on population weights, distance curvature, and catchment construction. Appendix C provides the weighting, covariance, and numerical details.

A full-rank local derivative of these stacked restrictions is needed to distinguish nearby parameter vectors. Profile calculations address weakly determined rural-distance responses and the nesting parameter. They do not prove global uniqueness of the economically relevant parameter vector or validity of the moment restrictions. Numerical agreement between implementations likewise verifies the calculation, not the identifying assumptions.

### 5.4 Fit and Scope

The current estimates reproduce the pooled exit target by construction, but the specification does not fit all the substitution evidence. The simulation-based overidentification test rejects at the five-percent level, with a p-value of 0.038. In particular, the model overpredicts attendance at the nearest surviving school, with the largest discrepancies in rural catchments. It also predicts a much smaller rural–urban exit difference than the event study. I therefore use the model for conditional access comparisons and report geography and common-distance-cost comparisons alongside it; I do not interpret its geographic exit predictions as additional findings.

Revealed choices do not provide a reliable separate price for teacher value-added in these data. Mean

school utility absorbs common responses to quality, and the noise in school-level teacher averages prevents an informative upper bound under the examined sign restriction. The baseline therefore estimates no household-specific quality coefficients. This limits policy exercises that change the information families have about effectiveness: those exercises require explicit response scenarios, rather than an estimated preference for quality. Similarly, applying the distance coefficients to a new closure list maintains the effective transportation and information environment represented by the historical estimates.

# 6 Results

## 6.1 School effects and the demand for access

The estimates distinguish variation in the contribution associated with a school from variation in its teachers' value-added. The covariance between estimates from separate school-year samples implies a standard deviation of school effects of about 0.30 test-score standard deviations. The corresponding dispersion in schools' average teacher value-added is between 0.06 and 0.08. Thus differences between schools cannot be described by teacher composition alone. These moments have different weighting and sampling properties from the variance of individual teacher effects; they are not a variance decomposition of students' scores. Appendix B reports the split-sample estimates, the leave-out correction under alternative common-shock calibrations, and the ANCOVA comparison.

Table 4 reports the school-choice estimates. Distance reduces the probability of choosing a school for every household group. The positive rural interaction implies a smaller effective distance cost for rural households: at the estimates, the rural-to-urban ratio is about 0.56 among non-poor households and 0.59 among poor households. The poverty interaction raises the magnitude of the distance coefficient, while the positive income-poverty coefficient increases the probability of choosing the public-school nest. These coefficients describe choices under the prevailing transport, information, and assignment environment. They do not separately identify the disutility of travel or the contribution of school transportation.

The nesting parameter is estimated at 0.847, with an inversion confidence set of [0.783, 0.904]. Its identification comes from the pooled closure-induced exit moment, whereas the destination-rank moments principally determine substitution within public schooling. The fit is uneven across these margins. The model overpredicts attendance at the nearest surviving school, particularly for rural origins, and the joint specification test rejects at the five-percent level. It also does not reproduce the rural–urban difference in exit rates. These failures limit the interpretation of finely disaggregated destination and exit predictions. The access comparisons below therefore include common-distance-cost specifications and parameter sensitivity,

**Table 4.** School Choice

| | Estimate | Standard error |
|---|---|---|
| Distance, $\widetilde{\lambda}_0$ | −2.785 | 0.577 |
| Rural distance interaction, $\widetilde{\lambda}_R$ | 1.219 | 0.637 |
| Poverty distance interaction, $\widetilde{\lambda}_P$ | −0.202 | 0.094 |
| Income-poverty inside shifter, $\pi$ | 1.411 | 0.062 |
| Nesting parameter, $\rho$ | 0.847 | 0.031 |
| *Effective rural/urban cost ratio* | | |
| Non-poor | 0.562 | 0.152 |
| Poor | 0.592 | 0.144 |

*Notes:* Distance is road kilometers raised to the power 0.8. The $\widetilde{\lambda}$ coefficients are in within-nest utility units; their utility-index counterparts equal $(1-\rho)\widetilde{\lambda}$. The base distance coefficient applies to urban non-poor households. Distance poverty is the administrative indicator; the inside shifter uses Census income poverty. Rank and composition moments use closure cohorts 2013–2017; the exit target uses cohorts 2013–2016. Standard errors come from the stacked estimating equations with the full derivative, conditional on the maintained population and geographic inputs. The reported calculation assumes zero covariance between the rank and exit targets. The inversion confidence set for $\rho$ is $[0.7833, 0.9044]$; profile inference is used for the rural shifter. The cost ratios compare rural and urban coefficients within the same poverty group at a common distance. No quality-preference coefficient is estimated in this specification. Appendix C reports the specification-fit limitations.

rather than relying on one coefficient estimate. They do not turn the model into a validated forecast of every postclosure enrollment flow.

## 6.2 The distribution of access losses

The contraction of the public-school network produces unequal access losses across household groups. For the 2013–2017 closure waves, the average loss per modeled displaced rural-poor child, expressed in kilometer equivalents, is 1.84 times that of an urban non-poor child (Table 5). The ratio ranges from 1.51 to 2.16 across the evaluated parameter values used for sensitivity analysis. Imposing a common distance coefficient and re-matching enrollment shares and public-school income composition gives ratios between 1.50 and 1.56. A substantial disparity thus remains in specifications with a common distance coefficient. This exercise does not by itself assign a causal share of the gap to geography, since mean utilities and the public-participation coefficient are adjusted as well.

Rural-poor households account for 50.0 percent of modeled displacement and 61.8 percent of the model's total access loss. Both shares use the population weights and initial attendance probabilities in the choice model. They are not shares of the observed relocation sample. Each household's loss is expressed as the additional distance that would produce the same reduction in its within-public-school inclusive value,

**Table 5.** Access Losses

| | 2013–2017 | 2013–2016 |
|---|---|---|
| *Rural-poor relative to urban non-poor, per displaced child* | | |
| Estimated demand model | 1.84 | 1.71 |
| Evaluated parameter range | [1.51, 2.16] | [1.45, 1.94] |
| Common distance cost, first reference | 1.50 | 1.45 |
| Common distance cost, second reference | 1.56 | 1.49 |
| *Rural-poor shares in the estimated demand model* | | |
| Modeled displacement | 50.0% | 52.2% |
| Total access loss | 61.8% | 62.0% |

Notes: Access losses are changes in the within-public-school inclusive value expressed in kilometer equivalents at a reference distance of 3.9 km. They are aggregated with population weights and predicted initial attendance at closing schools. Rurality and poverty are household types in the demand model. The ranges summarize evaluated demand-parameter values; they are not confidence intervals for the ratios and do not exhaust the confidence set. The sensitivity grid also retains a nesting-parameter value of 0.78, slightly below the reported lower endpoint of the inversion confidence set. The common-cost calculations first estimate the model with the rural interaction set to zero. They then assign every group either that estimate's non-poor or poor distance coefficient, re-match income composition through the participation coefficient, and re-invert mean utilities to match shares. They are restricted response scenarios rather than two unrestricted re-estimations of the common-coefficient model. Closure waves refer to the last operating year.

evaluated at the common reference distance defined in Section 4. These kilometer equivalents summarize access to the public-school menu. They are neither observed additional travel nor monetary compensating variation, and they do not incorporate the full welfare value of adjustment to the outside option. For a given loss in utility, a smaller absolute distance coefficient implies a larger kilometer equivalent. The baseline difference in this measure therefore reflects both the utility losses from the removed alternatives and the coefficients used to express those losses in distance units.

The comparison is similar when the 2017 closure wave is excluded. The rural-poor-to-urban-non-poor ratio is 1.71 at the estimates and lies between 1.45 and 1.49 in the common-cost specifications. The corresponding rural-poor shares are 52.2 percent of modeled displacement and 62.0 percent of access loss. This restriction removes the closure wave leading into the school year in which Hurricane María struck; it is reported separately from the main 2013–2017 calculation.

### 6.3 Changes in educational exposure

The school a child attends after closure need not have the same contribution to achievement as the school that closed. Table 6 compares the estimated effects of observed destinations and origins for displaced students with the required linked records. Effectiveness combines the school's contribution with the average value

**Table 6.** Destinations and Origins

| | 2013–2017 | | 2013–2016 | |
|---|---|---|---|---|
| | School effect | Effectiveness | School effect | Effectiveness |
| Rural-poor | −0.079 | −0.060 | −0.044 | 0.007 |
| | (0.036) | (0.032) | (0.047) | (0.046) |
| Rural non-poor | −0.032 | −0.005 | −0.003 | 0.062 |
| | (0.035) | (0.029) | (0.045) | (0.044) |
| Urban poor | 0.011 | 0.002 | 0.056 | 0.087 |
| | (0.027) | (0.030) | (0.045) | (0.043) |
| Urban non-poor | 0.018 | 0.034 | 0.096 | 0.145 |
| | (0.025) | (0.031) | (0.034) | (0.045) |
| Rural-poor minus urban non-poor | −0.097 | −0.094 | −0.140 | −0.138 |
| | (0.037) | (0.036) | (0.053) | (0.057) |

Notes: Destination minus origin in student test-score standard deviations. School effects are $\widehat{\phi}_j$; effectiveness adds average teacher value-added, $\widehat{\Psi}_{jt} = \widehat{\phi}_j + \widehat{q}_{jt}$. The latter compares the destination roster in the year after closure with the origin roster in the last operating year. Standard errors clustered by municipality are in parentheses. The observation is a displaced student record with observed origin, destination, poverty, rural exposure, and the required estimated effects. The rural comparisons weight each record by its origin's distance-weighted rural population share $A^R(o)$; urban comparisons use $1 - A^R(o)$. Poverty is the child's administrative indicator. The school-effect rows have 16,605, 3,377, 16,503, and 3,364 records, respectively, in 2013–2017, and 8,475, 1,574, 8,373, and 1,561 in 2013–2016 with positive weight. Because rural and urban weights can both be positive, these counts overlap and must not be added to obtain a unique sample size. These samples are distinct from the records used in the demand moments and the information exercise. The estimates describe exposure changes, not closure effects on students' scores.

added of its teachers. Each record receives its origin's rural-exposure weight in the rural comparison and the complementary weight in the urban comparison; poverty is the child's own administrative indicator. Thus these are comparisons weighted by rural exposure, rather than a partition of children by observed rural residence. They differ from the residential types used to aggregate the demand model.

