

A Framework for Assessing the Distributional Effects of School Closures:

Evidence from Puerto Rico

Reza Moradi*

October 14, 2026

Abstract

This paper studies the distributional consequences of school consolidation when families differ in access to effective alternatives. I develop and estimate a structural model of school choice and educational production using Puerto Rico's 2013–17 public-school closures. Linked student–teacher records and school enrollment choices allow me to separately estimate school and teacher contributions to achievement and families' responses to distance and the removal of public-school options. The results reveal substantial heterogeneity in the consequences of closure. The nearest surviving school stood about 3.4 km from rural children's closed schools, against about 2 km for urban children. Rural-poor children's change in school effects from the school closed to the school entered was about 0.08 test-score standard deviations less favorable than that of all other displaced children, mainly reflecting the schools closed. Their model-implied access loss per displaced child is 84 percent larger than urban non-poor children's, measured in km-equivalents; this conversion reflects rural families' lower estimated costs per km. At unchanged closure counts, proximity-based lists lower rural-poor distances to surviving schools, but their school-effect changes are not statistically distinguishable from the realized closures'. Keeping open 47 closures farthest from a surviving alternative avoids about 30 percent of rural-poor modeled access loss. The comparisons measure access and school contributions, not causal effects of closure on children's own achievement. These findings show why enrollment alone provides an incomplete basis for evaluating consolidation: distributional comparisons depend on both the effectiveness and accessibility of the alternatives available to displaced students.

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

Keywords: school closures; school consolidation; school effects; teachers' value-added; school choice; distance costs; rural schools; access to schools; 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

Demographic change poses fundamental challenges for the provision of public services. Declining birth rates and population movements alter not only the size but also the spatial distribution of demand for schools, hospitals, and other public infrastructure. Because these services rely on geographically fixed facilities and involve substantial operating costs, governments must adapt their existing networks to changing populations. School consolidation presents a particularly difficult trade-off. While declining enrollment creates immediate pressure to reduce expenditures, decisions about which schools to close can have consequences that extend far beyond the current generation of students. By changing children's access to education and the quality of schooling they receive, these decisions can affect lifetime opportunities and potentially transmit disadvantages across generations. What appears to be a short-run adjustment to demographic change may therefore have lasting consequences for economic inequality, particularly when disadvantaged families have limited access to alternative schools.

Existing work evaluates the consequences of school closures by studying how displaced students' enrollment and educational outcomes change. Such reduced-form evaluations provide essential evidence on the effects of consolidation policies and can inform future decisions about their design. However, their findings reflect the particular schools that closed, the alternatives available to displaced students, and the resulting allocation of educational resources. They cannot readily be extrapolated to alternative consolidation policies because changing which schools close also changes families' substitution opportunities, their enrollment decisions, and the allocation of teachers across the remaining schools. Evaluating these counterfactual policies therefore requires a framework that explicitly accounts for the interactions between school closures, household choices, and educational production. This paper develops such a framework to examine how alternative consolidation decisions shape the distribution of educational opportunities across students and communities.

I study this question in Puerto Rico, where a sharp decline in public school enrollment accompanied a fall of about 44 percent in the number of operating public schools between 2009–10 and 2019–20 (NCES, 2021b). This is a net change, not a count of closures; I study the 334 schools closed in 2013–17 and the 34,975 children in grades 1 to 7 whom they displaced. Puerto Rico combines the institutional framework of the US public education system with substantially lower incomes, widespread poverty, and considerable geographic disparities. This makes it a particularly informative setting for examining how the consequences of school consolidation differ across socioeconomic groups and locations, while offering lessons for other school systems facing demographic decline and fiscal pressures. The scale of the consolidation, together with detailed administrative records linking students, teachers, and schools, provides a rare opportunity to

examine how the consequences of closures depend on the organization of the remaining school network. I combine a model of household school choice, which characterizes substitution across schools in response to closures, with an additive test-score model that separates school-specific contributions to achievement from the contributions of teachers who can be reassigned. This distinction allows me to evaluate not only which educational opportunities disappear when schools close, but also how the remaining opportunities and resources are distributed across households.

To estimate the educational production component of the model, I exploit longitudinal administrative records linking students to their schools and subject-specific teachers. Under maintained assignment assumptions, movements of students and teachers across schools allow me to separately estimate school contributions to achievement and teacher value added, while accounting for persistent differences across students. This distinction is central to evaluating school closures because teachers can continue contributing to student learning after being reassigned, whereas school-specific contributions cannot simply be transferred with them. The resulting estimates allow me to characterize the educational inputs associated with individual schools and distinguish those that are potentially portable from those tied to the institution. The closure-list comparisons hold surviving-school inputs fixed; teacher reassignment is examined separately with the realized closures held fixed.

To estimate the school-choice component, I combine observed enrollment shares with information on where displaced students enroll after their schools close, alongside data on the geographic distribution and socioeconomic characteristics of households. Closing a school changes the set of alternatives available to families, but its consequences depend on more than the distance to the nearest remaining school. Families may differ in their willingness or ability to travel, their reliance on public education, and the attractiveness of the remaining schools. A separately estimated closure-induced change in enrollment outside the origin municipality's public schools provides additional information about substitution away from the local public-school network. Together, these sources allow me to characterize how families substitute among schools and measure the accessibility losses associated with removing different schools from the network. The model also supplies enrollment probabilities for children displaced under alternative closure lists. The records do not reveal who selected a displaced child's school, so the estimated cost of distance also reflects placement, busing, and information. The structural school-choice model keeps the attractiveness of schools to families separate from their contributions to student achievement, allowing me to evaluate accessibility and educational effectiveness as separate dimensions of school consolidation.

Differences in educational effectiveness across schools appear to reflect substantially greater variation in school-specific contributions than in the average value added of their teachers. Across the two school-year

splits, the estimated standard deviation of school effects is approximately 0.30 student test-score standard deviations, considerably larger than the dispersion of schools' average teacher value added, which ranges between 0.06 and 0.08. This finding has important implications for consolidation decisions. Moving students and teachers from one school to another does not necessarily preserve the educational inputs associated with the school that closes. Evaluating closure decisions therefore requires comparing the contributions of schools that disappear with those of the alternatives that remain available, rather than relying exclusively on enrollment or the characteristics of the teachers employed there.

The distributional consequences of school consolidation are equally striking. The nearest surviving school stood about 3.4 km from rural children's closed schools, against about 2 km for urban children. Rural-poor households experience model-implied losses of access to public schools that are 84 percent larger per displaced child than those of urban non-poor households. The raw distance and model comparisons use different populations and rural classifications. The priced gap depends largely on the conversion into distance equivalents: rural families have lower estimated costs per km, so a given utility loss translates into more km. In units of utility, the ordering of the two groups' losses is not determined over the confidence set for rural distance costs. The raw distance comparison separately shows the importance of the spatial distribution of alternative schools in determining who faces limited access after consolidation. Accessibility losses are expressed as equivalent increases in travel distance, rather than observed changes in commuting distances or monetary welfare losses.

Beyond differences in accessibility, school closures also change students' exposure to educational inputs. Comparing the schools attended by displaced students before and after closure, I find that changes in estimated school effectiveness are less favorable for rural-poor students than for urban non-poor students. The difference between these groups is approximately 0.09 test-score standard deviations when effectiveness includes both school-specific contributions and average teacher value added. The main comparison of school effects, excluding teachers, gives rural-poor children a change 0.084 test-score standard deviations (0.031) less favorable than that of all other displaced children. The larger estimated contrast lies in the schools that closed: relative to the municipality's average surviving school, the rural-poor contrast is 0.104 (0.031) for the school closed and 0.020 (0.016) for the school entered (Table 4). These estimates describe changes in the educational inputs associated with students' origin and destination schools, rather than causal effects of closure on their own achievement. Together, these findings reveal two distinct dimensions of inequality: rural-poor households face greater losses in access to public schooling and less favorable changes in the educational effectiveness of the schools their children attend.

Motivated by these findings, I examine whether alternative consolidation decisions could have reduced

the disproportionate burden borne by disadvantaged households. Specifically, I compare closure rules that vary which schools are selected while holding the number of closures fixed within municipalities, closure years, and enrollment-size categories. These comparisons give different levels of precision for geographic access and school contributions. Prioritizing schools with nearby surviving alternatives lowers the distance to the nearest surviving school for rural-poor children: from 3.43 km under the realized closures to 3.11 km under smallest first with an alternative within 4.8 km and 2.67 km under nearest alternative first. With children placed by the choice model, no alternative list's change in school effects for rural-poor children is statistically distinguishable from the realized closures', allowing for model error and adjustment for the five comparisons. Differences as large as 0.233 test-score standard deviations cannot be ruled out under the reported standard-error calculation (Table 7). The results separate what can be established about geographic access from what remains uncertain about school contributions, illustrating why neither school enrollment nor proximity alone measures both dimensions of consolidation. These comparisons hold closure counts rather than fiscal savings fixed and should not be interpreted as identifying an optimal consolidation policy.

Beyond deciding which schools close, the framework also permits an examination of how the remaining educational resources could be allocated. I first consider whether reallocating displaced teachers among compatible receiving-school positions could improve the educational inputs available to disadvantaged students. Within the observed pool of displaced teachers and compatible local posts, the point-estimated gains are small; the comparison holds the teacher pool and posts fixed. A separate assignment comparison asks, for each child with other children's schools held fixed, whether the child could have entered a more effective school without going farther from the closed school than the school entered, with seats capped at earlier maximum enrollment. I then examine how the allocation of displaced students might change if families placed additional weight on information about the effectiveness of surviving schools. These simulations suggest that changes in destination choices can generate improvements in educational exposure under an assumed response to information. The information scenario imposes no seat limit, and many simulated moves require enrollment above receiving schools' earlier maximum. Together, these exercises distinguish the consequences of selecting schools for closure from the scope for mitigating those consequences through subsequent reallocations of students and teachers.

This paper contributes to the literature on school closures and their consequences for students. Reduced-form studies, including Brummet (2014), Steinberg and MacDonald (2019), and Bobonis et al. (2026), examine how closures affect educational outcomes and enrollment decisions. These studies provide important evidence on the consequences of implemented consolidation policies. Brummet finds no lasting achievement harm for displaced students on average in Michigan, with gains for students leaving relatively

low-performing schools. Steinberg and MacDonald find no average change in displaced students' achievement in Philadelphia, but gains for those attending higher-performing schools. In joint work on Puerto Rico's consolidation episode, my coauthors and I report persistent reductions in public-system enrollment and heterogeneous achievement responses associated with the baseline performance of closed schools relative to nearby alternatives and the teacher inputs displaced pupils encounter. These findings make the schools students leave and enter central to interpreting realized closures; they do not directly identify what would happen under a different closure list. My analysis complements this literature by moving from evaluating realized closures to comparing how alternative configurations of the school network would change the distribution of educational opportunities.

The paper also contributes to the structural literature on school choice and the allocation of educational resources. [Epple et al. \(2018\)](#) provide a particularly relevant benchmark because they explicitly examine school closure decisions while allowing households to re-sort. They model a superintendent choosing closures to meet a capacity-reduction requirement, anticipating parents' choices among surviving schools and outside options. They compare alternative objectives subject to receiving-school capacity constraints. Their measure of school attractiveness reflects parental valuations rather than separate school and teacher contributions to achievement, and their conclusion proposes teacher reassignment as an extension. My approach combines a spatial model of household substitution with estimates of school-specific and portable teacher contributions, building on work such as [Mansfield \(2015\)](#). Mansfield uses teacher transfers to distinguish teacher contributions from other school-level factors; my contribution is to integrate that distinction into the closure problem, rather than propose a new fixed-effect estimator. Related research on the demand for effective schools, including [Walters \(2018\)](#) and [Abdulkadiroğlu et al. \(2020\)](#), distinguishes schools' attractiveness to families from their contributions to student achievement. Walters finds that students who would gain most from Boston charter schools are among those least likely to apply; in New York City, parents' preferences are unrelated to schools' estimated effectiveness once peers are held fixed. In my main specification, school attractiveness common to families is absorbed in school constants, and rural and poor families' additional responses to effectiveness are restricted to zero. Thus no school-effect estimate enters the access valuation; school effects instead evaluate the destinations that the choice model predicts, while the information exercise imposes an assumed response. Recent work also simulates closures of low-effect New York City high schools and multiple closure scenarios for one district ([Abdulkadiroğlu et al., 2026](#); [Chrzan and Pearman, 2026](#)). Together, these elements allow me to evaluate alternative consolidation policies along two distinct dimensions: changes in access to public schooling and changes in the educational inputs available to displaced students. In Puerto Rico, proximity-based lists lower rural-poor distances, but the school-effect

comparisons are too imprecise to determine whether these lists also yield more favorable changes in school effects.

The remainder of the paper proceeds as follows. Sections 2 and 3 describe the institutional setting and data. Section 4 presents descriptive evidence. Sections 5 and 6 introduce the model and estimation strategy. Section 7 presents the empirical findings, Section 8 evaluates alternative consolidation and allocation policies, and Section 9 concludes.

2 Institutional Background

This section describes the public school system, the closures of 2013–17, and what followed a closure for children and for teachers. One Department decided every closure, and teachers were placed by a process separate from the one that placed children. The analysis measures access within the closed school’s municipality; this is a maintained set of alternatives, not a restriction on every child’s observed enrollment.

2.1 Public Schools and the Closures of 2013–17


The Puerto Rico Department of Education (the Department) runs every public school on the island. It is both the state education agency and the only local education agency (PRDE, 2017a). The 78 municipalities are geographic units without school budgets of their own, so every closure studied here was a decision of one Department.

The system contracted sharply over the decade. Public enrollment fell by about a third between fall 2010 and fall 2018, and the number of operating public schools by about 44 percent between 2009–10 and 2019–20 (NCES, 2021a,b). Both are net changes over all grades and a longer period than this paper covers; neither is a count of closures. Appendix Figure A1 shows schools, enrollment and closures by year in the Department’s own records for grades 1 to 8.

I date a closure by the school’s last operating year, labeled by the calendar year in which that school year ends. A closure of 2017 is a school whose last year was 2016–17 and whose children enrolled elsewhere in 2017–18. Dated this way, 334 schools closed in 2013–17: 8, 70, 59, 41 and 156 in the five years (Section 3.1 gives the rule behind the count). The closures of 2017 alone displaced about half of the children studied in this paper, 17,907 of 34,975 (Section 3.3). Figure 1 maps the schools that closed and those still open in 2017.

All 334 closures were decided and carried out before Hurricane María made landfall on September 20, 2017 (National Hurricane Center, 2017). Which schools closed, what they looked like the year before, and

Figure 1. Schools closed in 2013–17 and schools open



Notes: Each mark is a public school that serves a grade with a standardized test. Filled: the 334 schools closed in 2013–17; a closure is dated by the school’s last operating year (8, 70, 59, 41, 156 in 2013 to 2017). Hollow: the 935 schools open in 2017. Coordinates are those of the school file, rounded to 0.001 degree; municipal outlines are from the Census Bureau’s cartographic boundary file. Table 1 holds the counts by year; its 1,274 open schools are those open in any year of 2013–17.

how far each stood from the nearest surviving school were therefore fixed before the storm. The school that a child displaced in 2017 entered was not. Those children attended their new schools in 2017–18, a year in which the storm damaged schools and interrupted instruction (Nelson et al., 2020). I do not study the closures that came after the storm.

Written criteria for choosing which schools to close appear in the record only from 2015. The two May 21, 2015 letters read list seven criteria, among them current and projected enrollment, infrastructure, operating costs and the school’s location, and name receiving schools.¹ A circular of November 2015 listed criteria for closing, consolidating or restructuring schools in order of priority. It allowed exceptions for difficult transport, academic programs and special education, and it required an assessment of the receiving school’s space, staff and transport (PRDE, 2015). It postdates the 137 closures of 2013–15. A circular of March 2017 replaced it, listing transport, location, building use and academic offerings among its considerations and requiring consultation before the Secretary decided (PRDE, 2017c). None of these documents gives the weight of any criterion, and I have found no document that set criteria for the closures of 2013 and 2014.

In the year before each closure, the Department held each school’s enrollment, grades served, location, average test scores and test-score growth. These are the only characteristics that the other sets of closures

¹ Department of Education letters to closing schools, May 21, 2015 (intraedu.dde.pr, “Cartas/Escuela-cierre”), for closures effective in school year 2015–16. The description rests on the two letters read.

in Section 8 use. Closure was not random with respect to these characteristics: Section 4.1 shows that it followed enrollment most closely.

2.2 After a Closure: Placement, Transport, Teachers

A child whose school closed could enroll in another public school, a private school, or no school. For the analysis, I define the set of alternatives as the surviving public schools of the closed school's municipality that serve the child's grade. Both distance measures use this maintained set, and the choice model of Section 5 groups enrollment outside it into a single outside option. Some recorded destinations fall outside this set (Section 4.4), so the access measure describes within-set access rather than every school a child could attend.

For the closures of 2017, the Department's online enrollment instructions of August 2017 let families request a school and a grade; school directors approved enrollments and could transfer a child to a school other than the one requested (PRDE, 2017d). The Department also named receiving schools for that year's closures, as its letters of May 2015 did for the closures of 2015 (Section 2.1); I have found no such document for the closures of 2013, 2014 and 2016. The records show where a displaced child enrolled, but not whether the family or the Department chose that school; Section 6.2 takes up what this means for the choice model.

In 2017 families requested transport on the same enrollment form, and the director approved or rejected the request once the enrollment was approved (PRDE, 2017d). What transport displaced children actually received, in 2017 or earlier, is not in the records or in any document I have found. I use 4.8 km, or three miles, as a reporting benchmark for the long-distance tail. It is not an established transport entitlement or policy threshold for Puerto Rico.

Teachers were placed by a separate process. Under the Department's personnel circular of January 2017, the teachers of closing schools were reassigned according to permanent status, experience and qualifications (PRDE, 2017b); I have not found the rules for earlier years. Nothing in either process required a closed school's children and its teachers to go to the same school. A comparison of the school a child left with the school the child entered therefore mixes what each school contributes with what its teachers carry from one school to another. Section 5 separates the two, after Sections 3 and 4 describe the records and what they show.

3 Data

The paper combines the Department’s administrative records with the Puerto Rico Community Survey. I describe each source in turn and then build the objects that combine them; Appendix Table A1 lists every variable with its source and construction.

3.1 Department of Education Records

The Department provided four extracts of its administrative records under a data use agreement. I link school-level information by school code and school year, and link the pupil records to teacher links by child identifier and school year.

The first extract, the enrollment extract, counts the children enrolled in each grade of each public school that serves grades 1 to 8, from 2012–13 to 2018–19. It gives the children in each grade of a closed school in its last year, each school’s size, and the earlier maximum enrollment that Section 8 uses as a count of seats. For 2013–17 it holds 5,949 school-years: the 334 closed schools in their last year, and 5,615 school-years, from 1,274 schools, in which a school stayed open.

The second extract, the pupil records, follows children enrolled in public schools from year to year. It records each child’s school, grade of enrollment, municipality and ZIP code of residence, the Department’s family-poverty record, and test scores in math, Spanish and English for grades 3 to 8. The family-poverty record is the Department’s flag for a low-income family. It is recorded only in years with a test score, and I fill the other years with the child’s most frequent status (Appendix A.5). Because the records follow children, they show the public school a displaced child entered the next year, or that the child left the municipality’s public schools.

The third extract, the teacher links, connects each child, in each year, to the child’s teacher in math, in Spanish and in English: 3,082,690 child-years from 2009–10 to 2022–23, of which I use those through 2018–19. One identifier follows a teacher across subjects and schools. The links therefore show where the 2,145 teachers of the schools closed in 2013–17 taught afterward, and they let Section 5 separate what moves with teachers from what stays with a school.

The fourth extract, the school file, describes each school: its coordinates in each year, the grades it serves,

the Department’s rural or urban designation, and its last operating year, by which I date every closure. Under this rule, 334 schools serving at least one grade with a standardized test closed in 2013–17.²

The records lack three things. First, homes are observed only to ZIP code. Second, no operating costs by school exist in the records or in any public file for these years, so Section 8 compares sets of closures by their number, except where Section 8.3 prices the closures spared (Appendix F). Third, there is no record of transport (Section 2.2).

3.2 Puerto Rico Community Survey

The Puerto Rico Community Survey is the Census Bureau’s American Community Survey for the island, and I use two of its public products. The first gives the school-age population of each block group. I call a block group rural when its population density falls below a threshold set to match the Department’s rural designation of schools, because the Census urban share equals 1.0 in the median block group and so cannot separate rural from urban areas on the island (Appendix A.5). The block groups locate the rural population near each school, which defines the four groups of Section 3.3. They are also the residences of the families in the choice model, linked to schools by routes on the OpenStreetMap road network.

The second product is the Census Bureau’s 2013–17 five-year public-use person file for Puerto Rico, obtained from its public PUMS archive. It records children’s enrollment (public school, private school, or none) and their family’s income against the Census poverty line. Among children aged 8 to 14 in public schools, 67.1 percent (0.4) are income-poor. Income-poor children are less often outside public schools: 8.7 percent were in a private school or not enrolled, against 42.6 percent of other children. Income poverty is not the Department’s family-poverty record, which marks about four in five children in the schools studied here as poor. Section 5.2 uses the two measures for different jobs. The survey is not linked to the pupil records child by child.

3.3 Displaced Children, the Four Groups, and the Two Measures

Displaced children are the children in grades 1 to 7 enrolled in a closed school in its last operating year. The enrollment extract gives 36,109 of them; the 34,975 in schools with a rural and a poor share, in 323 of the 334 closed schools, are the displaced children of this paper: 495, 6,877, 6,073, 3,623 and 17,907 for the

² The 334 (8, 70, 59, 41 and 156 by year) leave out five schools with a last operating year in the window that served no grade with a test and have no rows in the enrollment extract, for 339 in all. The 34,975 displaced children with group shares were enrolled in 323 of the 334. Section 8 chooses 315 of the 334 again under other sets of closures and holds the other 19 closed in every set; these 323 schools are 314 of the 315 and 9 of the 19.

closures of 2013 to 2017 (Table 1, Panel A; Appendix A.2). Because the counts come from enrollment, they include children who later left the public schools.

Table 1. Closures, displaced children, and schools the year before

<i>Panel A. Closures and displaced children</i>						
Last operating year	Schools closed (1)	Children displaced (2)	Percent of displaced children			
			Rural poor (3)	Rural non-poor (4)	Urban poor (5)	Urban non-poor (6)
2013	8	495	39.1	16.4	27.0	17.5
2014	70	6,877	40.8	9.8	42.8	6.7
2015	59	6,073	37.4	8.0	47.0	7.6
2016	41	3,623	47.1	9.0	39.5	4.5
2017	156	17,907	38.7	8.3	44.8	8.2
2013–17	334	34,975	39.7	8.7	44.0	7.6
<i>Panel B. Public schools in the year before</i>						
	Closed 2013–17 (7)	Open (8)	Closed minus open (9)			
Enrollment	165.2	290.0	–124.8 (5.6)			
Grades offered	5.19	5.36	–0.17 (0.10)			
Rural share of the catchment	0.52	0.43	0.09 (0.02)			
Poor share of pupils	0.84	0.82	0.02 (0.01)			
Km to the nearest other school	2.67	2.61	0.06 (0.10)			
Average test score, SD	–0.007	0.037	–0.044 (0.031)			
School-years	334	5,615	5,949			
School-years with an average test score	326	4,297	4,623			
Schools	334	1,274	1,281			

Notes: A closure is dated by its last operating year; displaced children: those in grades 1–7 of a closed school in that year; the 34,975 displaced children with group shares were enrolled in 323 of the 334 closed schools. (3)–(6): a school’s children split by its catchment’s rural share and pupils’ poor share. Panel B: school means in the year before; (7) holds a closed school once and (8) a school in every year of 2013–17 in which it did not close, so 327 of (8)’s 1,274 schools closed later and are also in (7); (9) is (7) minus (8), standard error clustered by municipality; counts cover each school once. For the closures of 2013, characteristics are those of the last operating year, since no file for 2011–12 exists, and there is no average test score. In (7), the rural share of the catchment exists for 324 schools. Km: nearest other school of the municipality serving a common grade. Average test score: mean of the school’s pupils.

The pupil records give a second flow: 32,514 displaced children, from kindergarten to grade 8, with a recorded next public school. They are not a subset of the 34,975, since the grades and the extracts differ (Appendix Table A2), but they give the schools that children entered. Children absent from the public-

school records have no observed school entered, so destination comparisons exclude them. The records with a school entered are therefore a selected group.

I divide displaced children into four groups: rural poor, rural non-poor, urban poor and urban non-poor. Rural is the share of the population near the closed school that lives in rural block groups, weighting each block group by its population and by $\exp(-d/2)$, where d is its distance in km from the school. Poor is the family-poverty record. In the enrollment counts, a school's children are split by the school's rural share and by its share of children with the poverty record, so group counts are not whole numbers and the poor and non-poor children of one school share its distances. The four groups are 39.7, 8.7, 44.0 and 7.6 percent of the 34,975, so rural-poor children number about 13,900. The pupil records use the same rural share and the child's own poverty record, so the groups keep their names in both flows. Appendix Table A3 compares the three constructions of rural used in the paper.

The first measure is access: the distance in km from the closed school to the nearest school in its set of alternatives. I compute it as the straight-line distance times the median ratio of road to straight-line distance for the closed school's class, 1.52 for rural and 1.42 for urban schools, from 279,759 routes on the road network. For this adjustment, a school is rural when at least half of its nearby population lives in rural block groups. I also report the share of children with no surviving school within 4.8 km. Both statistics start at the closed school, not at a child's home, so neither is the change in a child's trip to school. The pupil records carry their own distance between the school closed and the school entered, whose mean absolute difference from this construction is 0.47 km (Appendix A.4), so I compare distances only within one construction.

The second measure is a school's average test score: the mean score of its children with a test score in grades 3 to 8, in the year before the closure. Scores in math, Spanish and English are normalized within subject, grade and year to mean zero and standard deviation one, and SD denotes a standard deviation of these scores. Standardization removes subject-, grade- and year-specific differences in score means and scales; it does not by itself establish comparability across the transition from the Pruebas Puertorriqueñas de Aprovechamiento Académico (PPAA) to the META-PR test, first given in 2015–16 (PRDE, 2016, p. 4). I measure the average for the same schools as distance: the school closed, the nearest surviving school, the municipality's surviving schools weighted by enrollment, and the school entered. Its standard deviation across schools is 0.52 SD; the schools closed in 2013 have no average test score. An average reflects which children attend a school as well as what the school and its teachers add; Section 5 separates the two.

I study the closures of 2013–17 as one window: every result pools the five years, and tables by year appear only in the appendix. No displaced child's own test score is an outcome in this paper; what I compare are the schools that children left and entered. Section 4 describes both measures for the four groups.

4 Descriptive Evidence


This section asks which schools closed and what each group of displaced children lost, using only counts, shares, km and average test scores.

4.1 Which Schools Closed

Closed schools were small. In the year before closing, a closed school enrolled 165.2 children on average, against 290.0 in the school-years of schools that stayed open; the difference is -124.8 (5.6) (Table 1, Panel B).

Closure followed enrollment. In a conditional logit of closure within municipality and year, which also includes the school's distance to its nearest neighbor, its rural and poor shares and its enrollment trend, the coefficient on log enrollment is -3.02 (0.22) (Appendix Table B1). Figure 2, Panel A shows the pattern without a model: in the smallest of 20 equal bins of school-years, ordered by enrollment, 25.3 percent closed, against 5.3 percent of all school-years.

Figure 2. Enrollment, closure, and the two measures



Notes: School-years of public schools, 2013–17, by enrollment the year before (5,930 in A, 315 of them of schools closed that year; 5,861 in B; 4,605 with an average test score in C). A point is the mean over a bin of school-years. A, B: the same 20 bins of 280 to 313 school-years. B: km (a straight line times a road ratio) to the nearest other school of the municipality serving each of its grades, weighted by the school's enrollment in each grade; rural: a mostly rural catchment; a line: the area's slope on log enrollment with closure-year effects, through its mean. C: schools closed that year (filled) and not closed (hollow), by fifth of enrollment, a fifth with fewer than 10 closed schools pooled with the fifth below; scores in SD. An average reflects a school's children as well as the school. The 19 closures held in every list of Section 8 are not among these school-years (5,930 of the 5,949 of Table 1). Appendix Figure B2 gives the slopes.


Other differences between closed and open schools were small. The share of the nearby population living in rural block groups was 0.52 around closed schools against 0.43 around open ones, a difference of 0.09 (0.02). The share of children with the family-poverty record was 0.84 against 0.82, a difference

of 0.02 (0.01). The difference in average test score is -0.044 SD (0.031), and the nearest other school of the municipality stood 0.06 km (0.10) farther from closed schools; neither difference can be distinguished from zero. Enrollment carried the clearest pattern. Among schools of similar size, however, the schools that closed had lower average test scores than those that stayed open (Figure 2, Panel C). What each group of children lost depended on where the closed schools stood: among schools with higher or lower average test scores, and on the map.

4.2 Average Test Scores of the Schools Closed and Entered

Average test scores show a gap between rural-poor children and all other displaced children, mainly reflecting the schools that closed. I compare the average test score of the school a child left with that of three surviving schools: the nearest, the average of the municipality's surviving schools weighted by enrollment, and the school the child entered. Every comparison is the surviving school minus the school closed, so a negative number means a lower average score than the school closed. The rural-poor contrast is the value for rural-poor children minus the value for all other displaced children. Figure 3, Panel A draws, for each group, the average test score of the school closed and of the school entered.

Figure 3. The two margins in raw data, by group: average test scores of the schools closed and entered, and distance to the nearest surviving school



Notes: Closures of 2013–17; groups by the child's poverty record and the closed school's rural share. A: the average test score of the school closed and of the school entered, each the year before the closure, in SD; the 14,791 pupil records of re-enrolled children with both scores. Averages reflect a school's children as well as the school. B: km from the closed school (not the home) to the nearest surviving school of the municipality serving the grade (a straight line times a road ratio), and the percent of children with none within 4.8 km (3 miles); all 34,975 displaced children. Lines: 95 percent intervals, standard errors clustered on the closed school's municipality. Rural poor minus all other displaced children: closed 0.155 (0.056), entered 0.058 (0.045), entered minus closed -0.097 (0.042); km 1.29 (0.20). Table 2 holds the numbers.

The schools that rural children lost had average scores above the island's mean, and at the point estimate those that urban-poor children lost were below it (Table 2, Panel A). Among the 34,191 displaced children with an average test score for the school closed and the nearest surviving school, rural-poor children's closed schools scored 0.153 SD (0.061) above those of all other displaced children. Before any child moved, the municipality's surviving schools scored below the school closed for rural children and above it for urban-poor children at the point estimates, a rural-poor contrast of -0.105 (0.036). Against the nearest surviving school alone, the rural-poor contrast is -0.072 (0.053), which cannot be distinguished from zero.

At the point estimates, urban-poor children entered schools with higher average scores than the schools they left, while both rural groups entered schools with slightly lower average scores. Neither rural group's change is distinguishable from zero. On the 14,791 records of re-enrolled children with an average test score for both schools, the school entered minus the school closed is -0.036 SD (0.046) for rural-poor children and -0.047 (0.052) for rural non-poor children. For urban-poor and urban non-poor children it is 0.084 (0.037) and 0.045 (0.050). The rural-poor contrast is -0.097 SD (0.042): rural-poor children's change was less favorable than that of all other displaced children. The rural minus urban difference is -0.117 (0.049). Against urban non-poor children alone, who are 7.6 percent of the displaced, the difference of -0.081 (0.058) cannot be distinguished from zero. Rural-poor children's own raw change cannot be distinguished from zero; the raw result is a gap between groups.

The gap sits in the schools that closed more than in the schools entered. On the same 14,791 records, the schools that rural-poor children left scored 0.155 SD (0.056) above those of all other displaced children, and the schools they entered 0.058 (0.045) above. Average test scores also fall with school size: among open schools, the smallest fifth by enrollment averaged 0.193 SD and the largest fifth -0.040 (Figure 2, Panel C). A rule that closes the smallest schools therefore closes schools with above-average scores. Both facts are associations.

4.3 Access: Who Lost It and Why

Access was a rural loss. The nearest surviving school stood 3.43 km (0.16) from the closed school for rural-poor children and 3.37 km (0.22) for rural non-poor children, against 1.91 km (0.19) and 2.12 km (0.18) for urban-poor and urban non-poor children (Figure 3, Panel B; Table 2, Panel B). Rural children's nearest surviving school was 1.48 km (0.22) farther than urban children's.

The far tail is rural as well. For 19.6 percent of rural-poor and 20.4 percent of rural non-poor children, no surviving school lay within 4.8 km, against 5.1 and 6.1 percent of urban-poor and urban non-poor children; the share for rural children is 14.5 percentage points (3.2) higher. Within rural areas, the difference between

Table 2. Average test scores and distance: schools closed, nearest surviving, and entered

	Displaced children					Difference		
	All (1)	Rural poor (2)	Rural non-poor (3)	Urban poor (4)	Urban non-poor (5)	Rural minus urban (6)	Rural poor minus all other (7)	Rural poor minus urban non-poor (8)
<i>Panel A. Average test score, SD of pupil scores</i>								
Displaced children								
School closed	0.002 (0.052)	0.094 (0.039)	0.152 (0.060)	−0.117 (0.067)	0.043 (0.093)	0.199 (0.066)	0.153 (0.061)	0.051 (0.091)
Nearest surviving school minus school closed	0.022 (0.043)	−0.021 (0.051)	−0.035 (0.069)	0.081 (0.048)	−0.031 (0.092)	−0.089 (0.059)	−0.072 (0.053)	0.010 (0.096)
Municipality's surviving schools minus school closed	0.019 (0.033)	−0.044 (0.035)	−0.086 (0.053)	0.102 (0.037)	−0.013 (0.054)	−0.137 (0.039)	−0.105 (0.036)	−0.031 (0.059)
Displaced, with scores	34,191	13,600	2,954	15,098	2,539	34,191	34,191	16,139
Children who re-enrolled in a school of their set of alternatives								
School closed	−0.013 (0.048)	0.077 (0.043)	0.087 (0.048)	−0.117 (0.062)	−0.031 (0.059)	0.184 (0.063)	0.155 (0.056)	0.108 (0.060)
School entered	0.007 (0.037)	0.041 (0.028)	0.040 (0.032)	−0.033 (0.057)	0.014 (0.050)	0.067 (0.051)	0.058 (0.045)	0.027 (0.045)
School entered minus school closed	0.020 (0.036)	−0.036 (0.046)	−0.047 (0.052)	0.084 (0.037)	0.045 (0.050)	−0.117 (0.049)	−0.097 (0.042)	−0.081 (0.058)
Re-enrolled, with scores	14,791	6,217	1,210	6,353	1,011	14,791	14,791	7,228
<i>Panel B. Distance from the closed school, km</i>								
Displaced children								
To the nearest surviving school	2.66 (0.18)	3.43 (0.16)	3.37 (0.22)	1.91 (0.19)	2.12 (0.18)	1.48 (0.22)	1.29 (0.20)	1.31 (0.22)
Percent with none within 4.8 km	12.3 (2.3)	19.6 (3.0)	20.4 (4.8)	5.1 (1.6)	6.1 (1.9)	14.5 (3.2)	12.2 (2.5)	13.5 (3.1)
Displaced children	34,975	13,900	3,051	15,379	2,646	34,975	34,975	16,545
Children who re-enrolled in a school of their set of alternatives								
To the school entered	3.88 (0.17)	4.70 (0.16)	4.74 (0.21)	2.98 (0.14)	3.35 (0.16)	1.67 (0.20)	1.43 (0.20)	1.35 (0.20)
School entered minus nearest surviving school	1.33 (0.06)	1.37 (0.08)	1.61 (0.11)	1.23 (0.06)	1.44 (0.09)	0.15 (0.08)	0.06 (0.07)	−0.07 (0.10)
Percent entering the nearest surviving school	41.0 (2.2)	44.6 (2.3)	37.2 (3.6)	38.7 (3.0)	37.9 (3.2)	4.8 (3.3)	6.1 (2.8)	6.6 (3.6)
Re-enrolled children	15,025	6,345	1,221	6,436	1,023	15,025	15,025	7,368

Notes: Closures of 2013–17. Panel A: a school's average test score the year before the closure, in SD of pupil scores (none for the closures of 2013); the municipality row is the enrollment-weighted mean of the municipality's surviving schools. Re-enrolled: pupil records of displaced children found the next year in a surviving school of their set of alternatives (Section 2.2), grouped by the child's own poverty record. Panel B: distance as in Figure 3; the re-enrolled rows use the pupil records' km between schools, another construction. Groups as in Table 1; N are rounded sums of shares, which need not add up, and a difference column counts both sides. Standard errors in parentheses, clustered on the closed school's municipality; those of (1), (6), (7), which pool groups, let the mix of groups vary. Displaced children: 34,975 in 323 closed schools.

poor and non-poor children in distance to a surviving school is small and cannot be distinguished from zero: the difference is 0.06 km (0.12), a comparison that reflects only how poverty shares differ across schools, since a school's poor and non-poor children share its distance (Section 3.3). The access gap runs between rural and urban children, not between poor and non-poor children. Appendix Figure B1 draws the full distributions behind these means.