Changes in estimated effectiveness are less favorable for the rural-poor group than for the urban non-poor group. The difference between their destination–origin changes is −0.094 test-score standard deviations in 2013–2017, with a municipality-clustered standard error of 0.036. It is −0.138 (0.057) for waves through 2016. This relative disadvantage persists across the two windows even though the rural-poor group's own change differs: its average change in effectiveness is −0.060 (0.032) in the main window and 0.007 (0.046) in the earlier window. Neither within-group change is statistically distinguishable from zero at the five-percent level.

Separating the school and teacher contributions helps explain these changes. In 2013–2017, the rural-poor group's destination–origin change in school effects is −0.079 (0.036), compared with 0.018 (0.025) for the urban non-poor group. Their difference is −0.097 (0.037). Adding teacher contributions makes the rural-poor group's own decline smaller, but leaves a similar difference between groups. For waves through 2016, the rural-poor change in school effects is −0.044 (0.047) and the between-group difference is −0.140 (0.053).

Teachers can move when a school closes, so the contribution associated with the closed school alone does not measure the full change in educational inputs.

These are changes in exposure to estimated school and teacher contributions. They are not treatment effects of closure on the displaced students' own scores. Closure can also change peers, class size, disruption, and the match between a student and a school. Nor does the comparison identify whether a particular destination was feasible for every family. Keeping these distinctions explicit allows the access and production estimates to inform the same question about distributional incidence without adding quantities measured in different units or drawn from different populations.

Figure 2 summarizes the two comparisons. Its first panel shows that the rural-poor access disadvantage persists in the common-cost scenarios. Its second panel distinguishes changes within each group from the difference between the rural-poor and urban non-poor groups.

# 7 Counterfactual Policies

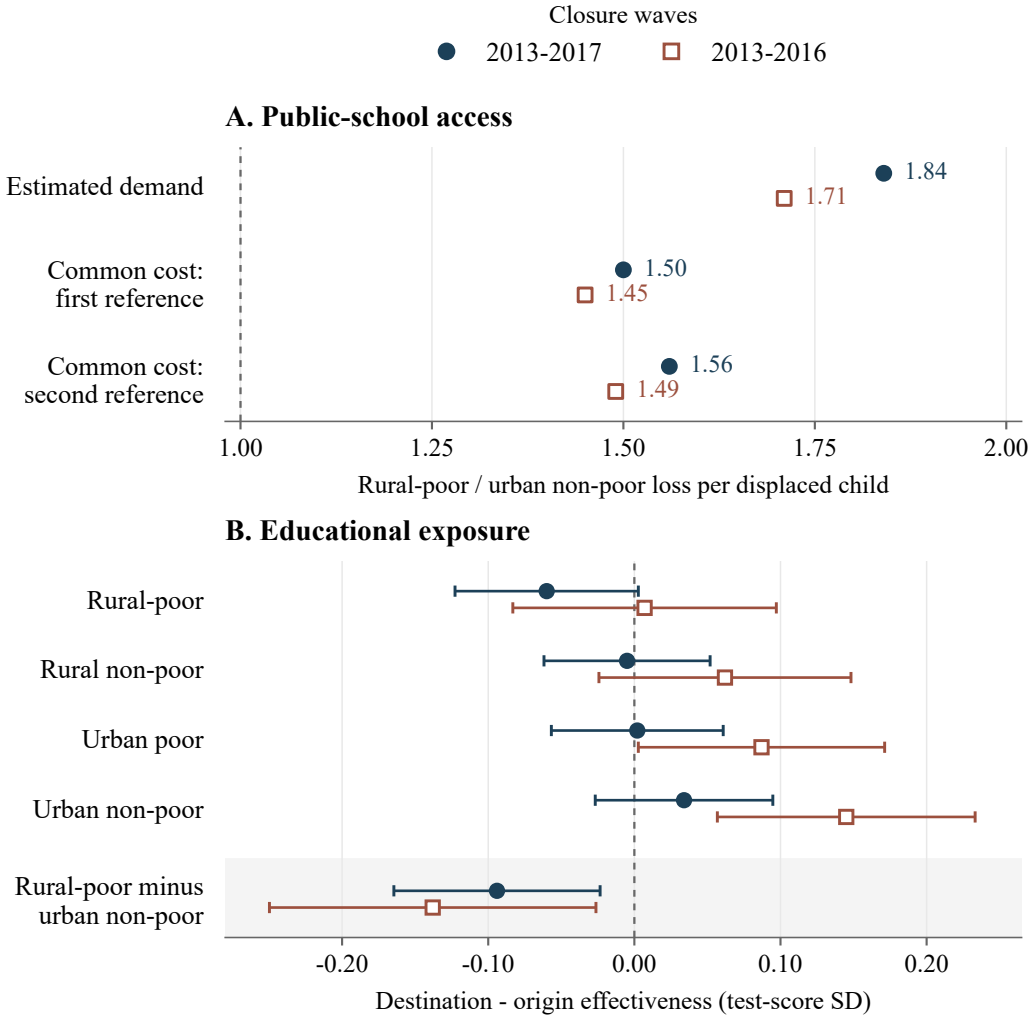
The counterfactuals vary three margins through which network contraction affects families: the set of schools removed, the allocation of displaced teachers, and substitution among the surviving schools. They distinguish changes in the school-specific inputs available locally from changes in how portable teacher inputs and students are allocated. I evaluate public-school access and exposure to educational inputs separately. A reduction in access loss need not imply an improvement in educational exposure, and closing the same number of buildings need not produce the same fiscal savings.

## 7.1 Comparing closure rules

The closure-rule exercise holds fixed the number of schools closed in each municipality, closure wave, and enrollment-size band, and changes which schools close. The size bands are terciles of the enrollment of the schools actually closed in that wave. Candidate rules use lagged enrollment, grade offerings, location, and, where specified, prior test scores or score growth. Rules based on estimated school effects provide an additional benchmark, but those estimates use information unavailable to the Department when it made the decisions. They should therefore be distinguished from rules that could have been implemented with contemporary administrative records.

I consider closing the smallest schools first, adding a proximity screen to that ordering, selecting closures with the smallest increase in enrollment-weighted distance to a surviving alternative, and using prior achievement to protect schools. The proximity screen requires each grade offered by a candidate school to remain

**Figure 2.** Distributional Incidence




*Notes:* Panel A compares average model-implied losses of public-school access per displaced child, expressed in kilometer equivalents. Rurality and poverty are household types in the demand model. The common-cost scenarios re-match income composition and school shares; their construction is reported in Table 5. Panel B reports destination minus origin effectiveness, including school and teacher contributions, for observed displaced students. Rural and urban comparisons use the origin's fractional rural exposure and its complement; poverty is the child's administrative indicator. Effectiveness compares the destination in the year after closure with the origin in its last operating year. Bars are normal-approximation 95% intervals from municipality-clustered standard errors. The shaded row reports the rural-poor minus urban non-poor contrast. These exposure changes are not causal effects on test scores. The panels use distinct populations and units; Table 6 gives the exposure definitions. Closure waves denote the last operating year.

available within a specified distance. When too few schools satisfy the screen, the procedure completes the required count using schools with the nearest remaining alternatives. The resulting policy is consequently a preference for proximity, rather than an unconditional distance guarantee. Each sequential rule updates the remaining network after a school is removed.

Comparisons are made one wave at a time, taking earlier actual closures as given. Summing these comparisons does not describe a policy implemented continuously over the whole program: keeping a school open in one year could change the choices available in a later year. The calculations also hold surviving schools' estimated contributions fixed. They do not model how different enrollment flows would change school organization, peer composition, staffing, or outcomes for incumbent students.

For production, the comparison asks how the schools removed by a rule compare with the surviving local menu. Let $E^g_{jc\ell}$ denote the number of children in group $g$ and grade $\ell$ enrolled at school $j$ before closure wave $c$, and let $C_c$ be a candidate set of closures. Define

$$Q_g(C) = \frac{1}{H_g^{\text{act}}} \sum_c \sum_{j \in C_c} \sum_\ell E^g_{jc\ell} \left[ \widehat{\phi}_j - \overline{\phi}_{a(j),c,\ell}(C_c) \right], \tag{15}$$

where $a(j)$ is school $j$'s municipality and $\overline{\phi}_{a,c,\ell}(C_c)$ is the enrollment-weighted average school effect of surviving schools in municipality $a$ serving grade $\ell$. The denominator $H_g^{\text{act}}$ is the number of children in group $g$ displaced by the actual closures, and is held fixed across rules. Thus a difference in $Q_g$ compares totals on a common scale; it does not reward a rule merely for displacing more children. A positive value means that the schools removed have larger estimated school effects than their surviving local alternatives. This is an index of school contributions removed relative to a benchmark, rather than a prediction of which school each child would enter or a treatment effect on achievement.

The corresponding access comparison aggregates the index defined in Section 4 for each closure set. A separate geographical measure weights the distance from each closing school to its nearest surviving grade-eligible school by the children displaced. This second measure does not require estimated demand, but it measures separation between schools rather than a change in travel from children's homes.