The rural distance gap is consistent with rural schools lying farther apart. In the closure regressions, there is no detectable additional association between closure and isolation after controlling for enrollment. A rural school's nearest other school of the municipality serving its grades was 3.89 km away on average, against 1.93 km for an urban school. The rural mean is the larger in each of the 20 enrollment bins of Figure 2, Panel B. Small schools are also more often rural: about two thirds of the school-years in the smallest bin are rural, against about three in ten in the largest.

Across the island, small schools are more isolated; within a municipality, the relation cannot be distinguished from zero. Km to the nearest other school falls with log enrollment with a slope of -0.306 (0.117) given year effects, and of -0.059 (0.124) within a municipality and year (Appendix Figure B2 gives the slopes by area). No association between distance and closure was detected after controlling for enrollment: in the closure logit of Section 4.1, the coefficient on log km to the nearest other school is -0.09 (0.11). That logit measures a school's distance to its nearest neighbor before any closure, one school at a time; the random lists of Section 8.2 measure the distance to the nearest surviving school after a year's closures, which depends on which neighbors close together.

These patterns link the rural access gap to the geography of small schools; they do not identify the effect of school size, which no design in this paper varies. Because the measure follows the set of alternatives, a nearer school across a municipal line does not count; a version that admits any municipality is not reported.

4.4 Where Displaced Children Went

Most displaced children who re-enrolled did not enter the nearest surviving school, and the school they entered was farther by a similar distance in every group. Among 15,025 displaced children in the pupil records who re-enrolled in a surviving school of their set of alternatives, 41.0 percent (2.2) entered the nearest one (Table 2, Panel B). The shares were 44.6, 37.2, 38.7 and 37.9 percent for rural-poor, rural non-poor, urban-poor and urban non-poor children. Of all re-enrolled displaced children in the choice model's extract of the pupil records, 19.5 percent entered, outside the set of alternatives, the school most common among the closed school's other grades, and 11.9 percent another school outside it; neither access measure counts these schools (Appendix A.2). On the same 15,025 records, the school entered was 1.33 km (0.06)

farther from the closed school than the nearest surviving school, and this extra distance was only 0.15 km (0.08) larger for rural than for urban children. Where children enrolled thus added little to the rural gap. These distances use the pupil records' own km between schools, so I compare them only with each other.

Some children left the municipality's public schools, and children of rural schools left more often. An event study with the estimator of Sun and Abraham (2021) compares children displaced by the closures of 2013–16, in grades 3 to 7, with matched children whose schools stayed open. In the first year after a closure, the probability of leaving the municipality's public schools rose by 8.5 percentage points (1.5) (Appendix A.6). Of this rise, 3.2 points (0.9) are children no longer in the public records and 5.3 points (1.2) are children enrolled in another municipality's public school. The rise was 12.0 points (2.5) for children of schools the Department designates as rural and 3.2 points (1.2) for children of urban schools. Children absent from the public records may have entered a private school, left the island or stopped attending; the records do not say which.

Each of these facts asks something of the models that follow. Average test scores show that rural-poor children's change was less favorable than that of other displaced children, but not why: an average mixes the children a school enrolls with what the school and its teachers add. Smaller schools also have higher averages, and teachers were reassigned apart from children (Section 2.2). Separating each test score into a part that belongs to the child, a part that moves with the teacher and a part that stays with the school requires the score model of Section 5.1. Km to the nearest surviving school, in turn, does not say what the lost option was worth to a family: most children did not enter the nearest school, and some left the municipality's public schools. Valuing the lost option requires a choice model in which the cost per km can differ by type of family and an outside option matches the estimated exit (Section 5.2).

5 A Model of Test Scores and School Choice

This section defines the two margins for any set of closures, realized or not. The score model of Section 5.1 separates a school's contribution to test scores from those of its students and teachers, and so defines the school effects that the second margin compares. The choice model of Section 5.2 describes where displaced families enroll; it prices the option a closure removes and places the children of closures that did not happen. Section 5.3 defines both margins with the two models.

5.1 Student, Teacher and School Effects

Let i index students, s subjects (mathematics, Spanish and English), t school years, k teachers and j schools. Write $k(i, s, t)$ for the student's teacher in subject s and year t , and $j(i, t)$ for the school the student attends. The score model is

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

where Y_{ist} is the test score normalized within subject, grade and year. The student effect a_i is common to the three subjects, μ_{ks} is teacher k 's value-added in subject s , and ϕ_j is the school effect. A subject intercept would be absorbed by the teachers' value-added, so the equation has none. The error is the sum of a student-subject term, a classroom-year shock and a school-year shock, $e_{ist} = v_{ist} + \zeta_{c(i,s,t)} + \eta_{j(i,t),t}$. Scores are therefore correlated within a classroom and within a school-year. All effects are in SD. School effects are common to all students; the paper does not estimate effects that differ by type of student.

The school effect is the part of a school's contribution to test scores that does not move with its teachers. When a school closes, its teachers are reassigned to other schools and take their value-added with them (Section 2.2). A displaced child leaves the closed school's effect behind and receives the effect of the school entered. For school j in year t , let $n_{jt,ks}$ count the scored student-subject observations taught by teacher k in subject s , and let $n_{jt} = \sum_{k,s} n_{jt,ks}$. The school-average value-added and the effectiveness of the school are

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

Effectiveness is what an attending child receives from a school and its teachers together. A teacher's value-added is portable: it does not depend on the school where the teacher works. Section 7.1 bounds a component specific to a teacher and a school.

The school effect is constant over the years of the sample. It bundles everything that stays with a school when its teachers change, including its organization, its class sizes and its peers, and equation (1) does not separate these parts. Under every set of closures, I hold the school effects of the surviving schools fixed, although their intake changes. This is a maintained restriction; the paper does not estimate how a school responds to the children it receives.

Teachers enter the score model so that the school effect is net of them. I do not model how teachers would be reassigned under other sets of closures; Appendix E.4 describes a reassignment of displaced teachers with the realized closures held fixed. School effects and value-added describe performance in the three subjects

tested, and effects on test scores need not match effects on other outcomes, for teachers (Jackson, 2018) or for schools (Beuermann et al., 2023).

5.2 School Choice after a Closure

The choice model gives the probability that a displaced family enters each surviving school or leaves the municipality's public schools. It has two uses. First, it prices the option a closure removes, the value to a family of the school that closed, in km. Second, Section 8 uses its probabilities to place the children of closures that did not happen, after a check on the realized closures (Section 6.2).

A market m is a municipality, a grade and a year. Its choice set $\mathcal{J}_m$ contains the public schools of the municipality that serve the grade, the set of alternatives of Section 2.2. The outside option, $j = 0$, pools private school, not enrolling, and a public school in another municipality. The records observe homes only to ZIP code, so families are placed at census block groups, each carrying the survey's population weights (Section 3.2). For a family i in market m , the utility of school $j \in \mathcal{J}_m$ is

$$u_{imj} = \delta_{mj} + (1 - \rho) \tilde{\lambda}_i d_{ij}^p + \pi I_i + \varepsilon_{imj}, \quad \tilde{\lambda}_i = \tilde{\lambda}_0 + \tilde{\lambda}_R R_i + \tilde{\lambda}_P P_i, \quad (2)$$

and the utility of the outside option is $u_{im0} = \varepsilon_{im0}$. Here d_{ij} is the routed distance in km from the block group's centroid to school j , and the power $p = 0.8$ is fixed before estimation. R_i indicates a rural block group, P_i is the Department's family-poverty record, and I_i is income poverty as the survey measures it. The constant δ_{mj} is common to all families in the market and is recovered from enrollment shares. It holds everything about a school that all families value alike, effectiveness included.

The errors follow a nested logit with one nest for the municipality's public schools and the outside option alone (Cardell, 1997; Train, 2009), and $0 \leq \rho < 1$ is the nest parameter. The coefficients $\tilde{\lambda}$ are in within-nest units, which keeps the cost per km comparable as ρ changes. With $\tilde{\delta}_{mj} = \delta_{mj}/(1 - \rho)$, the choice probabilities are

$$r_{ij} = \frac{\exp(\tilde{\delta}_{mj} + \tilde{\lambda}_i d_{ij}^p)}{D_i}, \quad D_i = \sum_{l \in \mathcal{J}_m} \exp(\tilde{\delta}_{ml} + \tilde{\lambda}_i d_{il}^p), \quad s_{i0} = \frac{1}{1 + e^{\pi I_i} D_i^{1-\rho}}, \quad s_{ij} = (1 - s_{i0}) r_{ij}.$$

Here r_{ij} is the probability of school j for a family that stays in the municipality's public schools, and s_{ij} is its unconditional probability. I integrate both over block groups and poverty with the survey's weights.

Poverty enters through two measures that do different jobs. The family-poverty record shifts the cost per

km. Income poverty shifts attachment to public schools: π is the added pull of the municipality's public schools, against the outside option, for income-poor families.

The cost per km, $\tilde{\lambda}_i$, is an effective cost. It reflects the cost of travel together with the busing, placement and information arrangements that families faced after a closure (Section 2.2), and the data do not separate these. Pricing the lost option assumes that enrollment shows what a school is worth to a family under those arrangements. Placing children under other closures assumes that the same arrangements would have operated. A richer specification adds $(b_R R_i + b_P P_i) \Psi_j$ to the within-nest utility, so that rural and poor families may respond to effectiveness differently from other families. The response common to all families stays in δ_{mj} . The specification used in this paper sets $b_R = b_P = 0$; Section 6.2 explains the restriction, and Appendix D reports the richer specification.

5.3 The Two Margins for Any Set of Closures

A set of closures names the schools that close in each year. Its displaced children are the children in grades 1 to 7 of a closed school in its last operating year. For the realized closures, they are the 34,975 children of Section 3.3. Each margin is defined per displaced child of the set, by group, and the number of children the set displaces is reported beside it.

Access is measured in two ways. The first needs no model. It is the km from the closed school to the nearest surviving school of the municipality serving the child's grade, and beside it the share of children with no such school within 4.8 km (Section 3.3). The second prices the lost option with the choice model. Let C_m be the closed schools of market m and $r_{iC} = \sum_{j \in C_m} r_{ij}$ the probability that family i attended one of them. Removing C_m lowers the family's expected best utility within the nest by $\log D_i - \log D_i^{(-C_m)} = -\log(1 - r_{iC})$. Dividing by r_{iC} gives the loss per displaced child, ℓ_i , which I express in km:

$$\ell_i(C_m) = \frac{-\log(1 - r_{iC})}{r_{iC}}, \quad |\tilde{\lambda}_i| \left[(d_* + x_i)^p - d_*^p \right] = \ell_i(C_m), \quad d_* = 3.9 \text{ km}. \quad (3)$$

The km-equivalent x_i is the extra distance, beyond a school d_* away, that would cost the family as much as the lost option. Because distance enters with the power p , I solve for x_i instead of dividing ℓ_i by the cost per km. The km-equivalent is a value, not a journey. It includes the value of the family's own match with the closed school, which does not depend on the map: since $\ell_i \geq 1$, every displaced family loses at least one unit of within-nest utility. At the estimates the closed schools' lower constants offset most of that part, so the km-equivalent loss is mostly its distance part (Appendix D.5). The km-equivalent values only the options inside the municipality's public schools. Conclusions drawn from school choices can reverse

when families' beliefs are inaccurate, as Kapor et al. (2020) show for beliefs about admission chances; the km-equivalent values the lost option under the information families had. Because the specification used here sets $b_R = b_P = 0$, no estimate of a school effect enters the access numbers. Markets in which no public school survives have no finite loss and are left out; no market of the realized closures is of this kind.

For a displaced child i who left the closed school $o(i)$ and entered school $d(i)$, the change in school effects is

$$\Delta_i = \phi_{d(i)} - \phi_{o(i)} = (\phi_{d(i)} - \bar{\phi}_i) - (\phi_{o(i)} - \bar{\phi}_i), \quad \bar{\phi}_i = \frac{\sum_{j \in \mathcal{S}_i} E_j \phi_j}{\sum_{j \in \mathcal{S}_i} E_j}. \quad (4)$$

Here $\mathcal{S}_i$ is the set of surviving schools of the child's municipality that serve the child's grade, and E_j is their enrollment in that grade before the closure. The first bracket is the entered-school term and the second the closed-school term; each compares a school with the average surviving school of the municipality. A negative change means that the school entered compares less favorably with the school closed. The change compares schools; it is not a change in any child's test score. Replacing ϕ with Ψ gives the change in effectiveness, which adds the school-average value-added. It counts the closed school's teachers in its last operating year and the entered school's teachers in the year after.

The school entered is recorded only for the realized closures, so the paper uses three versions of the change in school effects. All three have the same sign: a school minus the school closed, negative when less favorable. The recorded change uses the school a child entered in the pupil records, for the realized closures (Section 7.2). The change before children re-sort uses the average surviving school of the municipality, $\bar{\phi}_i$, minus the school closed: the closed-school term with its sign reversed. It needs no destination, so it is defined for any set of closures, and it changes only through which schools close and which survive. The change placed by the choice model, or placed change, replaces the recorded school entered with the schools the model places each child in, weighted by the model's probabilities. It counts the children in markets the model covers. Section 8.1 states which children each comparison of Section 8 counts.

The two margins stand side by side. Km and SD are never added, weighted against each other or put in a ratio, because the paper has no price that would convert one into the other. Together, the two definitions are the framework that Section 7 applies to the realized closures and Section 8 to closures that did not happen. Section 6 first shows what variation in the closures identifies each model.

6 Empirical Strategy

The score model and the choice model draw on different variation from the same closures. Students and teachers who change schools identify the school effects and the teachers' value-added (Section 6.1). The

schools that displaced children entered identify the cost per km, and the exit effect of Section 4.4 sets the nest parameter (Section 6.2).

6.1 School and Teacher Effects

I estimate equation (1) by least squares with the student effects absorbed. A student observed with different teachers compares those teachers with the student held fixed. Students who change schools, and teachers who teach in more than one school, link the school effects and the teachers' value-added across schools, as in the two-way fixed-effects model of Abowd et al. (1999). Closures add such moves: displaced children change schools, and the teachers of a closed school are reassigned to others. The effects are identified within the connected set of schools that these moves link. A school effect is estimated for 97.3 percent of the 334 closed schools (Appendix C.2).

Assumption 1 (Assignment and test-score shocks). *Let $\mathcal{A}$ collect the assignments of students to teachers and schools in every year of the sample. Conditional on the student, teacher and school effects, the error in equation (1) has mean zero:*

$$\mathbb{E}[e_{ist} \mid \mathcal{A}, a, \mu, \phi] = 0 \quad \text{for every observed } (i, s, t).$$

Assumption 1 allows sorting on the permanent effects. A Department that closes persistently low-scoring schools, or families who choose schools with large school effects, do not bias the estimates. The assumption excludes moves and closures that respond to transitory shocks, a central concern in value-added estimation (Rothstein, 2010). A closure removes the option to stay; it does not assign a child at random to a school or a teacher.

Three threats bear on the comparison of the school closed with the school entered, and each has its own version of the estimates. First, the Department may have closed schools after a bad year, which would put a transitory shock into the closed school's effect. I check the closing schools' residuals in the years before closure (Section 7.1). I also re-estimate the school effects for each year of closures without its last operating year and the year before, the closure-decision years. Second, displaced children's own later test scores help estimate the effects of both the school they left and the school they entered, so a shock common to those children would enter both sides of the change. The version without displaced pupils' records leaves them out from the year of the move on. Third, closed schools are small, and their effects carry more noise. I estimate every effect on two halves of the school-years, split so that no school-year shock enters both halves, and use the covariance across halves. I call the estimates on all school-years the full-sample estimates.