Table 7 reports the previous run's numerical comparisons for the 2013–2017 closure waves. These estimates are preliminary: calculations using corrected grade enrollment and historical-enrollment headroom are pending. The correction can change both the selected closure sets and the weights used to aggregate their consequences. Appendix D describes the correction and the associated limitations. The results below describe the previous run rather than a verified ranking under the revised enrollment construction.

Closing the smallest schools first gives different answers on the two access measures. For rural-poor

**Table 7.** Closure-Rule Comparisons

| | Nearest-school loss (%) | Model access loss (%) | School effects removed (SD) |
|---|---|---|---|
| *Panel A: Rural-poor children* | | | |
| Smallest first | −1.1 | 3.1 | 0.125 |
| Smallest first, 4.8 km screen | −25.4 | −10.5 | 0.101 |
| Smallest first, 3.2 km screen | −36.4 | −15.6 | 0.075 |
| Nearest alternative first | −45.4 | −24.9 | 0.054 |
| Smallest first, 4.8 km, protect score growth | −23.8 | −7.4 | 0.066 |
| Lowest average scores | −0.7 | 7.1 | −0.063 |
| Lowest average scores, 4.8 km screen | −20.4 | −8.0 | −0.055 |
| *Panel B: All displaced children* | | | |
| Smallest first | −6.4 | 5.9 | 0.112 |
| Smallest first, 4.8 km screen | −18.8 | −4.4 | 0.095 |
| Smallest first, 3.2 km screen | −27.3 | −8.7 | 0.078 |
| Nearest alternative first | −38.4 | −14.4 | 0.059 |
| Smallest first, 4.8 km, protect score growth | −16.7 | 0.3 | 0.082 |
| Lowest average scores | 4.7 | 16.1 | −0.100 |
| Lowest average scores, 4.8 km screen | −8.7 | 3.8 | −0.092 |

Notes: Preliminary results from the previous run for closure waves 2013–2017, recorded on September 28, 2026. Calculations using corrected grade enrollment and historical-enrollment headroom are pending. The correction can change both the schools selected and the weights used to evaluate them, so the reported rankings and magnitudes require validation. The first two columns are percentage differences in total access loss from the actual closures. The final column is the rule's school-effects-removed index minus the actual index, in student-score standard deviations per child displaced by the actual closures, using a common denominator across rules. Positive values indicate that the rule removes larger school contributions relative to surviving local schools. The actual rural-poor index is 0.053, or 0.063 under the held-out evaluation used for rules selecting on achievement. Size-band closure counts are fixed within municipality and wave. Proximity screens include a fallback to complete those counts and are not strict guarantees. Rules based on average scores were added after initial results were seen. The table records point estimates only; it does not establish statistical improvements, a feasible frontier, or equal fiscal savings.

children, it reduces the nearest-school measure by 1.1 percent but raises total model-implied access loss by 3.1 percent relative to actual closures. Adding a 4.8-kilometer proximity screen reduces the two measures by 25.4 and 10.5 percent, respectively. These rules also increase the rural-poor school-contribution index by 0.125 and 0.101 test-score standard deviations per child displaced by the actual closures. A smaller school is therefore not necessarily a close substitute for the other schools in its municipality, nor does its enrollment summarize its estimated contribution to achievement.

The nearest-alternative rule yields a larger access reduction in the previous run. Rural-poor total model-implied access loss falls by 24.9 percent, and the nearest-school measure falls by 45.4 percent. The school-contribution index increases by 0.054 standard deviations. For all displaced children, the corresponding model access reduction is 14.4 percent and the index increase is 0.059 standard deviations. This rule thus reduces access losses while removing larger estimated school contributions relative to the surviving local menu. The positive school-contribution differences are changes in the index defined in (15), not predicted declines in students' test scores. All percentage comparisons refer to totals, whereas the school-contribution

comparisons use actual displacement as their common denominator; the alternative rules need not displace the same children.

The exploratory rule combining low prior average scores with a 4.8-kilometer proximity screen produces a different distribution of changes. For rural-poor children, total model-implied access loss falls by 8.0 percent and the school-contribution index falls by 0.055 standard deviations. For all displaced children, however, model-implied access loss rises by 3.8 percent, while the school-contribution index falls by 0.092. The municipality-bootstrap interval for the rural-poor school-contribution difference includes zero (Appendix D). These point estimates consequently do not establish an improvement for every group or a statistically established improvement in both rural-poor outcomes. They also preclude a general claim that every rule trades access against school contributions.

These comparisons hold closure counts fixed, not fiscal savings. They describe selected rules under the stated restrictions and do not establish an optimal closure plan, a feasible frontier, or that the actual closure set cannot be improved upon.

### 7.2 Reassigning displaced teachers

The teacher exercise instead holds the actual closures and the students served by each receiving post fixed. It reassigns displaced teachers among the posts that displaced teachers actually filled, within municipality, year after closure, subject profile, and grade band. The two grade bands are grades 3–6 and 7–8. The eligible pool consists of 626 displaced teachers with one observed open-school destination. Keeping this pool and these posts fixed isolates the allocation of teacher value-added from changes in staffing levels or the set of open schools. Of the 434 assignment blocks, 303 contain a single teacher. Thus 48.4 percent of the pool has no within-block reassignment margin, which limits the scope for gains under these restrictions.

The assignment maximizes predicted teacher value-added delivered to the students in these posts, counting rural-poor students twice. The weight of two is a distributional objective specified for the exercise, not an estimated preference parameter. The assignment respects the stated subject and grade restrictions but abstracts from teacher preferences and transfer priorities. Its gains are therefore conditional on a reassignment being administratively feasible under those restrictions.

Table 8 reports a gain of 0.0007 test-score standard deviations per displaced rural-poor child, with an interval of $[-0.003, 0.004]$. The gain is small and is not statistically detectable. This result concerns a narrow margin: reallocating the observed pool of displaced teachers across the observed receiving posts. It does not establish that teacher assignment throughout Puerto Rico is efficient or that changing the overall distribution

of teachers would have little effect. Within this exercise, however, teacher reassignment contributes little to reducing the disadvantage associated with the actual closures.

### 7.3 Information and destination choices

The final exercise asks how destination choices would change if displaced families placed additional weight on the estimated effectiveness of surviving schools. The baseline demand estimates do not identify either families' valuation of effectiveness or the information they possess. The exercise consequently specifies a response to information and examines its implications; it does not estimate the treatment effect of providing a report card.

The report card is constructed from estimates of $\Psi_{jc} = \phi_j + q_{jc}$, which include both the school contribution and the average value-added of its teachers in the closure year. One part of the data supplies the card and the other evaluates the effectiveness of the resulting allocation; the roles are then reversed and averaged. Families are assumed to receive and act on the card, seats are unlimited, and other families do not re-sort. All other school characteristics and preferences remain fixed. These assumptions define an optimistic information scenario. They do not establish an identified upper bound on the effect of any information intervention.

The response parameter is expressed as a willingness to travel for a one-standard-deviation increase in effectiveness. I report a moderate calibration of 1.1 km and a larger calibration of 13 km, with the same additional utility coefficient applied to all household groups. The conversion uses the rural-poor distance coefficient at the reference distance of 3.9 km, so the kilometer interpretation is specific to that group and reference point. The larger calibration is consistent with converting the elementary-school estimates in Harris and Larsen (2022) to student-score units; it is a sensitivity case rather than an estimate for Puerto Rico. Appendix D gives the conversion and the distinction between a common utility coefficient and a common willingness to travel.

At the 1.1 km calibration, the estimated improvement in the effectiveness of the modeled destinations of rural-poor children is 0.047 standard deviations, with a municipality-bootstrap interval of [0.037, 0.058]. The average modeled change in home-to-school distance is 0.004 km. At the 13 km calibration, the effectiveness gain is 0.316 [0.240, 0.385], with distance increasing by 1.60 km. These gains compare the model's postclosure allocations with and without the additional weight on the card. Observed displacement records supply the origin and poverty composition used to aggregate them; the baseline is not each child's observed destination.

These scenarios show that the effectiveness of destinations can respond to how families evaluate surviving schools, although the magnitude is sensitive to an unestimated information response. They do not show that

**Table 8.** Teacher Assignment and Information

| | Rural-poor | | Urban non-poor | |
|---|---|---|---|---|
| | Gain (SD) | Distance (km) | Gain (SD) | Distance (km) |
| *Teacher reassignment* | | | | |
| Rural-poor weight of two | 0.0007 | — | 0.0010 | — |
| | [−0.003, 0.004] | | [−0.005, 0.007] | |
| *Information scenarios* | | | | |
| 1.1 km calibration | 0.047 | 0.004 | 0.062 | 0.019 |
| | [0.037, 0.058] | [−0.025, 0.036] | [0.041, 0.087] | [−0.014, 0.057] |
| 13 km calibration | 0.316 | 1.603 | 0.395 | 1.514 |
| | [0.240, 0.385] | [1.205, 1.963] | [0.250, 0.540] | [1.105, 1.958] |

Notes: Closure waves 2013–2017. Teacher reassignment holds students, posts, and the displaced-teacher pool fixed and maximizes predicted value-added subject to municipality, year, subject-profile, and grade-band restrictions. Its gain is average teacher value-added delivered per displaced child; the brackets report the composite uncertainty interval for that exercise. Information gains are changes in modeled destination effectiveness relative to the model's no-additional-information allocation. Reported distance changes are modeled home-to-school distances. The report card uses one school-year estimation half and is evaluated on the other, then the roles are reversed. Aggregation uses 12,564 eligible displacement records from a frame of 15,025, with origin-school rural exposure and individual poverty. Information intervals resample municipalities, holding the selected cards and demand parameters fixed. The calibration converts a common utility coefficient using rural-poor distance costs at 3.9 km; it does not impose the same kilometer valuation for every group. Seats are unlimited and all targeted families act on the card. The teacher and information rows have different estimands and samples and are not added.

information alone would deliver the calculated gains. Admission limits, transport arrangements, changes in school inputs, and incomplete take-up could change the response substantially. A related exercise selects the most effective surviving school no farther from the origin than the destination actually attended. It separates the availability of alternatives from assumptions about preferences, but its enrollment-based seat screen requires the same reconciliation as the closure-rule comparisons. It is therefore retained as a preliminary exercise in Appendix D.

## 7.4 Fiscal accounting

The distributional comparison can be extended to alternatives evaluated at a common fiscal cost. For illustration, suppose a planner must achieve a specified savings target while choosing schools and assignments. Let $S(C,\alpha)$ denote net fiscal savings from a closure set $C$ and an assignment policy $\alpha$, relative to maintaining the existing school network. The assignment policy includes the placement of teachers and pupils. For a required saving $\overline{S}$, the feasible policies must satisfy

$$S(C,\alpha) \geq \overline{S}, \qquad (C,\alpha) \in \mathscr{F}, \tag{16}$$

where $\mathscr{F}$ imposes grade coverage, receiving-school capacity, and feasible teacher assignments. The relevant comparison is then the set of access and educational outcomes attainable under this constraint. Selecting one point from that set requires a social objective or explicit priorities over groups and outcomes. The present analysis does not estimate those priorities or solve this optimization problem.

The closure rules above hold counts fixed within municipality, year, and size bands. This is useful for comparing changes in selection within a similarly sized program, but it does not impose (16). Schools with similar enrollment can differ in avoidable facilities costs, staffing requirements, and transportation needs. Receiving-school expenditure can also change when pupils are redistributed. Consequently, neither equal closure counts nor equal numbers of displaced pupils imply equal fiscal savings. By contrast, the teacher exercise holds the pool of displaced teachers and destination posts fixed, so that their wage bill cancels under the maintained assumption that reassignment does not change compensation. Other implementation costs need not cancel.

Appendix E describes a separate accounting exercise for facilities and staffing. Its conditional component calculations are not net savings for each simulated closure rule. They cannot be used to convert the access comparisons into a benefit–cost ratio or to rank rules at a common budget. For that purpose, the most useful additional information would be a cost and capacity schedule for candidate receiving schools, the positions that would remain unfilled after reassignment, and the transportation routes required by each plan. With these inputs, the framework could compare policies at the same savings target while retaining the separate incidence of access and educational exposure.

# 8 Conclusion

This paper asks how the contraction of a spatial school network changes access and educational inputs across families with different opportunities to substitute among schools. A closure removes a school-specific contribution to production, while teachers can continue contributing at other schools. Families' choices then determine their exposure to the inputs that remain. Separating these margins makes it possible to examine access and educational exposure within one framework, without assuming that test-score units and travel distances have a common social value.

The Puerto Rican evidence shows a substantial rural-poor disadvantage in access. The model-implied loss per displaced child is 1.84 times that of urban non-poor children, and the difference persists when distance costs are held common. The comparative change in educational exposure is also less favorable for the rural-poor group. In observed transfers, the difference from the urban non-poor group in destination–origin

effectiveness changes, including teacher contributions, is about $-0.09$ test-score standard deviations. The difference remains negative when the final closure wave is excluded, although the rural-poor group's own change is then close to zero. These are changes in exposure to estimated educational inputs, not treatment effects on students' achievement.

The policy exercises distinguish changing which schools close from changing how the remaining schools and teachers are used. Within the restricted set of local vacancies filled by displaced teachers, reassignment offers little estimated improvement for rural-poor students. The information exercise produces larger gains in educational exposure under a specified response to quality, but requires strong assumptions about take-up and available seats. It therefore identifies a margin worth examining, rather than establishing the effect of a particular information program.

The broader implication is that the incidence of a contraction depends on both the inputs removed and households' substitution opportunities. Enrollment alone does not summarize either margin. A comparison of alternative networks at a common budget would additionally require comparable receiving-school costs, transportation costs, and feasible assignments. The present estimates establish neither an optimal closure plan nor a common valuation of access and achievement; they show why the allocation of educational inputs and families' access to those inputs must be considered together.

# References


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

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Berry, S., Levinsohn, J., and Pakes, A. (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112(1):68–105.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262.

Bobonis, G. J., Sotomayor, O. J., and Wagner, J. (2022). The effects of school consolidation on students and teachers: Evidence from an underperforming system. Working paper, University of Toronto, December 2, 2022.

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

Cardell, N. S. (1997). Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Econometric Theory*, 13(2):185–213.

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

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

Hall, A. R. and Inoue, A. (2003). The large sample behaviour of the generalized method of moments estimator in misspecified models. *Journal of Econometrics*, 114(2):361–394.

Harris, D. N. and Larsen, M. F. (2022). What schools do families want (and why)? evidence on revealed preferences from new orleans. Technical report, Education Research Alliance for New Orleans. Technical report, updated June 12, 2022.

Jackson, C. K. (2018). What do test scores miss? The importance of teacher effects on non–test score outcomes. *Journal of Political Economy*, 126(5):2072–2107.

Kapor, A. J., Neilson, C. A., and Zimmerman, S. D. (2020). Heterogeneous beliefs and school choice mechanisms. *American Economic Review*, 110(5):1274–1315.

Kline, P., Saggio, R., and Sølvsten, M. (2020). Leave-out estimation of variance components. *Econometrica*, 88(5):1859–1898.

Mansfield, R. K. (2015). Teacher quality and student inequality. *Journal of Labor Economics*, 33(3):751–788.

National Hurricane Center (2017). Hurricane Maria tropical cyclone update. National Oceanic and Atmospheric Administration, 6:35 a.m. AST, September 20, 2017.

NCES (2021a). Enrollment in public elementary and secondary schools, by region, state, and jurisdiction: Selected years, fall 1990 through fall 2030. National Center for Education Statistics, Digest of Education Statistics 2021, Table 203.20. Table prepared March 2022; historical enrollment from the Common Core of Data.

NCES (2021b). Public elementary and secondary schools, by level, type, and state or jurisdiction: 1990–91, 2000–01, 2009–10, and 2019–20. National Center for Education Statistics, Digest of Education Statistics 2021, Table 216.70.

Nelson, C. D., Prado Tuma, A., Marsh, T., Andrew, M., Anderson, D. M., Whitaker, A. A., Karoly, L. A., Murphy, R. F., Nanda, N., Ryan, J. L., Smith, T. D., and Chandra, A. (2020). The education sector in Puerto Rico after Hurricane Maria: Predisaster conditions, hurricane damage, and themes for recovery. Research report RR-2858-DHS, Homeland Security Operational Analysis Center operated by the RAND Corporation.

Newey, W. K. (1984). A method of moments interpretation of sequential estimators. *Economics Letters*, 14(2–3):201–206.

PRDE (2015). Política pública sobre las normas y los procedimientos para el cierre, la consolidación o la reestructuración de escuelas. Carta Circular 21-2015-2016, Puerto Rico Department of Education. November 17, 2015.

PRDE (2016a). Orientación a los estudiantes: Medición y evaluación para la transformación académica de puerto rico (meta-pr), año escolar 2015–2016. Technical report, Puerto Rico Department of Education. April 12.

PRDE (2016b). Perfil escolar: Año académico 2016–2017. Technical report, Puerto Rico Department of Education. December 2; see p. 4 on the introduction of META-PR.

PRDE (2017a). Ffy 2015 part b state performance plan and annual performance report. Technical report, Puerto Rico Department of Education. Clarifications version, May 3; pp. 1, 28, and 121.

PRDE (2017b). Política pública para la reubicación, reasignación, traslado y reclutamiento del personal docente de las escuelas, de los institutos tecnológicos, la escuela de troquelería y herramentaje y la PR Aviation Maintenance Institute en el departamento de educación. Carta Circular 31-2016-2017, Puerto Rico Department of Education. January 27, 2017.

PRDE (2017c). Procedimiento general para el rediseño de escuelas del sistema público de Puerto Rico. Carta Circular 33-2016-2017, Puerto Rico Department of Education. March 10, 2017.

PRDE (2017d). Proceso de matrícula en línea: Entrada y actualización de datos al nuevo sistema. Memorandum and attached director manual, Puerto Rico Department of Education. August 3, 2017; manual for school year 2017–18.

Rothstein, J. (2010). Teacher quality in educational production: Tracking, decay, and student achievement. *Quarterly Journal of Economics*, 125(1):175–214.

Steinberg, M. P. and MacDonald, J. M. (2019). The effects of closing urban schools on students' academic and behavioral outcomes: Evidence from Philadelphia. *Economics of Education Review*, 69:25–60.

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

Train, K. E. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press, Cambridge, 2nd edition. ISBN 9780521766555 (hardback); 9780521747387 (paperback).

Walters, C. R. (2018). The demand for effective charter schools. *Journal of Political Economy*, 126(6):2179–2223.

# A Data Construction and Samples

This appendix describes the calendar, linkage, and geographic definitions used in the paper, then gives additional details on the enrollment event study. The main distinction is between the administrative universe, the population represented by the choice model, and the pupils for whom educational inputs can be measured. Restrictions that are appropriate for one object need not define another.

## A.1 Academic years, closures, and grades

An academic-year identifier denotes the calendar year in which the academic year ends. Closure year $c$ is the final operating academic year of the origin school; the next enrollment is observed in $c+1$. Thus, closure year 2017 means final operation in 2016–17 and next-year attendance in 2017–18. The principal choice and counterfactual windows cover $c = 2013, \ldots, 2017$, while the earlier comparison covers $c = 2013, \ldots, 2016$. A table labeled by closure year is therefore not labeled by the year in which the displaced pupil's destination is recorded.