Noise also widens the spread of estimated effects. I correct the variances with the leave-out estimator of Kline et al. (2020), adapted to classroom-year and school-year shocks. The variances of those shocks are calibrated inputs, and Appendix C.3 reports how the corrected variances move with them. As a further check, I replace the student effects with lagged test scores, which imposes a different assumption on assignment (Chetty et al., 2014). On grades 4 to 8, the estimates of teachers' value-added from the two specifications correlate at 0.849 (Appendix C.4).

I do not run an out-of-sample test that forecasts displaced children's later test scores from the school effects and compares the forecast with a placebo, so every school-effect number in the paper is an in-sample estimate, labeled by its version.

6.2 The Choice Model: Identification and Checks

Enrollment shares alone cannot separate the cost per km from a school's constant: for any cost per km, the constants δ_{mj} can be chosen to match the shares exactly (Berry, 1994). What separates them is where displaced children went. Given the constants, the cost per km sets how fast the share of displaced children entering a surviving school falls with that school's distance rank. The moments are the shares of the 15,025 re-enrolled displaced children who entered the second, third, fourth or fifth, and sixth or farther surviving school, ranked by distance from the closed school. They are computed in each of the four groups, and the share entering the nearest school is implied. This uses the information in substitution choices emphasized by Berry et al. (2004), observed here after an actual closure. The model's families are the block-group populations of 3,023 markets, each a municipality, grade and year; 805 of these markets contain displaced children whose next school is recorded, and 885 have a closed school among the choices. The rural difference in the cost per km also rests on where rural families live across the catchments of closed schools, since the records give the rural share of a closed school's surroundings, not each child's home.

Attachment to public schools, π , is identified by the share of income-poor children among children in public schools in the survey, 67.1 percent (0.4). The nest parameter is set by a single moment: the model's exit when the closed schools leave the choice set must equal the exit effect of Section 4.4, 8.5 points in the year after a closure. The nest parameter is therefore just identified, and matching exit is not a test of the model. The exit effect is measured in grades 3 to 7 and on the closures of 2013–16, and it is applied to displaced children in grades 1 to 7 and to the closures of 2017, whose next school year followed Hurricane María; with the closures of 2017 added, the effect is 7.6 points (Appendix A.6).

Assumption 2 (Stable rankings). *Conditional on the school a family attended and on its characteristics in the choice model, a closure does not change how the family ranks the surviving schools and the outside*

option. A displaced family enters its best remaining option under the same draws of ε that made it choose the closed school.

Assumption 2 lets the Department choose closures on enrollment and other school attributes. It rules out closures that change information, transport or placement in ways the model does not represent. That matters most for who chose the receiving school, which is documented only for the closures of 2015 and 2017 (Section 2). If the Department in effect assigned children, the rank shares describe placement, and the cost per km is not a family's cost of distance. The records show more concentration than the model. For each closed school and grade, take the surviving school that the closed school's other grades in the same grade span entered most in excess of the model: 46 percent of the re-enrolled children for whom that school is defined entered it, against 21 percent in the model, and the excess appears in all four groups (Appendix D.2). Assignment to a receiving school would produce this pattern; so would bus routes or siblings in other grades, an error in where the model places families, or an error in the cost per km, and the records do not separate them. This is why Section 5.2 reads the cost per km as an effective cost.

The records say little about whether poor or rural families respond to effectiveness differently from other families. Among displaced children, the effectiveness of the school entered differs between poor and non-poor children by -0.008 SD (0.022) in urban areas and $+0.013$ (0.016) in rural areas. In the choice model, a rural difference in the response to effectiveness cannot be distinguished from the rural difference in the cost per km. I restrict both differences to zero, $b_R = b_P = 0$, and the response common to all families remains in the constants. I chose this restriction after seeing the results. In the richer specification, the poor difference is -0.063 (0.398), and the confidence set for the rural difference has no lower end over the range searched. A test that both differences are zero gives $p = 0.59$ (Appendix D.3). Under the restriction, differences between groups in where the model places children come from their locations, distance costs, attachment to public schools and the school constants, rather than different responses to effectiveness.

I estimate the model by sequential generalized method of moments (GMM) (Newey, 1984). Given the nest parameter, four parameters (the three coefficients of the cost per km and attachment) minimize a weighted distance between 17 moments and their targets: the 16 rank shares and the income-poor share. The nest parameter then solves the exit equation, and its confidence set comes from inverting that equation. The standard errors stack the two steps and use the full derivative, since the limiting distribution of GMM changes when the moments are not matched exactly (Hall and Inoue, 2003). Two untargeted quantities check the fit: the km from the closed school to the school entered, and exit by the class of the closed school (Section 7.3). Appendix D.1 gives the moments, the weights and the covariance.

Section 8 uses the model for its second job, placing the children of closures that did not happen, and I

check that use on the realized closures first. The model’s error is the school effect, or the effectiveness, of the school a child entered in the records minus that of the school the model places the child in. On the realized closures it is -0.041 SD (0.024) for rural-poor children and -0.026 (0.020) for all displaced children: at the point estimates, the model places children in schools with larger school effects than those they entered (Appendix Table D3, Panel A). I regress this error on the closed school’s size band, its distance to the nearest surviving school and the number of surviving schools, and I apply the regression to the closures that each set in Section 8 would make instead. For all displaced children and for rural-poor children, these features predict a change in the error of at most 0.007 SD. Changes in the error through other features are not measured. A placed contrast is read as a difference in where children go only if its size, less a tolerance for the model’s error, is distinguishable from zero given its standard error (Appendix E.5). The tolerance is specific to each set and at least 0.02 SD (Appendix Table D3). It also covers the change predicted when the closed school’s average test score and its enrollment are added to the features.

The check has four limits. First, the check and its thresholds were set after the model’s error on the realized closures was known. Second, the comparison uses the displaced children whose choices the choice model was estimated on, although the school effects of the schools they entered were not used in that estimation. Third, the schools entered after the closures of 2017 are those of 2017–18, the school year the storm disrupted (Section 2.1). Fourth, the standard errors hold the choice-model estimates fixed. Section 7 reports what both models estimate under these assumptions.

7 Results

The estimates match both raw facts of Section 4 and locate the gap in school effects in the schools that closed. This section reports the estimates: the school effects and the teachers’ value-added (Section 7.1), the change in school effects and where its gap between groups sits (Section 7.2), the choice model and the access loss it prices (Section 7.3), and the two margins together (Section 7.4).

7.1 School Effects and Teachers’ Value-Added

Schools differ more than their teaching staffs. Measured by the covariance of estimates on two halves of the school-years, the SD of school effects is 0.303 (0.017), against 0.075 (0.032) for school-average value-added (Table 3, Panel A). Individual teachers’ value-added has an SD of 0.203 (0.008); averaged over a school’s teachers, most of that spread cancels. With the leave-out correction instead, the SDs are 0.361 for school

effects and 0.089 for school-average value-added. The estimates use the school years 2011–19: 3,452,029 student-subject-year observations of 395,030 students in 1,274 schools, with 19,484 teacher-subject effects.

Table 3. School effects and teachers' value-added

	Across schools		Across teacher-subjects
	School effects	School-average value-added	Teachers' value-added
<i>Panel A. Standard deviation, in test-score SD</i>			
Across halves of the school-years	0.303 (0.017)	0.075 (0.032)	0.203 (0.008)
without displaced pupils	0.311	0.078	0.198
Leave-out corrected	0.361	0.089	0.225
<i>Panel B. Reliability across halves of the school-years</i>			
First split	0.565	0.087	0.29
Second split	0.561	0.059	0.24
Without displaced pupils	0.549	0.084	0.26
Observations (student-subject-years)		3,452,029	
Students		395,030	
Units: schools; teacher-subjects	1,274	1,274	19,484

Notes: Test scores of school years 2011–19; full-sample estimates of the model with student, teacher-subject and school effects. School-average value-added: the mean over a school's teachers. Across halves: from the covariance of estimates on two halves of the school-years; standard errors in parentheses, from a bootstrap over schools; none is computed for the other rows. The two splits are two random divisions of the school-years; Panel A uses the first. Without displaced pupils: their records after the closure left out. Leave-out corrected: net of estimation noise (0.34 to 0.38, 0.00 to 0.15 and 0.14 to 0.28 by column across noise inputs). Reliability: Spearman–Brown, of the full-sample estimate. School effects correlate 0.803 across periods and 0.78 across the two tests. Teacher-specific drift in value-added and a teacher-by-school component are each at most 0.10 SD.


School-effect rankings are more reproducible across splits than rankings based on school-average value-added. Across two random splits of the school-years, the reliability of the school effects is 0.565 and 0.561, against 0.087 and 0.059 for school-average value-added (Table 3, Panel B). Without displaced pupils' records from the year of their move on, the components are similar: the SD of school effects across halves is 0.311, and their reliability 0.549. Section 7.2 reports the change in school effects in the same version.

Estimated school effects are correlated across periods and tests. School effects estimated on different periods correlate at 0.803, and those estimated on the two tests the Department used correlate at 0.78. A teacher's value-added drifts by at most about 0.10 SD over the sample. A component specific to a teacher and a school has an SD of at most 0.10, with a point estimate of zero (Table 3, notes).

The pooled pre-closure test does not detect a fall in residual test scores. Figure 4 plots, for the closures of 2013–17, the closing schools' mean test-score residual relative to each school's own average, from four years before the last operating year to that year. A closure that followed a bad year would show as negative residuals in the last years. The joint test that the residuals of the last three years are zero gives $p = 0.16$;

measured against the municipality’s never-closed schools in the same years, it gives $p = 0.38$. The pooled test does not rule out differences across closure periods; Appendix Figure C1 reports those checks.

Figure 4. Checks of the school effects: test-score residuals before closure



Notes: Schools closed in 2013–17, by last operating year c (245 to 316 schools a point). Dots: the mean over schools of the school-year’s mean test-score residual less the school’s own average over its years; lines: 95 percent intervals, standard errors clustered by municipality. Residuals are from the score model (student, teacher-subject and school effects), full-sample estimates; they average zero within a school, so the figure shows timing, not a level. Joint test that the means at $c-2$, $c-1$ and c are zero: $p = 0.16$; against the never-closed schools of the municipality in the same year, $p = 0.38$. Figure C1 gives the closure periods separately. The score-model estimates are in Table 3.

7.2 The Change in School Effects and Where the Gap Sits

The change was least favorable for rural-poor children. With the full-sample estimates, the school entered minus the school closed is -0.079 SD (0.036) for rural-poor children and -0.032 (0.035) for rural non-poor children. It is $+0.011$ (0.027) and $+0.018$ (0.025) for urban-poor and urban non-poor children (Table 4). The paper’s result is the gap between rural-poor children and the rest: against all other displaced children, the rural-poor contrast is -0.084 (0.031). Against urban non-poor children it is -0.097 (0.037), and rural children’s change differs from urban children’s by -0.082 (0.036). In an additional comparison within rural areas, the change was 0.047 SD (0.020) less favorable for poor than for non-poor children, while the distance difference between the two rural groups was not distinguishable from zero (Section 4.3). Including the teachers’ value-added, the change in effectiveness gives a contrast of -0.065 (0.029) on the 19,164 records that have both.

These comparisons need the school a child entered, which only the pupil records give. Of the 32,514 children displaced in 2013–17 with a recorded next public school, 20,422 have their own poverty record and a rural share for the closed school. Of these, 19,982 have a school effect at both schools, and 19,927 have

a municipality average to compare with (Appendix Table A2). The groups follow each child's own poverty record and the closed school's rural share, and the records leave out children who exited the public schools, more often after rural closures (Section 4.4).

Table 4. The change in school effects, by group

	Displaced children					Difference		
	All (1)	Rural poor (2)	Rural non-poor (3)	Urban poor (4)	Urban non-poor (5)	Rural minus urban (6)	Rural poor minus all other (7)	Rural poor minus urban non-poor (8)
School effect relative to the municipality's average surviving school, SD								
School closed	-0.017 (0.027)	0.045 (0.035)	-0.013 (0.038)	-0.066 (0.031)	-0.077 (0.028)	0.102 (0.037)	0.104 (0.031)	0.122 (0.036)
School entered	-0.046 (0.015)	-0.034 (0.020)	-0.045 (0.026)	-0.056 (0.016)	-0.059 (0.016)	0.020 (0.020)	0.020 (0.016)	0.025 (0.021)
Change, school entered minus school closed, SD								
School effects	-0.029 (0.026)	-0.079 (0.036)	-0.032 (0.035)	0.011 (0.027)	0.018 (0.025)	-0.082 (0.036)	-0.084 (0.031)	-0.097 (0.037)
Effectiveness	-0.021 (0.026)	-0.060 (0.032)	-0.005 (0.029)	0.002 (0.030)	0.034 (0.031)	-0.057 (0.034)	-0.065 (0.029)	-0.094 (0.036)
Raw measures on these records								
Average test score, entered minus closed, SD	0.003 (0.037)	-0.057 (0.038)	-0.119 (0.068)	0.079 (0.040)	0.033 (0.073)	-0.140 (0.047)	-0.100 (0.044)	-0.090 (0.073)
Nearest surviving school, km from the closed school	2.73 (0.19)	3.52 (0.18)	3.49 (0.23)	1.93 (0.19)	2.13 (0.19)	1.55 (0.23)	1.32 (0.21)	1.38 (0.24)
Pupil records	19,927	8,071	1,825	8,487	1,544	19,927	19,927	9,615


Notes: Pupil records of children displaced by closures of 2013–17 (32,514 with a next public school, 19,927 with a school effect for both schools and surviving schools to average; Table A2). First two rows: each school's effect less the enrollment-weighted mean of the municipality's surviving schools. Change: entered minus closed; negative is less favorable. Effectiveness adds the school-average value-added (19,164 records). Raw measures on these records: average test scores (18,671 records) and km (19,925). Full-sample estimates of Table 3; Appendix C gives the versions without displaced pupils and without the closure-decision years. N: group sums of weights; a difference column counts both sides. Standard errors in parentheses, clustered by municipality; those of (1), (6), (7), which pool groups, let the mix of groups vary. Rural poor minus rural non-poor, change in school effects: -0.047 (0.020).

The two versions that address the threats of Section 6.1 leave the gap where it is. Without displaced pupils' records, the contrast is -0.097 (0.041) on 14,525 of the 19,927 records. Without the closure-decision years, it is -0.086 (0.042), against -0.083 (0.029) with the full-sample estimates on the same 19,308 records. The level is less stable than the gap. Without those years, the change for all displaced children is -0.079 (0.033), against -0.021 (0.023) with them, while the contrast barely moves. That is why the paper's object is the gap and not the rural-poor level (Appendix C.6). The raw counterpart of the gap agrees. On 18,671 of the same records, comparing the average test score of the school entered with that of the school closed gives a rural-poor contrast of -0.100 (0.044) against all other children. Against urban non-poor children, the smallest group, the raw contrast is -0.090 (0.073) and cannot be distinguished from zero.

The gap mainly reflects the schools that closed. Relative to the average surviving school of the munic-

ipality, the schools that rural-poor children left stood at $+0.045$ SD (0.035), and those of rural non-poor children at -0.013 (0.038). The schools of urban-poor and urban non-poor children stood at -0.066 (0.031) and -0.077 (0.028) (Figure 5). Every group entered schools that stood below the average surviving school, by 0.034 to 0.059 SD at the point estimates. The rural-poor contrast is thus 0.104 (0.031) in the closed-school term and 0.020 (0.016) in the entered-school term, and the change subtracts the first from the second.

Figure 5. Where the gap sits: the school closed and the school entered, by group



Notes: Closures of 2013–17; the 19,927 pupil records of Table 4, grouped by the child’s own poverty record and the closed school’s rural share. Each school’s estimated school effect (full-sample estimates, Table 3) less the enrollment-weighted mean over the municipality’s surviving schools, the dashed line. Triangle: the school closed; circle: the school entered; the arrow joins them; lines: 95 percent confidence intervals, standard errors clustered on the municipality of the closed school. Rural poor minus all other displaced children: 0.104 (0.031) for the school closed, 0.020 (0.016) for the school entered. Table 4 holds the numbers.

School size goes with about a fifth of the gap against urban non-poor children. Within a municipality, smaller schools have larger estimated school effects: -0.114 (0.038) per log point of enrollment, estimated jointly with grade span and catchment breadth across 1,257 schools. I multiply this gradient by the change in log enrollment from the school closed to the school entered. On the 18,923 records with both enrollments, the product differs between rural-poor and urban non-poor children by -0.021 (0.008), against a gap of -0.098 in the change itself on those records. The gradient is an association (Section 5.1), and I do not attribute the rest of the gap to school size or to any other characteristic.

7.3 The Choice Model and the Access Loss

Rural families have a lower estimated cost per km than urban families, and poor families a slightly higher one. The within-nest coefficient on km to the power 0.8 is -2.785 (0.577) for urban non-poor families. The

rural difference is +1.219 (0.637), and the poor difference is -0.202 (0.094) (Table 5, Panel A). Among non-poor families, the rural coefficient is 0.56 (0.15) of the urban one, a ratio whose interval excludes one, although the 95 percent interval of the rural difference itself just includes zero; in urban areas, the poor coefficient is 1.07 (0.04) of the non-poor one. Attachment to public schools for income-poor families is 1.411 (0.062). The nest parameter is 0.847 (0.031), with a confidence set from 0.783 to 0.904.