The school operation file and the pupil displacement file provide related information. The former identifies schools that cease operating; the latter links a pupil's origin and subsequent enrollment. A school can enter the first frame without contributing an observed destination to a particular pupil sample. Changes in school codes, grade offerings, and observed enrollment also affect whether an origin appears in a given municipality–grade–year choice set. For this reason, system totals, dated schools in the operation file, and origins represented in the destination sample are not interchangeable closure counts.

Grade codes are converted to actual grades before imposing sample restrictions. In the raw fields used to construct the analysis, codes 5–10 correspond to grades 3–8 and codes 3–10 to grades 1–8. This transformation is made once, at the input-field level; variables already expressed as actual grades are retained. Enrollment grade determines progression and the choice market, whereas test grade determines the score-normalization cell. A discrepancy between the two does not by itself justify dropping a pupil's teacher link.

The terminal-grade restriction serves a separate purpose. A move beyond the grades covered by the choice network does not reveal a second choice among the alternatives represented in that network. Such transitions are excluded from the destination moments. They should not be recoded as exit simply because the subsequent grade is absent from a restricted analytical file.

## A.2 Scores and instructional links

The achievement observation is a pupil–subject–year. Pupil identifiers and academic years link enrollment and scores to the student–teacher file; the subject-specific teacher fields are then reshaped to identify instructional exposure. Pupil counts, teacher counts, teacher–subject effects, and observation counts consequently have different denominators. The main production model requires usable scores and links as well as the variation

needed to separate student, school, and teacher–subject effects. The estimation and leave-out restrictions are described in Appendix B.

PRDE introduced META-PR in 2015–16 after the PPAA. Its 2016 instructions identify mathematics, Spanish, and English as tested subjects in grades 3–8 and 11; science was tested in selected grades (PRDE, 2016a). The administrative achievement panel used here includes mathematics, Spanish, and English in grades 3–8. In each subject–grade–year cell, I subtract the mean score and divide by its standard deviation. The change in assessment is therefore accommodated in the normalization, while the stability of school and teacher contributions across years remains an empirical and modeling issue. The normalization alone does not establish that the two tests measure an identical construct.

The ANCOVA comparison additionally requires prior achievement. Since grade 3 has no preceding tested grade in this panel, that specification uses grades 4–8 with the required lagged scores. It cannot be interpreted as re-estimating the pupil-fixed-effect specification on exactly the same population.

### A.3 Population weights and distances

PRCS block-group counts provide the geographic distribution of the school-age population. The demand market is a municipality–grade–year, and simulated households combine possible block-group residences with poverty types. Their weights describe the population represented in a market. Aggregation across markets also accounts for market size; an equally weighted average of markets would instead give a small grade cohort the same weight as a large one. The market-size normalization and its sensitivity are part of the choice-model specification in Appendix C.

The PRCS Public Use Microdata Sample supplies a separate participation moment. It includes public enrollment, private enrollment, and non-enrollment among children aged 8–14 in 2013–17. The proportions in Table 2 use the survey weights and classify income poverty using the Census threshold. This source measures population participation; it is not linked child by child to the administrative pupil file.

The pupil data identify residence only at a coarser geographic level than the block group. Conditional choice probabilities therefore integrate over possible residences rather than treat a ZIP centroid as the pupil's home. Population weights are used before attendance is observed; conditional on attending origin $o$, they are updated by the probability of choosing $o$. This distinction is important because the residential composition of an origin school need not equal the population composition of its entire municipality.

The available school-to-school and block-group-to-school routes are measured in road kilometers using OSRM on the OpenStreetMap network. Counterfactual distance approximations and their distinction from these routes are described in Appendix D. The nearest surviving school is defined within the specified geographic boundary and eligible grade set. Changing that boundary changes the comparison: the nearest school anywhere on the island need not be the nearest school in the origin municipality. Neither distance directly measures fares, time on a school bus, or the change in the pupil's commute from home. The estimated distance responses reflect the transport conditions present in the enrollment data.

## A.4 Rural exposure and poverty

The residential rural indicator is based on block-group density, calibrated to the Department's rural designation for schools. This classification supplies the rural indicator for simulated households. It is distinct from the Census urban-share measure and from treating the Department's school designation as an observed household characteristic.

For the destination moments, let $A^R(o)$ denote the population-weighted rural share around origin $o$, using distance weights proportional to $\exp(-d/2)$ when distance $d$ is measured in kilometers. The weighted population and the choice-model construction are described in Appendix C. An observed pupil at $o$ receives rural weight $A^R(o)$ and urban weight $1 - A^R(o)$ in the rank moments. These fractional assignments avoid imposing a precisely observed rural residence where none is recorded. The pupil-level school-effect comparisons use the same fractional construction. For a pupil with administrative poverty indicator $P_i$ and base weight $w_i$, rural-poor weight is $w_i A^R(o) P_i$ and urban-poor weight is $w_i[1 - A^R(o)]P_i$, with analogous weights for the non-poor. The two geographic groups can contain the same contributing records; their weighted means are not means for two disjoint sets of pupils. The exit results retain the origin school's administrative rural designation so that their definitions coincide with the event-study estimation sample. None of these constructions implies that every pupil attending a school in a rural catchment lives in a rural block group.

The administrative poverty flag is observed with the pupil records, including the tested-grade population. Where the flag is missing, the analysis uses the pupil's modal observed status across enrollment years. This retains a time-invariant classification for imputed observations and relies on poverty status being sufficiently persistent for the comparison. It does not supply a new observation of family income. The PRCS income-poverty measure is kept separate: it determines population composition and the participation shifter, while the administrative flag is used for the pupil-level poverty grouping and the distance interaction.

## A.5 Sample differences

The choice model uses a wider grade range than the achievement model because enrollment is observed before standardized testing begins. The current demand sample contains 22,690 displaced-pupil records, with 16,327 destinations inside the specified choice sets. Origins or destinations can fall outside these sets because of grade transitions, missing origin-market links, cross-municipality attendance, or exclusions imposed on markets whose school attributes change. Being outside a destination-moment sample does not mean that a pupil left the public system. The exit outcome is constructed from next-year enrollment before using the narrower destination-rank sample.

Counterfactual access statistics are weighted by the population appropriate to the measure. The model follows simulated households and their predicted initial attendance. Observed origin–destination changes follow pupils with usable school-effect estimates at both schools. Grouping and weighting are specified with each comparison, rather than replacing an observed pupil denominator with predicted displacement. The restrictions for the reassignment and information exercises are given in Section 7.

Repeated school enrollment totals require particular care. A school–year total is not a school–grade enrollment count even when it appears on a row identified by grade. I do not sum this column over grades or geographic rows when describing school size. Grade-specific weights and capacity comparisons require their own enrollment construction; the limitations of the available counterfactual weights are stated alongside the affected exercises.

### A.6 The enrollment event study

The exit analysis follows pupils around the closure of their origin school and compares them with matched pupils not displaced by a closure. The treatment cohorts are schools last operating in 2013–16. Pupils must be enrolled and resident in the origin municipality in the pre-displacement year and in a nonterminal origin grade, with the delivered estimation sample covering grades 3–7. Event time zero is the first postclosure year, and event time $-1$ is the reference period. The sample contributes 7,138 treated observations at event time zero.

Let the first outcome indicate that a pupil is not observed in Puerto Rico's public enrollment records in the following year. The second indicates observed public enrollment in a different municipality. These events are mutually exclusive. Their sum indicates departure from the origin municipality's public schools and is the outcome corresponding to the choice model's outside option. All three outcomes use the same underlying event-study construction, so their estimated effects add when evaluated on the same sample with the same weights.

I use coarsened-exact-matching weights and cohort-by-event-time interactions, aggregated following Sun and Abraham (2021). The specification absorbs pupil-by-matching-cohort and matching-cohort-by-calendar-year effects, and clusters standard errors at the administrative district level. Identification requires that, conditional on the matching and fixed effects, treated pupils would have followed the same trends as the comparison pupils without closure. It also requires that the reference period is not already affected by an anticipated closure. The earlier cohorts avoid the destination year spanning Hurricane María; this restriction does not by itself establish either assumption.

At event time zero, the effects are 3.20 percentage points on absence from the public panel (standard error 0.89), 5.32 points on cross-municipality public enrollment (1.21), and 8.52 points on their sum (1.47). The pooled effect on the sum is used as the demand model's exit target. It is an incremental response relative to the matched comparison, not the gross fraction of displaced pupils observed outside the municipality. Equating it to the model's closure-induced outside substitution requires the bridge assumption stated in Section 5.2.

The supporting checks do not eliminate all uncertainty about this comparison. The joint test of the estimated leads from event times $-4$ through $-2$ does not reject, but this is not evidence sufficient to establish parallel trends. A separate sample-comparability check required the treated means to reproduce an earlier benchmark within one percentage point, or for the difference to be explained by excluding terminal grades. It did not: the treated proportions for absence, crossing, and their sum were 0.084, 0.069, and 0.153, compared

with 0.105, 0.068, and 0.173 in the earlier benchmark. Removing the residence restriction changes them to 0.107, 0.089, and 0.196, which also fails the comparison as specified. I retain the resident sample because it matches the model's intended population, while reporting the failed comparison as a limitation rather than treating the two samples as interchangeable.

Alternative exit targets show the sensitivity to these definitions. Including the 2017 closure cohort yields an effect of 7.6 percentage points, and including all origin grades yields 6.8 points. Removing the residence restriction yields 10.6 points on a different population. The effects for grades 3–4 and 5–7 are close, at 8.57 and 8.47 points, but they cannot establish the corresponding effect for grades 1–2. Applying the pooled target to the wider choice frame therefore retains a population-transport assumption. The model's fit to geographic exit differences is assessed separately in Appendix C.

# B School and Teacher Effects: Estimation Details and Robustness

This appendix describes the correction for estimation error, the independent-sample comparisons, and the achievement robustness exercises. The primary equation is (1). The production program uses the scored panel from 2011 through 2019; the closure cohorts used in subsequent comparisons are narrower. Changing the closure window does not by itself change the window over which school and teacher effects were estimated.

## B.1 Estimation Sample and Normalization

The sample restriction retains the largest component of schools linked by a shared teacher–subject identifier and requires at least two scored records per student fixed effect. These are records, not necessarily two academic years. A teacher who teaches different subjects at two schools does not, through that change alone, link the same teacher–subject effect across schools. I alternate the connected component restriction and the minimum-record restriction until neither changes the sample, and reapply both after pruning observations with high estimated leverage. Student mobility enters the within-student estimating design, but the preliminary school-component screen is constructed from teacher–subject links. This screen does not by itself establish full rank of the joint student, teacher–subject, and school design.

The computation absorbs student effects by within-student demeaning. Writing $Z$ for the teacher–subject and school indicator matrix and $M_a$ for the demeaning operator, it solves the remaining normal equations with an isotropic ridge of size $10^{-8}$ times the mean diagonal of $Z'M_aZ$. This selects the minimum-norm normalization in unidentified directions, including centering the teacher–subject and school effects separately with unweighted means rather than omitting a reference teacher or school. Student effects are then recovered as student-specific means of the residual after subtracting the teacher and school contributions. The small ridge also regularizes near-singular directions; it does not supply economic identification where the design lacks it. Leverage calculations use random-projection approximations and a separate final set of probes. The exact

generalized-inverse formulas below state the underlying leave-out identity; the numerical implementation approximates that calculation.

## B.2 Variance Components

Write $y = X\beta + e$, with $X$ containing student, teacher–subject, and school indicators and $\beta = (a', \mu', \phi')'$. Under the minimum-norm normalization, let $H = (X'X)^{\dagger}X'$, $P = XH$, and $\hat{\beta} = Hy$; the dagger denotes the generalized inverse. A variance component is a quadratic form $\theta_A = \beta' A \beta$, with its weighting population encoded in $A$. Teacher–subject dispersion and the dispersion of school-level roster means are different quadratic forms and must not be interchanged. Conditional on the design and a mean-zero error,

$$\mathbb{E}[\hat{\beta}' A \hat{\beta} \mid X] = \beta' A \beta + \operatorname{tr}(B_A \Sigma), \qquad B_A = H' A H, \quad \Sigma = \operatorname{Var}(e \mid X). \tag{B.1}$$

With independent record-level errors and leverage $P_{rr} < 1$, the leave-out variance estimate $\hat{\sigma}_r^2 = y_r \hat{e}_r / (1 - P_{rr})$ estimates the error variance for record $r$ (Kline et al., 2020). An individual estimate can be negative even though the population variance is nonnegative.

Independence across records is not appropriate here. The implemented error structure is

$$e_r = \nu_r + \zeta_{c(r)} + \eta_{j(r),t(r)}, \tag{B.2}$$

where $c(r)$ is a teacher–subject–school–year classroom cell. The school-year term induces covariance even across different classrooms in the same school. For calibrated common-shock variances $v_\zeta$ and $v_{\eta,jt}$, the adapted correction is

$$\hat{\theta}_A = \hat{\beta}' A \hat{\beta} - \sum_r (B_A)_{rr} \hat{v}_{\nu,r} - v_\zeta \sum_c \mathbf{1}_c' B_A \mathbf{1}_c - \sum_{j,t} v_{\eta,jt} \mathbf{1}_{jt}' B_A \mathbf{1}_{jt}. \tag{B.3}$$

Here $\hat{v}_{\nu,r}$ is the record-level variance estimate after removing the contribution of the common shocks, and $\mathbf{1}_c$ and $\mathbf{1}_{jt}$ select their respective cells. The calculation includes student-effect leverage after absorbing student indicators. Leverages and trace terms are approximated with random projections; the reported implementation is an adaptation of the leave-out method, not exact deletion of every classroom or school-year cluster.

The reference classroom variance is $v_\zeta = 0.10018$. The scalar school-year reference, 0.02218, comes from an empirical covariance calculation adjusted for attenuation measured in simulated data. The current correction allows this variance to differ with school size:

$$v_{\eta,jt} = \operatorname{clip}\big(\exp\{-0.9143 - 0.6072 \log n_j^{\text{score}}\}, 0.010, 0.040\big), \tag{B.4}$$

where $n_j^{\text{score}}$ is mean scored-student enrollment per school-year and clip bounds its first argument by the two stated endpoints. The inputs are therefore part of the maintained noise model. Sensitivity to their values matters, especially for teachers' school-level mean contributions.

At the reference calibration, the corrected standard deviations are 0.361 for school effects, 0.225 for teacher–subject effects, 0.089 for pooled school-level teacher means, and 0.393 for their school-level sum. Across the examined common-shock calibrations, the first three range from 0.34 to 0.38, 0.14 to 0.28, and 0 to 0.15, respectively. These are sensitivity ranges, not confidence intervals. They do not incorporate uncertainty about the assignment restriction or the policy invariance of school effects.

### B.3 School-Year Splits

Let $\hat{\phi}_j^A$ and $\hat{\phi}_j^B$ be estimates from separate school-year halves. Under a common underlying school effect and uncorrelated estimation errors,

$$\mathrm{Cov}(\hat{\phi}^A, \hat{\phi}^B) = \mathrm{Var}(\phi). \tag{B.5}$$

Analogous cross-products apply to teacher effects and school-level teacher means. The current implementation divides school-years across halves. Earlier classroom-only splits could place the same school-year shock in both estimates and are not the independent comparison used here. Serially correlated residual shocks, common model omissions, or systematic changes in effects can still invalidate (B.5).

Two school-year splits produce school-effect standard deviations of 0.303 and 0.307, school-level teacher-mean standard deviations of 0.075 and 0.063, and teacher–subject standard deviations of 0.203 and 0.191. A version excluding displaced pupils from the effect-estimation sample gives 0.311 for school effects and 0.078 for school-level teacher means, with reported intervals $[0.275, 0.344]$ and $[0, 0.124]$, respectively. These comparisons corroborate substantial school-level variation and much less precise measurement of the teacher-composition component. Because coverage, weights, and the error treatment differ, they are not additional draws from the same estimator as the full-sample leave-out calculation.

The covariance between school effects and teacher composition is imprecise and changes sign across corrections. I do not give it a signed economic interpretation. In particular, a small covariance estimate does not establish that more effective teachers are allocated independently of school inputs.

### B.4 Lagged-Score ANCOVA

The alternative achievement equation replaces student fixed effects with controls for lagged achievement:

$$Y_{ist} = X_{ist}^{\mathrm{anc}\prime}\gamma^{\mathrm{anc}} + \mu_{k(i,s,t),s}^{\mathrm{anc}} + \phi_{j(i,t)}^{\mathrm{anc}} + e_{ist}^{\mathrm{anc}}. \tag{B.6}$$

The controls contain grade-interacted cubic terms in lagged mathematics, English, and Spanish scores and student demographics. The specification uses grades 4–8 in the 2011–2019 panel with the required prior scores. Grade 3 lacks a preceding tested grade. Teacher effects remain subject-specific and school effects remain common across subjects.

The required conditional mean-zero restriction now conditions on $X^{\mathrm{anc}}$ and current assignments. Assignment may depend on observed prior achievement, but those controls must adequately account for the

student heterogeneity related to current assignment and achievement. Measurement error in prior scores and omitted changes in family circumstances remain possible concerns. Because lagged scores also reflect prior educational inputs, this conditional value-added specification need not estimate the same object as the levels model.

The teacher–subject estimates available under both specifications have correlation 0.849. This correlation uses estimated effects and the respective estimation samples, without correction for estimation error. It cannot establish robustness of the corrected variance decomposition, the demand estimates, or the policy comparisons.

### B.5 Portability and Forecasting

Teacher portability is the restriction that a teacher–subject effect does not change with the school of employment. Comparisons of residual performance across spells, with classroom and school-year noise taken into account, do not yield a positive residual match-variance estimate in the reference calculation. The analysis considers a match standard deviation between zero and 0.10 as a sensitivity range. This is not proof that match effects are absent, and the range is not reported as a conventional confidence interval. Reassignment simulations that omit match effects remain conditional on that restriction.

The pooled forecast exercise relates changes in displaced students' achievement to predicted changes in school effectiveness. Its attenuation-adjusted slope is 0.940, with standard error 0.011. It is informative about the combined component $\Psi$, conditional on the forecasting design and correction. It does not validate the school–teacher decomposition separately, and it does not establish forecast unbiasedness: the reported slope differs from one. Policy-induced changes in peers, congestion, or school organization would require additional responses beyond the fixed-effect equation.

## C Demand Estimation: Moments, Calibration, and Validation

This appendix records the demand moments, sequential estimator, and limitations relevant to the access comparisons. The estimated specification is (3); its free parameters are $(\widetilde{\lambda}_0, \widetilde{\lambda}_R, \widetilde{\lambda}_P, \pi, \rho)$. The school–market mean utilities are recovered from shares at each candidate vector.