The model comes close on how far from the closed school displaced children enrolled, but not on which school they entered. The model's mean km from the closed school to the school entered, which is not targeted, lies 0.08 to 0.34 km below the data across the four groups, with data and model on the same 13,058 records and on the construction of Section 3.3, not the pupil records' own km of Table 2 (Table 5, Panel B). The share entering the nearest surviving school is over-predicted in all four groups, by 5.2 and 9.7 points for rural-poor and rural non-poor children and by 3.5 and 2.5 points for urban-poor and urban non-poor children. The share entering the second-nearest school is under-predicted. The over-identification test of the 17 moments gives $p = 0.038$, with the covariance of the targets treated as known, and two of the 16 rank moments, the second-nearest shares of the two rural groups, lie more than two of the displayed standard errors from their targets (Appendix Table D1). Exit, matched in the aggregate by construction, is misplaced by origin: after rural closures it is 12.0 (2.5) points in the data against 9.2 in the model, and after urban closures 3.2 (1.2) against 7.9, which is why the paper uses the model's exit only in the aggregate.

Priced by the choice model, the access loss is again a rural loss. Per displaced child, it is 1.84 (0.11) km-equivalents for rural-poor children and 1.91 (0.15) for rural non-poor children. It is 0.97 (0.08) for urban-poor and 1.00 (0.09) for urban non-poor children (Table 6). These figures cover the 28,024 displaced children of the 885 markets in which a closed school was among the choices. The model types families by their block group's rural indicator, not by the closed school's rural share, so its four groups do not coincide with those of the km measure (Appendix D.6). Their standard errors come from the delta method in the five coefficients, clustered by municipality, and hold the inputs listed in the table's notes fixed.

The gap between rural and urban children stands once access is priced. The loss of rural-poor children is 1.84 (0.16) times that of urban non-poor children, against a ratio of 1.62 (0.15) in km to the nearest surviving school. The difference between the two groups is 0.84 (0.12) km-equivalents. The priced ratio depends largely on the conversion itself: rural families' cost per km is smaller, so a similar loss of utility converts into more km. In units of within-nest utility the ordering of the two groups' losses is not determined over the confidence set for the rural difference in the cost per km (Appendix D.5), so the paper's claim about priced access is the ordering in km-equivalents, beside the ordering in km of Section 4.3, which needs no conversion. Across the variants of the choice model in Appendix D.5, the ratio lies between 1.78 and 1.85.

Table 5. The choice model: estimates and fit

<i>Panel A. Estimates</i>			
	Estimate	Standard error	
Distance, km ^{0.8} : within-nest coefficient			
Urban non-poor family	−2.785	0.577	
Rural family, difference	1.219	0.637	
Poor family, difference	−0.202	0.094	
Leaving the municipality’s public schools			
Nest parameter	0.847	0.031	
Income-poor family: attachment to public schools	1.411	0.062	
Implied ratio of the cost of distance			
Rural over urban family (non-poor)	0.56	0.15	
Poor over non-poor family (urban)	1.07	0.04	
Over-identification test (17 moments, 4 parameters): p		0.038	
Markets (municipality-grade-years)		3,023	
<i>Panel B. The model against the data</i>			
	Data	Standard error	Model
Entered the nearest surviving school, percent (implied by the targeted shares)			
Rural poor	44.6	2.3	49.7
Rural non-poor	37.2	3.6	46.9
Urban poor	38.7	3.0	42.2
Urban non-poor	37.9	3.2	40.4
Km to the school entered (not targeted)			
Rural poor	4.73	0.14	4.65
Rural non-poor	4.83	0.19	4.62
Urban poor	2.91	0.16	2.73
Urban non-poor	3.28	0.17	2.94
Left the municipality’s public schools after a closure, percentage points (not targeted)			
Rural closed schools	12.0	2.5	9.2
Urban closed schools	3.2	1.2	7.9
Re-enrolled children (shares)		15,025	
Re-enrolled records with a model value (km)		13,058	
Children in the exit estimates		7,138	

Notes: Nested logit of the school a displaced child entered; outside option: leaving the municipality's public schools. Panel A: a family's within-nest distance coefficient sums the rows that apply; the exponent is fixed. Poor: the Department's family-poverty record; income-poor, survey income poverty. All families respond alike to effectiveness, a restriction chosen after the results (Table D2). The nest parameter rests on one moment, the effect of a closure on leaving; the test's p-value is simulated from its first-order null distribution with the covariance of the targets treated as known. Ratios: delta method. Panel B: Model, predicted at the estimates; Data standard errors clustered by municipality; km from the closed school, a straight line times a road ratio, on the 13,058 re-enrolled records with a model value, both weighted by children; leaving: event study, closures of 2013–16, by the Department's rural tag (Model: by catchment).

Table 6. The access loss priced by the choice model

	Displaced children				Rural poor over urban non-poor (5)	Rural poor minus urban non-poor (6)	Children (7)
	Rural poor (1)	Rural non-poor (2)	Urban poor (3)	Urban non-poor (4)			
Km to the nearest surviving school	3.43 (0.16)	3.37 (0.22)	1.91 (0.19)	2.12 (0.18)	1.62 (0.15)	1.31 (0.22)	34,975
Km-equivalents in the choice model	1.84 (0.11)	1.91 (0.15)	0.97 (0.08)	1.00 (0.09)	1.84 (0.16)	0.84 (0.12)	28,024

Notes: Closures of 2013–17. First row: km from the closed school to the nearest surviving school of its municipality that serves the grade, as in Table 2; all displaced children. Second row: the fall in a family’s expected best utility among the municipality’s public schools (within the nest) when the closed schools leave its choice set, at the estimates of Table 5, expressed as the extra distance that costs the family as much; the 28,024 displaced children of the 885 markets with a realized closure. The two rows rest on different children. The model types families by their block group’s rural indicator, not by the closed school’s rural share (Appendix D.6). Standard errors in parentheses: first row, clustered by municipality; second row, delta method in the coefficients, clustered by municipality, with enrollment shares, the set of closed schools, population weights, the reference distance and the distance power held fixed.

Over the confidence sets for the rural difference in the cost per km, it lies between 1.4 and 2.1, above one at every point evaluated.


7.4 The Two Margins Together

Rural-poor children stood on the less favorable side of both margins, in two units that the paper keeps apart (Figure 6). In km, rural children’s nearest surviving school was farther than urban children’s (Panel A), and priced by the choice model rural-poor children’s loss is again the larger. In school effects, their change was 0.084 SD (0.031) less favorable than that of all other displaced children (Panel B).

The two panels rest on different records of the same closures. Panel A covers all 34,975 displaced children, grouped by the closed school’s rural and poor shares. Panel B covers the 19,927 pupil records of Table 4, grouped by each child’s own poverty record and the closed school’s rural share. The access gap is similar on both: rural-poor children’s nearest surviving school was 1.29 km (0.20) farther than all other children’s among the 34,975, and 1.32 km (0.21) farther on the pupil records (Tables 2 and 4).

Two facts carry into Section 8. Because rural schools lie farther apart, a set of closures that keeps isolated schools open can lower the distance that rural children face. Because the school-effect gap sits in the schools that closed, while every group entered schools below the average survivor by a similar amount, another set of closures changes the second margin mainly through which schools it closes. Section 8 compares sets of closures first at the schools they close, before children re-sort, and then with children placed by the choice model.

Figure 6. The two margins by group: distance to the nearest surviving school and the change in school effects



Notes: Closures of 2013–17. A: mean distance from the closed school to the nearest surviving school of the municipality that serves the child’s grade, a straight line times a road ratio by class of school; all 34,975 displaced children, grouped by each closed school’s rural and poor shares. B: the estimated school effect of the school entered minus that of the school closed, in SD of test scores; the 19,927 pupil records of Table 4, grouped by the child’s own poverty record and the closed school’s rural share; dashed line: zero, no change. The two panels rest on different records of one set of closures. Rural poor minus all other displaced children: 1.29 (0.20) km in A (1.32 (0.21) on B’s records), -0.084 (0.031) SD in B. Circles: group means; lines: 95 percent confidence intervals, standard errors clustered on the municipality of the closed school. Table 2, Panel B, and Table 4 hold the numbers.

8 Counterfactuals

This section varies three decisions in turn: which schools close at the Department’s counts, how many close, and, with the list fixed, where displaced children and teachers go. Section 8.1 states what every comparison holds fixed, and Sections 8.2 to 8.4 take the decisions in order.

8.1 Design and What Is Held Fixed

A set of closures changes two things for a group: how many of its children are displaced, and how each margin changes per child displaced. A group’s burden is the product of the two, and the tables print both. A list that displaces fewer rural-poor children is thus kept apart from a list that changes less for each of them.

Each list at the Department’s counts uses only what the Department held the year before a closure: each school’s enrollment, its place on the map and its average test score. The 19 closures held in every list stay closed, and the other 315 of the 334 closures are chosen again. In every municipality, year and size band, a list closes as many schools as the Department did; I call these the Department’s counts. The schools come from the comparable schools: those open that year and no larger than the largest school closed that year.

Size bands are the thirds of school size among that year's closed schools. Each year starts from the schools left open by the realized closures of earlier years, so a comparison describes one year's choice at a time.

Five lists vary whether the choice uses the map, average test scores, both or neither (Appendix Table E1). They are smallest first; smallest first with an alternative within 4.8 km; lowest average test scores first; the same with an alternative within 4.8 km; and nearest alternative first. Each score list is chosen on one half of the school-years and measured on the other. A random list draws, in every municipality, year and size band, as many comparable schools as closed there. Holding counts within size bands holds enrollment roughly fixed, one of the Department's stated criteria (Section 2.1), so random lists isolate which comparable schools closed. One thousand random lists form a reference: a spread across possible choices, not a sampling interval.

Using the enrollment weights of the realized closures, the choice model and school-effect estimates supply a placed value for 32,398 of the 34,975 displaced children: 93.7 percent of the rural-poor weight and 91.9 percent of the other weight. On this base, 1,628 children (4.7 percent) have no grade market in the choice model, and 858 (2.5 percent) lack the closed school's effect; small remaining exclusions account for the rest. Section 6.2 checks the placements on realized closures. On the separate pupil-record sample of Table 4, 18,834 of 19,927 records have a placed value. The 1,093 records without one (5.5 percent) differ from the rest: their recorded rural-poor contrast is -0.282 SD (0.115), against -0.074 (0.031) for the covered records.

The bases differ by measure. Distances average over the children each list displaces, and the placed change over those of them the model places. Differences from the realized closures in the change before children re-sort average over the children the realized closures displaced, so that lists displacing different children share one base; Figure 7 draws that change per child each list displaces. Throughout, the section compares rural-poor children with all other displaced children, the contrast of Section 7.2; for access, all other children mix rural non-poor children, whose loss matched rural-poor children's, with urban children (Appendix Table E3 gives the four groups for random lists).

The comparisons of closure lists hold the rest of the system as it was in the year of closure. Surviving schools keep their school effects, teachers and other children, and they do not respond to their intake. Families live where they lived and choose as the families whose enrollment shares built the model's school constants. The calculations assume that the move itself has no effect, and no seat limit is imposed: lists hold counts, not seats. The historical-enrollment proxy for seats removed uses grades 3 to 8, while displaced children are counted in grades 1 to 7; the two are not a capacity comparison. On this proxy, the realized

closures remove 35,455 seats and the five lists 29,272 to 35,007; the lists displace 33,242 to 37,339 children, against 34,975 (Table 7, Panel A). Savings enter only Section 8.3, and transport costs enter nowhere.

Table 7. Closure lists at the Department's counts: both margins, by group

<i>Panels A and B. The lists: seats, children displaced, and the km to the nearest surviving school</i>					
	A. Children and seats			B. Km to the nearest surviving school	
	Seat proxy grades 3–8 (1)	Children displaced (2)	Percent rural poor (3)	Rural poor (4)	All other (5)
Realized list	35,455	34,975	39.7	3.43	2.15
Smallest first	29,272	33,242	43.6	3.79	2.25
Smallest first, alternative within 4.8 km	30,134	33,698	39.8	3.11	2.07
Lowest average test scores first †	35,007	36,850	40.9	3.68	2.18
Lowest average test scores first, alternative within 4.8 km †	34,855	37,339	37.9	3.11	2.01
Nearest alternative first	32,979	35,217	34.3	2.67	1.50
<i>Panels C and D. Change in school effects, school entered or surviving minus school closed, SD</i>					
	C. Before children re-sort		D. Placed by the model		
	Rural poor (6)	All other (7)	Rural poor (8)	All other (9)	
Realized list	−0.054 (0.041)	0.046 (0.029)	−0.021 (0.044)	0.024 (0.033)	
Difference from the realized list					
Smallest first	−0.098 (0.039)	−0.089 (0.030)	−0.093 (0.036)	−0.068 (0.036)	
Smallest first, alternative within 4.8 km	−0.081 (0.035)	−0.084 (0.029)	−0.085 (0.034)	−0.068 (0.034)	
Lowest average test scores first †	0.048 (0.042)	0.094 (0.032)	0.072 (0.052)	0.141 (0.046)	
Lowest average test scores first, alternative within 4.8 km †	0.057 (0.036)	0.104 (0.034)	0.080 (0.047)	0.144 (0.048)	
Nearest alternative first	−0.054 (0.035)	−0.028 (0.030)	−0.062 (0.040)	−0.008 (0.034)	


Notes: Closures of 2013–17; lists as in Table E1. A: the seat proxy is closed schools' earlier maximum enrollment in grades 3 to 8; children displaced are in grades 1 to 7, so the columns are not a capacity comparison. B: km per child each list displaces, as in Table 2, which gives the realized list's standard error; none is computed for the other lists. C: surviving schools' enrollment-weighted school effects minus the closed school's, per child the realized list displaced; standard errors are the largest of four imputation rules (Table E10), a rule changed after the results were seen, under which smallest first's rural-poor difference (standard error 0.039, against 0.035) is no longer distinguishable from zero. D: per child the choice model places (93.7 percent of the realized list's rural poor); standard errors combine the delta method in the model's coefficients with the noise in the school effects, unchecked by simulation (Appendix E.5). †: added after the results were seen; ranked on one half of the school-years and measured on the other; B pools the two directions, C and D average them. No multiple-comparison adjustment rejects a list in C or D.

Four parts of this section were added after the results were seen: the two score lists, the placement of children under other lists, the fewer-closures exercise and the rule that sets the standard errors of the change before children re-sort. Appendix Table E2 lists the reported lists, orders and counts, along with additional documented variants and their status. Statements of direction are adjusted for multiple comparisons within sets specified before each set's own results were seen (Appendix E.5).

8.2 Different Schools at the Department's Counts

Random lists of comparable schools reproduce the rural gap in distance (Figure 7, Panel A). Across 1,000 random lists, the median distance to the nearest surviving school per child displaced is 3.68 km for rural-poor children, 3.60 for rural non-poor, 1.96 for urban-poor and 2.13 for urban non-poor children. The realized closures give 3.43, 3.37, 1.91 and 2.12. They lie at the 1.2nd and 2.2nd percentiles of the random lists for the two rural groups, and at the 28.2nd and 47.9th for the two urban groups (Appendix Table E3). Closing comparable schools at these counts produces the rural gap.

Figure 7. The realized list among random lists at the Department's counts



Notes: Closures of 2013–17. A random list closes, in every municipality, year and size band, as many schools as closed there, drawn from comparable schools; the 19 held closures close under each. Diamond: the realized list; bar: the 5th to 95th percentile of 1,000 random lists, a spread across lists, not a sampling interval; tick: their median. Both panels are per child each list displaces. A: km from the closed school to the nearest surviving school of the municipality. B: the enrollment-weighted school effects (full-sample estimates, Table 3) of the municipality's surviving schools minus the school closed, before children re-sort; negative is less favorable. All other: the three groups other than rural poor. Table E3, Panels A and D, holds related figures; its middle 95 percent runs from the 2.5th to the 97.5th percentile. B's levels are not read in the text: they move with the version of the school effects (Section 7.2).

Random lists also give a gap of similar size in the change in school effects before children re-sort (Figure 7, Panel B). Under the realized closures, this change is 0.099 SD less favorable for rural-poor children than for all other displaced children. Across random lists, the gap averages 0.093 SD per child displaced, and the middle 95 percent of lists lie between 0.041 and 0.141 (2.5th to 97.5th percentiles; Figure 7's bars mark the 5th to 95th).

At the Department's counts, the lists that use the map lower distance for rural-poor children. Their nearest surviving school is 3.43 km away under the realized closures, 3.11 km under smallest first with an alternative within 4.8 km and 2.67 km under nearest alternative first (Table 7, Panel B). The share with none within 4.8 km falls from 19.6 percent to 9.1 and 10.7 percent (Appendix Table E1). Summed over

the rural-poor children each list displaces, distance falls by 12.5 percent [2.5, 21.0] and 32.4 percent [23.0, 41.2], with 95 percent intervals in brackets (Appendix Table E1, notes). In the choice model, the access loss per rural-poor child is 1.84 (0.11) km-equivalents under the realized closures and 1.79 (0.11) and 1.65 (0.12) under the two lists that use the map but not scores, figures that carry no statement of direction (Appendix Table E4). Under smallest first, which uses neither the map nor scores, the figures are 3.79 km and 2.20 (0.09) km-equivalents. Lists also change who is displaced: nearest alternative first displaces 35,217 children, 34.3 percent of them rural poor, against 34,975 and 39.7 percent (Panel A).

Before children re-sort, no list's change in school effects for rural-poor children can be distinguished from the realized closures' after adjusting for the five comparisons (Table 7, Panel C). The realized change is -0.054 SD (0.041) per rural-poor child displaced. This figure and the gap of 0.099 SD above are per child the realized closures displaced, the 34,037 children of the 315 closures chosen again, whereas Section 7.2's closed-school term, $+0.045$ for rural-poor children and 0.104 for the contrast, is on the 19,927 pupil records of Table 4. At the point estimates, the differences from the realized closures are -0.098 (0.039) for smallest first, -0.081 (0.035) with an alternative within 4.8 km and -0.054 (0.035) for nearest alternative first, and $+0.048$ (0.042) and $+0.057$ (0.036) for the two score lists. A closed school without a school effect counts as zero in these figures, and this coverage differs between lists (Appendix E.1).

Placed by the choice model, the children each list would displace show the same result (Table 7, Panel D). The realized closures give rural-poor children a placed change of -0.021 SD (0.044); the recorded change of Section 7.2, -0.079 (0.036), covers other children, and its difference from the placed change includes the model's error (Section 6.2). The lists' contrasts with the realized closures run from -0.093 to $+0.080$ SD, and none can be distinguished from zero given the model's error, for which each list carries a tolerance of 0.020 to 0.033 SD. Conditional on the delta-method and split-half standard errors, and allowing for the model's error, no list's placed change for rural-poor children differs from the realized closures' by more than 0.233 SD (95 percent, simultaneous over the five lists; Appendix Table E5). The bound is nearly three times the gap of Section 7.2, so the data cannot rank the lists on this margin.

Re-sorting adds little. For the size and map lists, nearly all of the placed contrast comes from which schools close: for smallest first, -0.110 SD (0.034) comes from the schools closed and $+0.017$ (0.014) from the schools entered (Appendix Table E5). On lists, then, this section's conclusions concern access, where the map lowers distance for rural-poor children. Here and in Section 8.3, the standard errors for placed children combine the delta method in the model's coefficients with the noise in the school effects, and they have not been checked by simulation; such a check could leave a figure without an interval but cannot change a statement of direction, because none was made (Appendix E.5).

These comparisons support only a narrow statement about the Department’s choice: among comparable schools at its counts, its closures left rural children nearer a surviving school than almost all random choices. The Department did not compute school effects, and none of these facts shows what it weighed.


8.3 Closing Fewer Schools

Closing fewer schools is a different decision: it keeps some realized closures open and gives up their savings. From the 315 closures chosen again in Section 8.2, with the held closures still closed, I spare closures one at a time in two orders. “Farthest from a surviving school first” spares first the closure whose children, in its worst-served grade, were farthest from a surviving school serving that grade. The Department could have computed it from the map. “Largest model access gain first” spares at each step the closure that most reduces the choice model’s access loss for all children. It requires the model and is chosen and measured at the same estimates. One thousand random orders of the same closures give a reference band, the 5th to 95th percentiles at each count. I read the paths at 32, 63 and 95 closures spared, a tenth, a fifth and three tenths of the 315. I also read them at 47, where the first order brings the share of rural-poor children with no surviving school within 4.8 km close to zero (Figure 8).

Keeping open the 47 closures farthest from a surviving school nearly empties the far tail. The share of rural-poor children with no surviving school within 4.8 km falls from 19.6 to 0.14 percent, and the remaining closures displace 30,441 children instead of 34,975. The 47 closures spared avoid 30.4 percent [27.7, 33.2] of rural-poor children’s access loss in the choice model and 16.6 percent [14.3, 19.0] of all other children’s, about a quarter in all (Figure 8, Panel B). For rural-poor children the share avoided is 20.9, 37.6 and 49.5 percent at 32, 63 and 95 closures spared. “Largest model access gain first” avoids more of the model’s loss, the measure it is chosen on: 43.2 percent [40.3, 46.2] for rural-poor children at 47. It leaves 10.4 percent of them with no surviving school within 4.8 km. For all children, both orders avoid more access loss than the 95th percentile of the random orders at each count.

The change in school effects before children re-sort does not move by an amount the data can distinguish (Figure 8, Panel C). With 47 closures spared in the order “Farthest from a surviving school first,” the remaining closures’ change differs from the realized closures’ by -0.024 SD (0.022) for rural-poor children and -0.010 (0.009) for all other children, per child the realized closures displaced. At 32, 63 and 95 the rural-poor differences are -0.013 (0.009), -0.015 (0.023) and -0.005 (0.034). None can be distinguished from zero after adjusting for the four counts. Against random orders of the same closures, this path lies below their median at all four counts, and at 47 at their 0.4th percentile, a spread across orders rather than sampling error (Appendix E.3). With children placed by the choice model, per child the realized closures displaced,

Figure 8. Closing fewer schools: savings forgone, access loss avoided, and the change in school effects



Notes: Horizontal axis: closures of 2013–17 kept open, of 315. “Farthest from a surviving school first” spares first the closure whose children were farthest from one; “Largest model access gain first” spares the closure that most reduces the choice model’s access loss (the Department could not compute it). The exercise, its orders and its bands were added after the results were seen (Table E2). Gray: 5th to 95th percentile of 1,000 random orders (a spread, not sampling error). A: low end \$12,197 a closure (every order); the high end adds \$121.16 a child and \$31,544 a closure for staff posts. B: 28,024 children of the choice model; access counts each child once, while savings recur yearly. Bars: pointwise 95 percent intervals at 32, 47, 63 and 95; shaded: simultaneous 95 percent bands (C: first order only; Figure E1). C: standard errors are the largest of four imputation rules (Table E10), a rule changed after the results were seen; more than half of the variance is imputed at 95 under the first order and at 7 of 8 estimates under the second.

the rural-poor differences are -0.010 (0.011), -0.016 (0.020), -0.009 (0.022) and -0.005 (0.034) at 32, 47, 63 and 95; none can be distinguished from zero given the model's error (Appendix Table E6).

Keeping schools open forgoes savings that recur every year (Figure 8, Panel A). I count a closure's savings at the two ends of a range of cost assumptions and treat both ends as fixed amounts. The low end is \$12,197 a year, the fixed part of a school's facilities spending net of what a closed building still costs. The high end adds \$121.16 per child enrolled in the closure's last year and \$31,544 for staff posts (1.06 posts per closure), about \$60,600 at the median closed school of 139 children. All 315 closures save \$3.84 million to \$19.39 million a year on this account. Keeping open the 47 farthest from a surviving school forgoes \$0.57 million to \$2.93 million a year, about 15 percent at either end. The high end assumes that a released post is not refilled (Appendix F).

Dividing savings forgone by access loss avoided gives a break-even value: the value of one km-equivalent of access for one child for one year at which the two are equal. With 47 closures spared under "Farthest from a surviving school first," it is \$55 [\$52, \$58] at the low end and \$279 [\$265, \$294] at the high end, for all children. Across 32 to 95 closures spared, it runs from \$52 to \$67 at the low end and from \$269 to \$334 at the high end (Appendix Table E6). The intervals carry the sampling error in access only. Access counts each displaced child once, for the school year of displacement, while savings recur every year; children who would have entered the closed schools in later years are not counted. Savings are also gross of any added cost of transport, which would lower the break-even values.³

8.4 Holding the Closures Fixed

Once the list is fixed, the estimated gains within the observed distance are modest and smaller for rural-poor children than for urban non-poor children at the point estimates. Larger gains require a wider distance limit. These exercises hold every closure as it was and change the school a displaced child enters or the posts displaced teachers fill.

Reach. Without going farther from the closed school than the school they entered, rural-poor children could have entered somewhat more effective schools; their estimated gain is smaller than urban non-poor children's at the point estimates (Table 8; Appendix Table E9). For each displaced child's pupil record, I find the most effective surviving school with a seat, capped at earlier maximum enrollment, no farther from the closed school than the school entered, holding every other child's school fixed. The school is chosen on one half of the school-years, and its effectiveness, the school effect plus the school-average value-added, is

³ Appendix E.3 adds two orders that rank closures by past test scores and the map, which the Department held.

measured on the other. For rural-poor children the school found is 0.048 SD more effective [0.028, 0.073] and 0.31 km nearer the closed school; 79 percent of their records keep the school entered. For urban non-poor children the gain is 0.067 [0.031, 0.102], at 0.37 km nearer. For rural-poor children, one more km raises the gain to 0.064 [0.040, 0.090]. Any surviving school of the municipality raises it to 0.204 [0.146, 0.269], at 3.1 km [2.4, 3.9] more. The exercise uses 12,564 pupil records.

Table 8. Closures held fixed: schools within reach

	Rural poor
Gain in effectiveness, SD	
No farther than the school entered	0.048 [0.028, 0.073]
Within one more km	0.064 [0.040, 0.090]
Anywhere in the municipality	0.204 [0.146, 0.269]
Change in km to the school entered	
No farther than the school entered	−0.31
Within one more km	−0.24
Anywhere in the municipality	3.1
Children	5,369

Notes: The realized closures and every other child's school are held fixed. Each record of a displaced child is assigned to the most effective surviving school with a seat among the schools no farther from the closed school than the school the child entered, within one more km, or anywhere in the municipality; seats are capped at each school's earlier maximum enrollment. Effectiveness: school effects plus school-average value-added; schools are picked on one half of the estimates and scored on the other. Km: from the closed school, the assigned school less the school entered. Km, 95 percent bootstrap intervals: [−0.33, −0.16] within one more km; [2.4, 3.9] anywhere in the municipality. Last row: the group's sum of weights; 12,564 pupil records are assigned in all. In brackets: 95 percent bootstrap intervals over municipalities.

Teachers. Within the observed pool of displaced teachers and the posts they took, the point-estimated gains are small. Of the 2,145 teachers the closures displaced, 626 taught a subject with a standardized test at exactly one surviving school the next year and have an estimate of value-added. I reassign them among those posts, within municipality, subject profile and grade band, to maximize the value-added children receive, with rural-poor children counted twice, a weight I chose. The average value-added rural-poor children receive increases by 0.0007 SD, and 201 posts change teacher (Appendix Table E9, which reports point estimates without sampling intervals for this exercise). The comparison keeps teachers and posts fixed and does not establish whether the assignments respect seniority rules or teachers' preferences.⁴

Across the three decisions, access moves with which schools close and how many; the data cannot rank the resulting changes in school effects. With closures fixed, choosing the school a displaced child enters

⁴ Appendix E.4 also considers information about receiving-school effectiveness under an assumed response to a report card.

offers modest estimated gains within the observed distance, and reassigning the observed teachers among the observed posts gives a small point-estimated gain.

9 Conclusion

Puerto Rico's closures of 2013–17 fell differently on the two margins. Access was a rural loss: the nearest surviving school stood about 3.4 km from rural children's closed schools against about 2 km for urban children's. About one rural child in five had no surviving school within 4.8 km, against about one urban child in twenty. The reason is where small schools stand: closure followed enrollment, small schools are more often rural, and rural schools lie farther apart. The change in school effects was about 0.08 SD less favorable for rural-poor children than for all other displaced children. That gap mainly reflects the schools that closed. Relative to their municipality's average surviving school, rural-poor children left schools that stood higher than other displaced children's, while displaced children in every group entered schools below that average by similar amounts.

The three decisions move the two margins differently. Which schools close moves access: at the Department's counts, random lists of comparable schools reproduce the rural gap, and lists that use the map lower the distance rural-poor children face. On school effects the data cannot rank the lists: with children placed by the choice model, no list's change for rural-poor children can be distinguished from the realized closures', and differences up to 0.233 SD cannot be ruled out. How many schools close decides the far tail: keeping open the 47 closures farthest from a surviving school nearly empties it, while forgoing savings that recur every year. Once the list is set, more effective schools within the observed distance offer rural-poor children gains that are smaller than urban non-poor children's at the point estimates; larger gains require a wider distance limit. On the Department's own choice, the comparisons support only a narrow statement. Its closures left rural children nearer a surviving school than almost all random choices of comparable schools, and nothing here shows what it weighed.

This paper does not follow the children who left the public schools or those already enrolled in receiving schools. Nor does it study what the closures did to any child's own test scores, or the closures that followed Hurricane María. What travels is a way of reading a contraction. When a shrinking system closes schools by enrollment, where its small schools stand decides which children end up far from a school. The map shows this before any list is drawn. How the change in school effects differs across groups depends on how the schools that close compare with the schools around them, which enrollment shows only in part. Kept apart, each in its own units, the two margins show which decision moves which.

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.
- Abdulkadiroğlu, A., Pathak, P. A., and Walters, C. R. (2026). Who gets what in education: Can school matching improve student achievement? NBER Working Paper 34936, National Bureau of Economic Research.
- Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.
- Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262.
- Berry, S. T., 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.
- Beuermann, D. W., Jackson, C. K., Navarro-Sola, L., and Pardo, F. (2023). What is a good school, and can parents tell? Evidence on the multidimensionality of school output. *Review of Economic Studies*, 90(1):65–101.
- Bobonis, G. J., Moradi, R., Sotomayor, O. J., Velez, J., and Wagner, J. (2026). The effects of school consolidation on students and teachers: Evidence from an underperforming system. Working paper.
- 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.
- Chrzan, M. L. and Pearman, F. A. (2026). A sandbox for hard choices: Using simulation to explore school closure scenarios and their consequences. EdWorkingPaper 26-1440, Annenberg Institute at Brown University.
- Epplé, 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.
- 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 Sølvesten, 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.
- 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 (2016). Perfil del Departamento de Educación de Puerto Rico (año académico 2015–2016). Technical report, Puerto Rico Department of Education. December 2.
- PRDE (2017a). FFY 2015 Part B State Performance Plan and Annual Performance Report. Technical report, Puerto Rico Department of Education. Clarifications version, May 3.
- 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 Herramientaje 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.
- Walters, C. R. (2018). The demand for effective charter schools. *Journal of Political Economy*, 126(6):2179–2223.

A Data

This appendix documents the data behind Sections 2 to 4: the variable inventory (Appendix A.1), the two flows of displaced children (Appendix A.2), schools, enrollment and closures by year (Appendix A.3), the construction of distance (Appendix A.4), the groups and the survey (Appendix A.5), and the event study of exit (Appendix A.6).

A.1 Variable Inventory

Table A1 lists each variable the paper uses, the extract it comes from, how it is built and where it is used. Five conventions apply throughout. First, a school year is labeled by the calendar year in which it ends, and a closure by the school's last operating year: a closure of 2017 is a school that last operated in 2016–17, whose children enrolled elsewhere in 2017–18. Second, grades in the raw files are coded two above the actual grade, and I convert them once, before any restriction. Third, the grade of enrollment defines displacement and the choice model's markets, while the grade of the test defines the subject, grade and year within which a score is normalized; 0.3 percent of children took the test in a grade other than their grade of enrollment. Fourth, a school's total enrollment appears on each of its grade rows in the enrollment extract, so I take it once per school-year and never sum it across grades. Fifth, scores are merged on the child and the school year, never on the grade. The child identifier matches between the pupil records and the teacher links for 98.1 percent of children; school codes match for 87 percent of schools, the rest being high schools and schools outside the enrollment extract.

A.2 Samples: Two Flows of Displaced Children

Table A2 sets out the two flows of displaced children; they are not nested. The first starts from the enrollment extract: 36,109 children in grades 1 to 7 of closed schools in their last year, of whom 1,134 belong to schools without a rural or poor share, leaving the 34,975 displaced children of the main text. The second starts from the pupil records: 32,514 displaced children, from kindergarten to grade 8, with a recorded next public school. Removing the 11,548 without a family-poverty record and the 544 whose closed school has no rural share leaves 20,422. Requiring an estimated school effect at the school closed (343 removed), at the school entered (75) or at both (22) leaves 19,982, and requiring the average surviving school of the municipality (55) leaves 19,927. On the 19,925 of these records that match their closed school in the enrollment extract, the mean km to the nearest surviving school is 3.52, 3.49, 1.93 and 2.13 for the four groups, against 3.43, 3.37, 1.91 and 2.12 for all 34,975 displaced children.

The records of where re-enrolled children went (Section 4.4) come from the extract of the pupil records prepared for the choice model. Of its 22,690 records of displaced children who re-enrolled in a public school, 15,573 re-enrolled in a surviving school of their set of alternatives and have a family-poverty record; 15,025 of these fall in the two poverty classes used to form the four groups, and 548 do not. Of the 22,690, 19.5

Table A1. Variables: sources, construction and use

	Source	Construction	Used in
Last operating year	Department's school file	Calendar year in which the school's last operating school year ends	Sections 2 to 8
School enrollment	Enrollment extract	Total enrollment, taken once per school-year	Table 1, Figure 2, lists
Location	Department's school file	Coordinates by year	Distances, Figure 1
Rural designation	Department's school file	The Department's rural or urban designation of the school	Exit by area; Table A3
Rural share near a school	Census survey, block groups	Population-weighted rural share around the school (Table A3)	The four groups
Family poverty	Pupil records	The child's poverty flag, its most frequent value across years	The four groups; cost per km
Income poverty	Census microdata, 2013–17	Income below the Census poverty threshold	Attachment to public schools
Rural block group	Census survey, block groups	Log population density below –8.681	Families of the choice model
Grade	Pupil records; teacher links	Coded two above the grade, converted once; enrollment grade and test grade kept apart	Displacement; normalization of scores
Test score	Pupil records	Normalized within subject, grade and year	Score model; average scores
Average test score of a school	Pupil records	Mean normalized score of its pupils in the year before the closure year	Table 2, lists
Displaced child	Enrollment extract linked to school file	A child in grades 1 to 7 of a closed school in its last operating year	Table 1; distance
Next public school	Pupil records	The child's school of enrollment the next year	Schools entered; Table 4
Km to the nearest surviving school	School file	Straight line times the road ratio of the school's class (1.52 rural, 1.42 urban), within the municipality, serving the grade	Table 2, Figure 3, lists
Km between schools, pupil records	Pupil records	Km from the school closed to the school entered, as recorded	Table 2; rank shares
Km from a block group to OpenStreetMap a school		Km along the road network, routed	Choice model
Exit	Pupil records, matched	Absent from the public records the next year, or enrolled in another municipality's public school	Nest parameter; Section 4.4
Teacher of a subject	Teacher links	One teacher per student, year and subject	Score model
Seats	Enrollment extract	Earlier maximum enrollment up to the closure year	Lists' rules

Notes: A school year is labeled by the calendar year in which it ends. Census survey: the Puerto Rico Community Survey, by block group. Counts and means of each variable are in Tables 1 and 2.