### C.1 Destination Moments

For origin $o$ in market $m$, let $\mathscr{B}_{mo}$ be the set of eligible surviving inside destinations with the required geographic coverage. Other schools closing in the same market are excluded. Let $b(o, j)$ give the destination's distance-rank category measured from $o$, among surviving schools. For each household integration cell $z$, the implemented single-origin transition kernel is

$$K_{moj}(z;\theta) = s_{mj}(z;\mathscr{J}_m \setminus \{o\},\theta) - s_{mj}(z;\mathscr{J}_m,\theta). \tag{C.1}$$

Under the stable-ranking experiment that removes only $o$, this is the joint probability of initially choosing $o$ and subsequently choosing $j$. School mean utilities stay at their origin-year values.

For cell $c$, write $p_c$ for its administrative poverty status and let $\gamma_c(o)$ equal $A^R(o)$ for rural cells and $1-A^R(o)$ for urban cells. With $a_c(o,z) = \gamma_c(o)\mathbf{1}\{P = p_c\}$ and $\omega_{mo} = n^{\text{in}}_{mo}/S^{\text{obs}}_{mo}$, the model rank share is

$$M_{cb}(\theta) = \frac{\sum_{m,o} \omega_{mo} \int a_c(o,z) \sum_{j \in \mathscr{B}_{mo}} \mathbf{1}\{b(o,j) = b\} K_{moj}(z;\theta)\, dF_m(z)}{\sum_{m,o} \omega_{mo} \int a_c(o,z) \sum_{j \in \mathscr{B}_{mo}} K_{moj}(z;\theta)\, dF_m(z)}. \tag{C.2}$$

Here $n^{\text{in}}_{mo}$ is the displaced-origin count used for the inside-destination moments. The data share applies the same poverty, rural-exposure, survivor, and coverage restrictions to observed destinations. The four retained bins are ranks 2, 3, 4–5, and 6 or higher in each of four cells. The rank-one share is redundant.

The origin weights scale the model using observed enrollment and the relevant displaced count. After conditioning on inside re-enrollment, an origin's model contribution also depends on its predicted inside-retention probability. Thus, the weights do not make every origin's contribution exactly proportional to its displaced count for all parameter values. Excluding out-of-menu destinations further relies on the nested-logit conditional substitution structure.

Equation (C.1) is the actual rank-moment implementation. It is distinct from deleting all concurrent closers in (13). Even if proportional within-nest diversion holds conditional on $z$, removing additional schools can change the distribution of $z$ among retained inside choosers. Masking the other closers therefore need not reproduce simultaneous elimination after integration. The rank estimates maintain this approximation in markets with multiple closers. The exit and access calculations instead delete the whole closing set. This distinction is relevant to interpreting specification fit.

## C.2 Composition and Exit

The composition moment is

$$M^{\text{comp}}(\theta) = \frac{\sum_m M_m \int I(1 - s_{m0}(z;\theta))\, dF_m(z)}{\sum_m M_m \int (1 - s_{m0}(z;\theta))\, dF_m(z)}. \tag{C.3}$$

Its PRCS target is 0.6706, with standard error 0.0036. This income-poverty indicator differs from the administrative poverty indicator used to construct destination cells.

For exit, let $C_m$ be all closing inside schools and let $n^E_m$ count treated records in the exit-target population. Set $\omega^E_m = n^E_m / \sum_{o \in C_m} S^{\text{obs}}_{mo}$. The structural target is

$$A(\theta) = \frac{\sum_m \omega^E_m \int [s_{m0}(z; \mathscr{J}_m \setminus C_m) - s_{m0}(z; \mathscr{J}_m)]\, dF_m(z)}{\sum_m \omega^E_m \int \sum_{o \in C_m} s_{mo}(z; \mathscr{J}_m)\, dF_m(z)}. \tag{C.4}$$

The sum uses eligible 2013–2016 closure cohorts. The estimate $\hat{\tau} = 0.0852$ has standard error 0.0147,

decomposing into 0.0320 for departure from public schooling and 0.0532 for transfer to another municipality's public schools. The respective standard errors are 0.0089 and 0.0121; covariance matters when the two components are added. This is an excess-exit effect from the matched event study, not the observed postclosure exit rate.

The mapping $A(\hat{\theta}) = \hat{\tau}$ maintains the counterfactual equivalence explained in Section 5.3. If an entire inside menu closes, the model's conditional outside-diversion probability is one regardless of its parameters. Such markets contribute no identifying curvature for $\rho$ even when included in the target's population.

## C.3 Weighting and Inference

At each $\rho$, the criterion uses the sixteen rank moments and the composition moment. Rank discrepancies receive inverse municipality-bootstrap variance weights. The composition discrepancy receives 100 times its inverse-variance weight. The pooled mean destination distance and quality-related contrasts are reported as diagnostics rather than additional substantive identifying targets. In the implementation, diagnostic moments retain a numerical weight of $10^{-8}$, which has negligible influence on the criterion. The nesting parameter solves the separate exit equation after the inner parameters are re-estimated at each candidate $\rho$.

Let $g(\vartheta, \rho)$ collect the seventeen fitted discrepancies, $G_\vartheta = \partial g / \partial \vartheta'$, and define

$$h(\theta) = \begin{pmatrix} G_\vartheta' W g(\vartheta, \rho) \\ A(\vartheta, \rho) - \hat{\tau} \end{pmatrix}. \tag{C.5}$$

The sandwich calculation uses the full derivative $\partial h / \partial \theta'$, including derivatives multiplying nonzero moment discrepancies. This differs from retaining only the Gauss–Newton derivative. The distinction matters when the model does not fit the population moments exactly (Newey, 1984; Hall and Inoue, 2003).

The rank-target covariance comes from 999 municipality-cluster bootstrap draws. The composition variance and exit variance are added as separate blocks; the calculation assumes zero covariance between these blocks. In particular, the rank and event-study samples are related, so their zero covariance is a maintained inference approximation, not a consequence of independent data. Population integration weights, market-size construction, geographic measurements, distance curvature, and the moment weighting matrix are treated as fixed in the reported covariance.

The inversion confidence set for $\rho$ is $[0.7833, 0.9044]$. Profile inversion is also used for the rural distance shifter, whose estimate is close enough to the numerical search boundary that an unconstrained symmetric interval can be misleading. The search and boundary diagnostics were amended after some results were available: the revised checks use extended profile inversion and explicitly examine boundary solutions. These numerical checks do not supply an economic restriction on preferences. Reported policy ranges cover the evaluated parameter profiles; they are not simultaneous confidence intervals for every policy outcome.

## C.4 Fit and Sensitivity

The inner criterion is 26.90 for seventeen targeted moments and four parameters conditional on $\rho$. Its simulation-based overidentification p-value is 0.038. The rank-one shares, which are omitted from the target vector but implied by the other categories, illustrate the discrepancy:

| Group | Observed | Model |
|---|---|---|
| Urban, non-poor | 0.380 | 0.404 |
| Urban, poor | 0.387 | 0.422 |
| Rural, non-poor | 0.373 | 0.469 |
| Rural, poor | 0.446 | 0.498 |

These are specification diagnostics, not additional validation moments independent of the fitted rank shares.

The matched event-study exit effects are 0.120 for rural origins and 0.032 for urban origins, with standard errors 0.025 and 0.012. The model predicts 0.0915 and 0.0786. Thus, the pooled match conceals a substantial error in the geographic allocation of exit, particularly for urban origins. This is why the paper does not use predicted exit incidence to evaluate alternative closure lists.

The examined demand sensitivity specifications include earlier closure cohorts, alternative distance exponents, a distance-mean target, and alternative market-size constructions. The admitted specifications preserve the larger rural-poor access loss and do not restore the earlier claim that rural families have a higher effective cost per kilometer. Some broader-bandwidth and market-size specifications fail the numerical admissibility checks; their exclusion is a limitation of the completed sensitivity analysis, rather than evidence supporting the baseline.

## C.5 Quality Preferences

A common coefficient on teacher value-added cannot be distinguished from freely recovered school mean utilities without another restriction. A projection of those utilities on teacher quality would combine preferences, omitted amenities, and measurement error. Even a sign restriction on the covariance of latent quality and omitted amenities would not by itself resolve measurement error in the quality regressor.

In the school-year split analysis, the lower confidence limit for the within-market latent teacher-quality variance is nonpositive. Consequently, the examined ratio bound on the common quality coefficient is unbounded. The baseline access specification sets household-specific quality interactions to zero; it does not estimate them precisely at zero. Neither a weak estimated association nor these uninformative bounds establish that families place little value on effective teaching. Information counterfactuals must therefore be read as conditional response scenarios.

# D Counterfactual Definitions and Additional Results

## D.1 Populations, constraints, and comparisons

The three policy exercises answer different questions. Closure rules change the set of schools available in a given wave. Teacher reassignment changes the teachers assigned to fixed posts after the actual closures. The information exercise changes how the displaced population chooses among surviving schools. The exercises therefore need not have the same denominators. Access uses predicted attendance and population weights; observed destination comparisons use linked displacement records; teacher reassignment uses the eligible displaced-teacher pool and the students served by its receiving posts. A comparison of the magnitudes must retain these distinctions.

The closure-rule calculations impose the realized count within each municipality, year, and size band. A separate rule imposing counts only at the education-region level permits closures to move between municipalities and has a larger policy scope. Random closures under the corresponding count restrictions are reference distributions, not experiments showing that actual closures were randomly assigned. A sequential rule updates the surviving network after each closure but does not necessarily solve a global optimization problem.

The geographical comparisons use different distance constructions. Closure rules and their nearest-school comparisons use great-circle distance between schools multiplied by the median ratio of road to great-circle distance among observed routes. The ratios are approximately 1.52 for origins with rural exposure $A_j^R \geq 0.5$ and 1.42 for other origins. This common construction applies to every pair, so the greater availability of recorded routes for actual closures does not give those closures a different distance measure from candidate closures. The reach exercise instead uses recorded school-to-school road distances where available and fills missing pairs using the corresponding exposure-class ratio. The threshold used to select a distance ratio does not define the outcome groups, which retain fractional rural-exposure weights. Information scenarios measure changes in home-to-school distance using the demand model's block-group-to-school distances. Thus those distance changes differ from both the origin-school distances in the reach exercise and the kilometer equivalents used to value public-school access.