Table A2. Samples: two flows of displaced children

<i>Panel A. Children displaced by closures of 2013–17</i>	
	Children
Children in a closed school's grades 1–7 in its last operating year	34,975
<i>Panel B. Pupil records of displaced children</i>	
	Records
With a next public school recorded	32,514
With the child's poverty record and the closed school's rural share	20,422
With a school effect for both schools	19,982
With surviving schools to average in the municipality	19,927
<i>Panel C. Children in the choice model</i>	
	Children
Children of the 885 markets with a realized closure	28,024

Notes: Closures of 2013–17. Panel A: all displaced children (Table 1); distance measures use them. Panel B: pupil records of displaced children, each condition applied to the line above; the last line is the sample of Table 4. Panel C: the children of the choice model's access measure (Table 6). The two flows are not nested: pupil records count a child once per record found the next year, children count every child of a closed school.

percent entered, outside the set of alternatives, the school most common among the closed school's other grades, and 11.9 percent another school outside it. The 15,025 come from 310 closed school-years in 74 municipalities. The choice model covers 28,024 of the 34,975 displaced children; Appendix D reconciles the two counts.

The footnote in Section 3.1 reconciles the counts of closed schools in print (334, 339 and 323). Nine of the 19 closures held in every list of Section 8 had displaced children, 938 in all, and all nine closed in 2015, 2016 or 2017.


A.3 Schools, Enrollment and Closures by Year

Figure A1 shows, by school year from 2012–13 to 2018–19, the number of schools that serve grades 1 to 8 and their enrollment in the Department's enrollment extract, with the closures of 2013–17 by last operating year. These series count schools in the enrollment extract, not all public schools, and enrollment once per school-year. They differ from the published totals of Section 2.1, which cover all grades and a longer period.

A.4 Distance

Km to the nearest surviving school (Section 3.3) multiplies the straight-line distance between the closed school and each surviving school by the median ratio of road to straight-line distance for the closed school's class: 1.52 for rural schools and 1.42 for urban schools. A school is rural for this adjustment when at least half of its nearby population lives in rural block groups. The ratios come from the 279,759 routes from block groups to schools that the choice model uses, computed with the Open Source Routing Machine (OSRM) on

Figure A1. Schools, enrollment, and closures by year



Notes: A and B: schools serving grades 1–8 and their enrollment in the administrative files. C: public schools closed, by last operating year, from the closure file; the paper studies the 334 closures of 2013–17. The bars count schools that served at least one grade with a standardized test (grades 3–8) in their last year; counting every school with a last operating year, whatever its grades, gives 8, 72, 59, 42 and 158. D: children displaced by the closures of each year, counting the 34,037 children of the 315 closures that Section 8 chooses again (the 938 children of the 19 held closures are left out). Counts of a census: no interval. Public schools fell from 1,509 in 2009–10 to 847 in 2019–20 and public enrollment from 473,735 in fall 2010 to 307,282 in fall 2018 (NCES, 2021b,a). Table 1 holds the counts of 2013–17.

the OpenStreetMap road network. The nearest school is the nearest surviving school of the closed school’s municipality that serves the child’s grade, and the share beyond 4.8 km uses the same construction.

Coordinates come from the school file for each year. Over the period, 129 schools were re-geocoded by more than 10 m, 116 of them first in 2017; 16 of the 129 are closed schools, all closed in 2017. Twenty-one schools moved by more than 1 km, and 19 of them never closed. The choice of coordinates barely moves distances from closed schools: the enrollment-weighted straight-line km from a closed school to the nearest open school of its municipality is 1.25 with the coordinates used here, 1.29 with each school’s first coordinates and 1.26 with its last, and 3 percent of closed schools move by more than 0.1 km.

The pupil records carry their own km between the school closed and the school entered. Section 4.4 uses it for the school entered and for the nearest surviving school on the same records, so that each comparison stays within one construction. The two constructions agree within 0.01 km on 1.0 percent of the 15,025

re-enrolled records, and the mean absolute difference between them is 0.47 km. I therefore never subtract a distance on one construction from a distance on the other, and I report no comparison of average test scores between the school entered and a nearest school found on the other construction.

Two other distances appear in the paper. Table 1 and Figure 2, Panel B measure km from a school to the nearest other school serving a common grade, on the same straight-line-times-ratio construction. In both, that school lies in the same municipality. Figure 2, Panel B takes, for each grade the school serves, the nearest other school of the municipality serving that grade, and weights these distances by the school's enrollment in each grade. The choice model measures distance from each block group to each school along the routes above. None of these distances measures fares, time on a school bus, or a child's trip from home.

A.5 The Groups and the Survey

The paper uses three constructions of “rural,” compared in Table A3. The first, which defines the four groups in both flows, is the rural share near a school: the share of the population around the school that lives in rural block groups, with each block group weighted by its population and by $\exp(-d/2)$ in its distance d in km from the school. A bandwidth of 2 km separates schools by the rural status of the children who attend them better than a narrower or a wider one (an area under the curve of 0.866, against 0.856 at 1 km and 0.792 at 5 km); because that target is built from the same block-group indicator, this is a check of the bandwidth, not an outside validation. The second is the block-group indicator itself: a block group is rural when its log population density is below -8.681 , a threshold set to match the Department's rural designation of schools (accuracy 0.788; area under the curve 0.849). The Census urban share cannot serve, since it equals 1.0 in the median block group. This indicator classifies the families of the choice model. The third is the Department's own designation of a school as rural or urban, which the exit estimates by origin use (Appendix A.6). Figure 2, Panel B calls a school rural when the rural share near it is at least one half.

Table A3. Three constructions of rural

	Definition	Check	Used for
Rural share near a school	Share of the population around the school that lives in rural block groups; block groups weighted by population and by $\exp(-d/2)$ in their km d from the school.	Area under the curve against the rural status of the children attending: 0.866 at a 2 km bandwidth, 0.856 at 1 km, 0.792 at 5 km; the target uses the block-group indicator.	The four groups in both flows; a share of at least one half also defines the school class for the road adjustment.
Block-group indicator	A block group is rural when its log population density is below -8.681 , the threshold that best matches the Department's designation of schools.	Accuracy 0.788; area under the curve 0.849; all 2,527 block groups of the choice model classified.	Families of the choice model, by block group of residence.
Department's designation	The Department's own designation of each school as rural or urban.	The target of the block-group indicator's threshold.	Exit estimates by the closed school's area.

Notes: The Census urban share cannot serve in Puerto Rico: it equals 1.0 in the median block group. Area under the curve: the probability that a school attended by rural children ranks above one attended by urban children.

The Department records the family-poverty flag only in years when a child has a test score. For the other

years I assign the child's most frequent status across the Department's enrollment records from 2009–10 to 2022–23. Among the 292,691 children with any record, 80.7 percent never change status, and the assigned status agrees with the recorded flag 96.8 percent of the time where both exist. In the enrollment flow, a school's displaced children are split by the school's share of children with the flag; in the pupil flow, each child keeps the child's own record. A school's share is never used in place of a child's record.

The survey's public-use microdata for 2013–17 describe children aged 8 to 14. Among them, 56.2 percent are income-poor and 76.5 percent attend a public school; among public-school children, 67.1 percent (0.4) are income-poor. Of income-poor children, 6.9 percent attend a private school and 1.8 percent are not enrolled, 8.7 percent outside the public schools in all. Of children who are not income-poor, 40.8 percent attend a private school and 1.8 percent are not enrolled, 42.6 percent in all. Non-enrollment is thus the same in both groups, and private schooling makes the difference. In the block-group counts, rural block groups hold 45.5 percent of the population, 47.3 percent of children enrolled in public schools in grades 1 to 8, and 38.2 percent of children enrolled in private schools.

A.6 Exit around a Closure

The exit estimate of Section 4.4 comes from an event study that follows children around the closure of their school and compares them with matched children whose schools did not close. Treated children attended a school whose last operating year fell in 2013–16, were enrolled the year before in a grade from 3 to 7 that was not their school's highest, and lived in the school's municipality. Event time zero is the first year after the closure and event time -1 is the reference; 7,138 treated children are observed at event time zero. Matching uses coarsened exact matching. The estimator interacts cohort with event time and aggregates as in Sun and Abraham (2021), with child-by-matching-cohort and matching-cohort-by-year effects, and standard errors are clustered by the Department's administrative districts. Identification requires that, given the matching and the fixed effects, treated children would have followed the comparison children's trend without the closure, and that the year before the closure was not already affected by it.

Two outcomes are mutually exclusive: the child does not appear in the public enrollment records the next year, or appears in a public school of another municipality. Their sum is leaving the municipality's public schools, which corresponds to the choice model's outside option. At event time zero the effects are 3.2 percentage points (0.9) on absence from the public records, 5.3 points (1.2) on enrollment in another municipality, and 8.5 points (1.5) on their sum. The estimated effect is smaller in later years: on leaving, it is 1.5 points (1.1) one year later and 0.9 points (1.5) two years later. By origin, the effect on leaving is 12.0 points (2.5) for children of schools the Department designates as rural, about two thirds of it enrollment in another municipality, and 3.2 points (1.2) for children of urban schools.

The leads at event times -4 to -2 are close to zero: 0.0 points (0.7), 0.5 (0.5) and 0.3 (0.4), with a joint p -value of 0.595. The treated children's raw rates of absence, enrollment in another municipality and their sum are 8.4, 6.9 and 15.3 percent, against 10.5, 6.8 and 17.3 percent in an earlier tabulation; without

the residence restriction they are 10.7, 8.9 and 19.6 percent, which reconciles absence but not the other two. The samples therefore differ, and I keep the resident sample because it matches the population of the choice model.

Other samples give effects of similar size: 7.6 points with the closures of 2017 added, 6.8 points with all grades, and 10.6 points without the residence restriction, which describes a different population. The effects for grades 3 and 4 and for grades 5 to 7 are 8.6 and 8.5 points. Applying the pooled effect to the choice model's grades 1 to 7 assumes that it also holds in grades 1 and 2.

Online Appendix

A Framework for Assessing the Distributional Effects of School Closures: Evidence from Puerto Rico

Reza Moradi

B Descriptive Evidence

This appendix adds three exhibits to Section 4: the predictors of closure, the distributions behind the group means of Figure 3, and closure and distance by area.

B.1 Predictors of Closure

Table B1 relates closure to the characteristics the Department held the year before, pooled over 2013–17. In the conditional logit within municipality and year, log enrollment carries the pattern, with a coefficient of -3.02 (0.22). Given enrollment, no other characteristic in the logit can be distinguished from zero: the coefficient is -0.09 (0.11) on log km to the nearest other school, -0.10 (0.34) on the rural share near the school, 0.37 (0.77) on the share of children with the family-poverty record, and -0.41 (0.49) on the trend in enrollment. The linear probability model with the same characteristics gives -0.143 (0.011) on log enrollment; in it, schools whose enrollment was growing were also less likely to close, with a coefficient of -0.083 (0.027) on the trend. The columns by year are a robustness check, and no statement in the paper rests on them.

B.2 The Distributions behind the Group Means

Figure B1 draws, for each group, the distribution across displaced children of the two quantities whose means Figure 3 and Table 2 report. Panel A is km from the closed school to the nearest surviving school. The rural distributions sit well to the right of the urban ones: for 81.6 percent of rural-poor and 81.0 percent of rural non-poor children the nearest surviving school was more than 1.6 km (one mile) away, against 42.0 and 51.0 percent of urban-poor and urban non-poor children. Panel B is the average test score of the nearest surviving school minus that of the school closed, on the 34,191 children with both. These distributions are wide and overlap across groups. Across rural-poor children the standard deviation is 0.634 SD, and the nearest surviving school scored below the school closed for 52.9 percent of rural-poor children against 41.5 percent of urban-poor children. Sampling noise, from the particular children with scores at each school, accounts for about 2.3 percent of the variance of this difference.

B.3 Enrollment and the Two Measures, by Area

Figure B2 draws closure and distance against enrollment separately for rural and urban schools, where a school is rural when the rural share near it is at least one half. Closure falls with enrollment in both areas: in

Table B1. Predictors of closure*Panel A. Conditional logit within municipality and year, log-odds*

	2013–17 pooled (1)	By last operating year			
		2014 (2)	2015 (3)	2016 (4)	2017 (5)
Log enrollment the year before	–3.02 (0.22)	–3.41 (0.45)	–3.56 (0.47)	–2.60 (0.80)	–2.78 (0.32)
Log km to the nearest other school	–0.09 (0.11)	–0.06 (0.20)	–0.20 (0.21)	0.30 (0.39)	–0.22 (0.17)
Rural share near the school	–0.10 (0.34)	–0.12 (0.55)	–0.66 (0.72)	0.07 (0.77)	0.13 (0.48)
Share with the family-poverty record	0.37 (0.77)	0.93 (1.62)	1.32 (1.44)	0.08 (2.63)	0.43 (0.78)
Closures	315	67	52	38	151
Schools in the choice sets	2,970	738	616	614	915


Panel B. Linear probability with the same fixed effects, share closed

	2013–17 pooled (6)	By last operating year			
		2014 (7)	2015 (8)	2016 (9)	2017 (10)
Log enrollment the year before	–0.143 (0.011)	–0.131 (0.017)	–0.109 (0.018)	–0.077 (0.018)	–0.274 (0.029)
Log km to the nearest other school	–0.001 (0.004)	–0.002 (0.010)	–0.002 (0.006)	0.011 (0.008)	–0.014 (0.015)
Rural share near the school	0.000 (0.014)	0.007 (0.030)	–0.025 (0.029)	0.011 (0.026)	0.011 (0.049)
Share with the family-poverty record	0.019 (0.029)	0.024 (0.070)	–0.003 (0.062)	0.019 (0.049)	0.049 (0.066)
Closures	315	67	52	38	151
Schools in the choice sets	5,930	1,257	1,190	1,130	1,087


Notes: Closure of a public school on its characteristics the year before, among the schools open at the start of the year, closures of 2013–17. Panel A: conditional logit within municipality and year, in log-odds per unit of the regressor; Panel B: linear probability with the same fixed effects. The pooled column and the years 2016 and 2017 also control for the change in log enrollment over the two years before, which the earlier years lack; indicators for a missing regressor are included and not shown. Standard errors in parentheses, clustered on the municipality. Likelihood-ratio test of the km row of the pooled logit: $p = 0.388$; 169 municipality-years. The 19 closures held in every list of Section 8 are not among these school-years (5,930 of the 5,949 of Table 1).

Figure B1. The two margins across displaced children

A. Distance to the nearest surviving school




B. Average test score: nearest surviving school minus school closed



Notes: Children displaced by the closures of 2013–17 (34,975, from 323 closed schools). A curve is a kernel density across a group's children. Panel A: distance from the closed school, not from home, to the nearest surviving school of the municipality serving the grade (the straight line times the median road ratio of the school's class). Panel B: that school's average test score minus the closed school's, each the year before the closure, in SD of pupil scores (34,191 children with an average score for both schools; closures of 2013 have none). Raw averages reflect a school's pupils as well as the school; sampling noise is 2.3 percent of the variance of the difference. Off the axes: 2 to 3 percent of children in Panel A, 2 to 4 in Panel B. The columns are alike: a closed school's poor and non-poor children share its distance and schools. Panel B's means differ less than its peaks; Table 2 holds means and standard errors.

the smallest of the 20 bins, 27.2 percent of rural and 21.7 percent of urban school-years closed, and almost none closed in the largest. Rural schools stand farther from their nearest other school at every size. In each of the 20 bins the rural mean exceeds the urban one, with rural means between 3.45 and 4.49 km and urban means between 1.57 and 2.62 km. Across all rural schools the mean is 3.89 km, against 1.93 km across urban schools, and the slopes on log enrollment are -0.050 (0.183) and 0.291 (0.149).

Figure B2. Enrollment, closure and the nearest other school: rural and urban schools



Notes: School-years of public schools, 2013–17, by enrollment the year before (5,909 in A, 5,861 in B). A point is the mean over a bin of school-years: the 20 bins of Figure 2, divided by area. Rural: schools whose rural share nearby is at least one half; urban: other schools. A: share of school-years that ended in a closure. B: km (a straight line times a road ratio) to the nearest other school serving a common grade; each line has its area's slope on log enrollment, with closure-year effects, through the area's mean: rural -0.050 (0.183), urban 0.291 (0.149), standard errors clustered by municipality.

C School and Teacher Effects

This appendix gives the details behind Sections 6.1, 7.1 and 7.2: what identifies each ingredient of the two models (Appendix C.1), the connected set (Appendix C.2), the variance components (Appendix C.3), the lagged-score check (Appendix C.4), the residuals before closure by period (Appendix C.5), the two versions of the estimates that address the threats of Section 6.1 (Appendix C.6 and C.7), and further results on the change in school effects (Appendix C.8 to C.10). The score model is estimated on the school years 2011–19; the closures studied are those of 2013–17.

C.1 What Identifies What

Table C1 lists each ingredient of the two models: the student effect, the teachers' value-added, the school effect, the cost per km by family type, the nest parameter, attachment to public schools and the responses to effectiveness. For each it gives the fact of Section 4 that calls for it, the variation that identifies it and the check that bears on it, with the place where each appears.