The enrollment-based feasibility screen compares displaced enrollment with the historical headroom of surviving schools. Historical headroom is the maximum recorded enrollment up to the closure year minus current enrollment. Where the actual allocation would violate the screen, it is relaxed to accommodate that allocation. This is a restriction on the analytical comparison, rather than a direct measure of building capacity, classroom availability, or feasible transport. It cannot be interpreted as an exact fiscal or physical-capacity constraint.

The initial calculations use enrollment both to form size bands and to aggregate distances and school contributions. They treated a school-level enrollment field repeated across grade rows as a school-grade population. The revised construction instead separates school totals used for size from student counts by grade used for the displaced population. Whether the direction and magnitude of the previous-run comparisons

survive this correction requires regeneration, since the correction can affect both the schools selected and the weights used to evaluate them. Table 7 in the main text reports those previous-run results as preliminary, and the abstract and introduction summarize them with the same qualification. Their rankings and trade-offs remain conditional on the earlier enrollment construction. This reservation also applies to the random-closure comparisons and to claims that none of the evaluated rules improves both outcomes.

The two rules ordering schools by average prior test scores were added after the initial comparisons had been examined. Their results should be interpreted as exploratory. For rules selecting on an estimated quantity, one estimation half determines the selected schools and the other evaluates their effects, and the roles are reversed. This reduces selection on estimation noise but does not make an exploratory rule prespecified. The conditional split-sample standard errors and the municipality bootstrap capture different sources of variation. For example, the rural-poor school-effect comparison for the score-based rule with the proximity screen is $-0.055$ with a split-sample standard error of 0.023, but its municipality-bootstrap interval is $[-0.116, 0.003]$. Thus even before the enrollment reconciliation, an unqualified claim of a statistically established improvement on that margin would be inappropriate.

## D.2 Information response and evaluation

Let $a^A_{jc}$ denote a posterior report card constructed from the school-year estimation half $A$, and let $\widehat{\Psi}^B_{jc}$ be the effectiveness estimate used for evaluation in the other half. The information scenario adds $\Delta b\, a^A_{jc}$ to the baseline within-nest utility of a surviving public school. For a calibration $w$, measured in kilometers per student-score standard deviation, the primary convention sets

$$\Delta b(w) = w\, p \left|\widetilde{\lambda}_0 + \widetilde{\lambda}_R + \widetilde{\lambda}_P\right| d_*^{p-1}, \qquad p = 0.8, \quad d_* = 3.9. \tag{D.1}$$

This uses the marginal distance cost of the rural-poor type at the reference distance. The same $\Delta b$ applies to all types in the reported exercise. A convention assigning every type the same willingness to travel would instead use its own distance coefficient in (D.1); it is a different experiment.

For each origin and poverty group, the model supplies the residential distribution conditional on having attended that origin. Let $P^{\text{info}}_{ij}(w)$ be the resulting probability of choosing survivor $j$, conditional on initial attendance at a closing school and with choice restricted to surviving public schools. $P^{\text{info}}_{ij}(0)$ is the no-additional-information probability for the same displaced type. The reported effectiveness change averages

$$\sum_{j \notin C_m} \left[P^{\text{info}}_{ij}(w) - P^{\text{info}}_{ij}(0)\right] \widehat{\Psi}^B_{jc} \tag{D.2}$$

over the displacement records' origin and poverty composition, with the corresponding calculation using $B$ for selection and $A$ for evaluation. This expression clarifies that information gains compare two modeled allocations. Observed destinations determine eligibility for the record sample, but are not substituted for $P^{\text{info}}_{ij}(0)$ in the calculation.

The full response grid is $w \in \{0, 0.25, 0.5, 1.1, 2, 4, 8, 13\}$. The main text reports 1.1 and 13 as calibrations, without treating either as an estimate of the Puerto Rican response to information. For the larger calibration, Table 2A of Harris and Larsen (2022) reports 2.024 miles per school-level standard deviation of value-added. Their note 24 gives an approximate school-level standard deviation of 0.25 in student-score units. The conversion is therefore $2.024 \times 1.609344/0.25 \simeq 13.0$ km per student-score standard deviation. Differences in information, transportation, choice sets, and populations limit its portability. The 1.1 calibration is retained as a moderate response scenario; its use does not require a claim that it is the median response in the literature.

The scenario assumes complete take-up, unlimited seats, fixed school attributes, and no re-sorting by other families. It adds a response to the card to baseline mean utilities, which can already reflect school quality and its correlates. Neither baseline knowledge nor the structural preference for quality is separately identified. Accordingly, the calculated gains are conditional on the specified incremental response and should not be described as an identified bound on the effect of information provision. The reported bootstrap intervals resample municipalities while holding the cards and demand parameters fixed. They do not incorporate all uncertainty in the estimated response of destinations to information.

For waves through 2016, the reported rural-poor gains at the same two calibrations are 0.060 [0.044, 0.079] and 0.337 [0.264, 0.426] standard deviations. The corresponding distance changes are 0.007 and 1.594 km. These results use 5,973 eligible displacement records from a frame of 7,050, rather than the main window's 12,564 from 15,025. They preserve the substantive sensitivity of the exercise to the response parameter.

### D.3 Alternatives within reach

The reach exercise does not specify families' preferences. For each observed displaced child, it searches over surviving public schools in the origin municipality that serve the relevant grade and lie no farther from the origin than the destination attended, or within one or two additional kilometers. It selects the most effective assignment using one estimation half and evaluates it using the other. These are distances from the closed school, not from the child's home. A more effective school in this set is therefore not proof that the family could have reached it at the same private travel cost.

The capacity-screened version allocates records jointly across schools. Available places are calculated from historical headroom, net of other displaced inflows, with a floor equal to the observed inflow of the sample to preserve the feasibility of the realized allocation. Hence it is a relaxed historical-enrollment screen, not an independently measured hard capacity. The same enrollment-unit issue that affects closure-rule comparisons affects these screened results.

In the preliminary 2013–2017 calculations, the rural-poor gain is 0.103 [0.065, 0.146] standard deviations at no extra origin-to-school distance and 0.131 [0.090, 0.173] with one additional kilometer. Allowing any surviving school in the municipality gives 0.426 [0.310, 0.535] with the screen and 0.469 [0.353, 0.580] without it. The screened figures await enrollment reconciliation and do not enter the paper's principal findings.

The unscreened figure is a descriptive assignment benchmark under unrestricted seats, rather than a prediction of take-up or a feasible placement policy.

# E Fiscal Accounting

## E.1 The Accounting Boundary

For a given closure set and assignment policy, write the annual change in the school system's budget as

$$S(C,\alpha) = F(C,\alpha) + W(C,\alpha) - K(C,\alpha) - \Delta T(C,\alpha) + E(C,\alpha). \tag{E.1}$$

Here $F$ is the reduction in facilities expenditure net of maintaining closed sites, $W$ is the reduction in expenditure on positions, $K$ is transition expenditure, and $\Delta T$ is the change in transportation expenditure. The term $E$ records expenditure changes associated with enrollment leaving the system. A positive value of $E$ is an accounting saving, not a measure of the social benefit of public-school exit. One-time transition expenses enter in the year incurred; asset sales and changes in property values belong in a separate capital account.

Reassignment yields payroll savings only if released posts remain unfilled or replace other funded positions. Similarly, average expenditure per pupil is not the cost avoided when a child transfers. Instructional resources can follow pupils, whereas facilities savings depend on receiving-school capacity. The accounting therefore requires assumptions about both reassignment and receiving costs.

## E.2 Conditional Component Calculations

The available fiscal calculations use expenditure records for surviving schools and personnel flows for the observed closure program. They are reported here to make the accounting assumptions explicit, and are not recomputed for the counterfactual closure lists. A regression of annual facilities expenditure on enrollment for 127 surviving schools in 2019–20 gives an intercept of \$18,657 and a slope of \$121.16 per pupil. The standard errors are \$5,200 and \$9.64, respectively. Predicting expenditure at the median closing school's enrollment of 139 gives \$35,498. Deducting the historical maintenance allowance of \$6,460 leaves \$29,038 per closed school-year.

This calculation applies a cross-sectional spending relationship outside its estimation sample and across years. Its interpretation as an avoided cost also requires assumptions about costs at receiving schools. In the archived accounting exercise, allowing for receiving-school utilization produces a conditional annual facilities range of \$15.2–17.7 million when extrapolated to 610 recorded closures. This is a historical accounting population, which includes later closures and differs from the 2013–17 counterfactual sample. The calculation should not be read as an estimate of savings realized by that earlier sample or as a policy-invariant cost function.

The corresponding personnel calculation uses the 2016–18 closure cohorts, for which personnel-flow

coverage is available. The revised count is 512 excess exits when negative school-level counts are set to zero, or 478 when they are retained, over 467 observed closing schools. Extrapolation and wage assumptions yield a $0–19.9 million annual positions component for the 610-school population. The lower endpoint allows released posts to be refilled; the upper endpoint assumes that the relevant released posts generate savings. Personnel observed in a later payroll file are a selected sample, so these calculations require an additional extrapolation assumption.

The combined $15.2–37.6 million annual range excludes transportation, transition, enrollment, and capital effects. It is not a confidence interval or bounds on net savings. Applying it to alternative policies requires consistent enrollment and capacity measures and a policy-specific cost mapping; it cannot supply a common budget or establish dominance across closure rules.