Table C1. What identifies what

	Why the model needs it	Identified by	Check
Student effect	An average test score mixes the children a school enrolls with what the school adds (Section 4.2).	A student's scores across years, teachers and schools.	Lagged scores in place of the student effect (Appendix C.4).
Teachers' value-added	Teachers were reassigned apart from children (Section 2.2).	A student taught by different teachers, the student effect held fixed; teachers who teach in more than one school.	Lagged-score specification; bounds on drift and on the teacher-school match (Table 3).
School effect	The gap in average test scores sits in the schools that closed (Section 4.2).	Students and teachers who change schools, within the connected set they link; closures add such moves (Table C2).	Residuals before closure (Figure 4); without the closure-decision years (Table C4); without displaced pupils (Table C5); two halves of the school-years (Table C3).
Cost per km, by family type	Rural children were left farther from a surviving school, and most re-enrolled children did not enter the nearest one (Section 4.3, Section 4.4).	Given the school constants, the shares of re-enrolled displaced children entering the second to sixth-or-farther surviving school, by group.	Km to the school entered, not targeted (Table 5); fit of the rank shares (Table D1).
Nest parameter	Children of closed schools left the municipality's public schools, more often from rural schools (Section 4.4).	Set so that the model's exit equals the exit effect of the event study; just identified.	Exit by the class of the closed school, not targeted (Table 5).
Attachment to public schools	Income-poor and other families leave public schools at different rates.	The share of income-poor children among public-school students in the survey.	The fit of that share (Table D1).
Responses to effectiveness, by family type	Poor and non-poor displaced children entered schools of similar effectiveness (Section 6.2).	Not separately: the rural difference cannot be told apart from the rural cost per km. Both differences are set to zero, a restriction chosen after the results were seen.	The richer specification and a test that both differences are zero (Table D2).

Notes: The two models of Section 5: student, teacher and school effects in test scores, and the choice model of where displaced children enroll. The table adds no number to the text.

C.2 The Connected Set

The estimation sample keeps the largest set of schools linked by a shared teacher and subject, and it requires at least two scored records for each student effect. I alternate the two restrictions until neither changes the sample, and I apply both again after removing observations with high estimated leverage. The computation absorbs the student effects by demeaning within student and solves the remaining normal equations with a small ridge, which selects the minimum-norm normalization; teacher-subject effects and school effects are each centered at an unweighted mean of zero. The ridge stabilizes nearly singular directions of the design. It cannot identify effects that no move links.

Coverage is nearly complete (Table C2). A school effect is estimated for 97.5 percent of the 1,281 schools of 2013–17 in the enrollment extract and, in both halves of the school-years, for 78.8 percent of them. Among the 334 closed schools the two shares are 97.3 and 72.8 percent. The schools with an estimate hold 98.2 percent of all children and 97.5 percent of the children of closed schools. On the two splits of the

school-years used in the paper, 1,018 and 1,037 of the 1,274 schools have an estimate in both halves; a third split gives 1,013.

Table C2. The estimation sample and the connected set

Panel A. Units identified

	Schools (1)	Teacher- subjects (2)
Estimation sample, full	1,274	19,484
Identified in both halves, first split	1,018	8,722
Identified in both halves, second split	1,037	8,912

Panel B. Schools of 2013–17 with a school effect

	Schools (3)	With an effect, % (4)	In both halves, % (5)	Children covered, % (6)
All schools	1,281	97.5	78.8	98.2
Closed	334	97.3	72.8	97.5
Never closed	947	97.6	81.0	98.3
Rural	529	97.7	76.0	98.6
Urban	741	97.3	81.2	98.0
Closed, rural	165	98.2	73.3	98.2
Closed, urban	159	96.2	73.0	96.7

Notes: Panel A: the model of test scores with pupil, teacher-subject and school effects, school years 2011–19 (3,452,029 student-subject-years of 395,030 pupils); a split assigns whole school-years to two halves and the model is estimated on each; a unit counts when its effect is identified in both. Panel B: schools of 2013–17; children covered are those of the grades with a standardized test whose school has an estimated effect. School effects are identified on the set of schools connected by teachers and pupils who move; the movers' graph has algebraic connectivity 0.0388, and 98.8 percent of the set stays connected when any one mover is removed.

C.3 Variance Components and the Two Splits

Write equation (1) in stacked form as $y = X\beta + e$, with $\beta = (a', \mu', \phi')'$, and let $H = (X'X)^\dagger X'$ under the minimum-norm normalization, $P = XH$ and $\hat{\beta} = Hy$. A variance component is a quadratic form $\theta_A = \beta' A \beta$, with its weights in A . Given the design and a mean-zero error,

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

With independent errors, $\hat{\sigma}_r^2 = y_r \hat{e}_r / (1 - P_{rr})$ estimates the error variance of record r (Kline et al., 2020). Here errors are not independent: $e_r = v_r + \zeta_{c(r)} + \eta_{j(r),t(r)}$, with a classroom-year shock and a school-year shock. For calibrated variances v_ζ and $v_{\eta,jt}$ of the two shocks, the corrected component is

$$\hat{\theta}_A = \hat{\beta}' A \hat{\beta} - \sum_r (B_A)_{rr} \hat{v}_{v,r} - v_\zeta \sum_c \mathbf{1}'_c B_A \mathbf{1}_c - \sum_{j,t} v_{\eta,jt} \mathbf{1}'_{jt} B_A \mathbf{1}_{jt}, \quad (\text{C.2})$$

where $\hat{v}_{v,r}$ is the record's variance net of the two shocks and $\mathbf{1}_c, \mathbf{1}_{jt}$ select a classroom-year and a school-year. Leverages and traces are approximated by random projections, so the calculation adapts the leave-out estimator; it does not delete every classroom or school-year cluster exactly.

The classroom-year variance is set at $v_\zeta = 0.10018$. The school-year variance falls with school size, $\log v_{\eta,jt} = -0.9143 - 0.6072 \log n_j$, where n_j is the school's mean number of scored students per year, and is bounded above and below; its reference value is 0.02218, from an empirical covariance adjusted for attenuation measured in simulated data. At these inputs the corrected SDs are 0.361 for school effects, 0.225 for teachers' value-added, 0.089 for school-average value-added and 0.393 for effectiveness. Across the calibrations of the noise inputs examined, the first three run from 0.34 to 0.38, from 0.14 to 0.28 and from 0 to 0.15; these are ranges over calibrations, not confidence intervals (Table C3).

The second measure uses two halves of the school-years. If $\hat{\phi}^A$ and $\hat{\phi}^B$ are estimated on separate halves, the school effect is common to both and their estimation errors are uncorrelated, then $\text{Cov}(\hat{\phi}^A, \hat{\phi}^B) = \text{Var}(\phi)$. Splitting by school-year keeps the same school-year shock out of both halves; serially correlated shocks would still bias the covariance. The two splits give SDs of 0.303 and 0.307 for school effects, 0.075 and 0.063 for school-average value-added, and 0.203 and 0.191 for teachers' value-added. The two measures differ in weights and in how they treat noise, so neither is a draw of the other. The covariance between school effects and school-average value-added is imprecise and changes sign across corrections, and I give it no sign.

C.4 Lagged-Score Check

The lagged-score specification replaces the student effects with controls for earlier scores:

$$Y_{ist} = X'_{ist} \gamma + \mu_{k(i,s,t),s}^L + \phi_{j(i,t)}^L + e_{ist}^L, \quad (\text{C.3})$$

where X_{ist} holds grade-interacted cubic terms in the previous year's mathematics, Spanish and English scores and the student's demographic characteristics. It uses grades 4 to 8 of the 2011–19 panel, since grade 3 has no earlier test score. Its assumption on assignment conditions on earlier scores instead of a permanent student effect (Chetty et al., 2014), and because earlier scores reflect earlier schooling, its effects need not estimate the same object as equation (1). The teacher-subject effects of the two specifications correlate at 0.849, without correction for noise.

C.5 Residuals before Closure, by Period

Figure C1 repeats the check of Figure 4 separately for the closures of 2013–16 and those of 2017. It is a robustness check on the pooled check of Section 7.1, and no statement in the paper rests on it.

Table C3. Variance components: the leave-out correction and the two splits

<i>Panel A. Standard deviation, leave-out correction</i>			
	School effects (1)	School-average value-added (2)	Teachers' value-added (3)
At the noise inputs used	0.361	0.089	0.225
Lowest across the noise inputs examined	0.34	0.00	0.14
Highest across the noise inputs examined	0.38	0.15	0.28
<i>Panel B. Standard deviation, covariance across two halves of the school-years</i>			
	School effects (4)	School-average value-added (5)	Teachers' value-added (6)
First split	0.303 (0.017)	0.075 (0.032)	0.203 (0.008)
Second split	0.307 (0.017)	0.063 (0.036)	0.191 (0.010)
<i>Panel C. Without displaced pupils' records, first split</i>			
	School effects (7)	School-average value-added (8)	Teachers' value-added (9)
Standard deviation, across the two halves	0.311	0.078	0.198


Notes: School effects and teachers' value-added of Table 3, SD of pupil scores. A: the leave-out correction for estimation noise at the classroom-year and school-year noise inputs used, the school-year variance falling with school size (log variance -0.9143 minus 0.6072 times log scored pupils); the lowest and highest across the inputs examined are a range, not a confidence interval. B: the covariance of the estimates on two halves of the school-years, two ways of splitting them; standard errors in parentheses. Effectiveness, the school effect plus school-average value-added: 0.393 after the correction. The variance of school effects is 2.6 times that of teachers' value-added.

C.6 Without the Closure-Decision Years

This version re-estimates the school effects separately for each wave of closures, leaving out the wave's last operating year and the year before. It applies the filters of the main estimates and a coverage floor: a comparison enters only if at least 85 percent of its displaced records keep an estimate at both schools. The version covers 300 of the 312 closed schools in the comparison and 19,308 pupil records, 96.9 percent of the 19,927 of Table 4; its school effects correlate with the full-sample ones at 0.89 to 0.91. The 2013 wave falls below the floor. Table C4 sets each figure beside the full-sample estimate on the same 19,308 records, with standard errors clustered by municipality.

Section 7.2 reports the contrast and the level for all displaced children; the rest of the table agrees. The rural-poor level moves from -0.071 (0.033) with the full-sample estimates to -0.130 (0.047) without the decision years, while the contrast with urban non-poor children moves only from -0.091 (0.036) to -0.107 (0.054). The contrast with all other displaced children, -0.086 (0.042), has a 95 percent interval from -0.168 to -0.004 and a wild bootstrap p -value over municipalities of 0.050. It sits in the closed-school

Figure C1. Checks of the school effects, by closure period



Notes: Schools closed in 2013–17, by last operating year c (92 to 163 schools a year for the closures of 2013–16; 144 to 153 for those of 2017). Dots: the mean over schools of the school-year's mean test-score residual less the school's own average over its years; lines: 95 percent intervals, standard errors clustered by municipality. Residuals are from the model of test scores with pupil, teacher-subject and school effects, full-sample estimates; they average zero within a school, so the figure shows timing, not a level. Joint test that the means at $c-2$, $c-1$ and c are zero: $p = 0.919$ for 2013–16 and $p = 0.012$ for 2017; against the never-closed schools of the municipality in the same year, $p = 0.834$ and $p = 0.040$. Figure 4 pools the periods. The score-model estimates are in Table 3.

term, 0.120 (0.040) against 0.109 (0.029) with the full-sample estimates, and not in the entered-school term, 0.034 (0.021) against 0.026 (0.016).

C.7 Without Displaced Pupils

This version leaves out displaced children's records from the year of their move on, so that their later test scores do not help estimate the school effect of either the school closed or the school entered. Sections 7.1 and 7.2 report its main figures, and Table C5 gives the rest. The variance components change little (Panel A). This version covers 14,525 records, compared with 19,927 in Table 4; its pooled standard errors hold group sizes fixed, as stated in Table C5. The contrast of -0.097 (0.041) between rural-poor and all other displaced children, with $p = 0.024$, sits in the closed-school term, 0.103 (0.041), against 0.007 (0.021) in the entered-school term (Panel B).


C.8 The Change in School Effects by Closure Year

Table C6 and Figure C2 give the change in school effects by group, and its two contrasts, for the closures of each year. They are a robustness check on the pooled estimates of Section 7.2, and no statement in the paper rests on them.

Table C4. The change in school effects without the closure-decision years

	Change, entered minus closed		School closed		School entered	
	Full sample (1)	Without decision years (2)	Full sample (3)	Without decision years (4)	Full sample (5)	Without decision years (6)
All displaced children	−0.021 (0.023)	−0.079 (0.033)	−0.019 (0.025)	0.023 (0.036)	−0.041 (0.015)	−0.056 (0.019)
Rural poor	−0.071 (0.033)	−0.130 (0.047)	0.046 (0.033)	0.094 (0.046)	−0.025 (0.020)	−0.036 (0.027)
Rural non-poor	−0.023 (0.035)	−0.067 (0.054)	−0.007 (0.038)	0.020 (0.054)	−0.030 (0.026)	−0.047 (0.036)
Urban poor	0.018 (0.025)	−0.043 (0.035)	−0.072 (0.030)	−0.030 (0.041)	−0.054 (0.016)	−0.073 (0.018)
Urban non-poor	0.021 (0.025)	−0.023 (0.041)	−0.079 (0.028)	−0.053 (0.047)	−0.059 (0.016)	−0.077 (0.018)
Rural poor minus all other	−0.083 (0.029)	−0.086 (0.042)	0.109 (0.029)	0.120 (0.040)	0.026 (0.016)	0.034 (0.021)
Rural poor minus urban non-poor	−0.091 (0.036)	−0.107 (0.054)	0.125 (0.035)	0.148 (0.052)	0.034 (0.021)	0.041 (0.027)
Pupil records	19,308					

Notes: Pupil records of children displaced by closures of 2013–17 whose closed and entered schools have both estimates: 19,308 records, 96.9 percent of the weight of Table 4, from 300 of 312 closed schools. Without the decision years: the model re-estimated leaving out, for every school, the two school years in which that wave's closures were decided; both versions use the same records. Each school is measured relative to the enrollment-weighted mean of the municipality's surviving schools. Standard errors in parentheses, clustered on the closed school's municipality.

Figure C2. The change in school effects by closure year: rural-poor children and the contrast

Notes: Displaced children with a recorded next school, by the closed school's last operating year; the bottom row pools the years. Full-sample estimate of the school effect of the school entered minus that of the school closed, test-score SD; negative is a loss. A: rural-poor children. B: rural-poor minus urban non-poor children. Dashed: zero. Lines: 95 percent intervals, standard errors clustered on the municipality of the closed school. Rural-poor children behind each year's estimate, as sums of group shares: 139, 1,834, 1,181, 1,067, 3,870 (2013–17). Table C6 holds the numbers.

Table C5. School effects estimated without displaced pupils' records

<i>Panel A. Standard deviation and reliability</i>							
	School effects (1)	School-average value-added (2)	Teachers' value-added (3)				
SD across halves, full sample	0.303	0.075	0.203				
SD across halves, without displaced pupils	0.311	0.078	0.198				
Reliability, full sample (first split)	0.565	0.087	0.29				
Reliability, without displaced pupils	0.549	0.084	0.26				
<i>Panel B. The change in school effects, without displaced pupils' records, SD</i>							
	All (4)	Rural poor (5)	Rural non-poor (6)	Urban poor (7)	Urban non-poor (8)	Rural poor minus all other (9)	Rural poor minus urban non-poor (10)
School closed	−0.006 (0.034)	0.055 (0.046)	0.026 (0.054)	−0.066 (0.038)	−0.032 (0.047)	0.103 (0.041)	0.087 (0.056)
School entered	−0.066 (0.019)	−0.062 (0.024)	−0.055 (0.031)	−0.072 (0.023)	−0.065 (0.025)	0.007 (0.021)	0.003 (0.028)
Change, entered minus closed	−0.059 (0.033)	−0.117 (0.045)	−0.081 (0.044)	−0.006 (0.036)	−0.033 (0.055)	−0.097 (0.041)	−0.084 (0.063)

Notes: Panel A: as Table 3, with the model estimated on halves that leave out displaced pupils' records from the year of the move on. Panel B: the pupil records of Table 4 (14,525 pupil records with both schools' estimates, 14,523 in column (10)), each school relative to the enrollment-weighted mean of the municipality's surviving schools. Standard errors in parentheses, clustered on the closed school's municipality; in the pooled columns they hold the groups' sizes fixed.

Table C6. The change in school effects by closure year

	Displaced children				Difference		Pupil records (7)
	Rural poor (1)	Rural non-poor (2)	Urban poor (3)	Urban non-poor (4)	Rural poor minus all other (5)	Rural poor minus urban non-poor (6)	
2013	−0.099 (0.178)	0.028 (0.078)	−0.015 (0.131)	0.170 (0.090)	−0.147 (0.174)	−0.269 (0.187)	291
2014	−0.047 (0.069)	−0.013 (0.062)	−0.063 (0.090)	−0.024 (0.066)	0.006 (0.087)	−0.023 (0.077)	4,254
2015	−0.098 (0.089)	−0.007 (0.091)	0.179 (0.049)	0.199 (0.063)	−0.251 (0.088)	−0.297 (0.103)	3,181
2016	0.029 (0.066)	0.017 (0.074)	0.108 (0.061)	0.072 (0.053)	−0.061 (0.054)	−0.044 (0.063)	2,323
2017	−0.117 (0.048)	−0.060 (0.052)	−0.034 (0.033)	−0.042 (0.032)	−0.078 (0.045)	−0.075 (0.053)	9,933
2013–17, pooled	−0.079 (0.036)	−0.032 (0.035)	0.011 (0.027)	0.018 (0.025)	−0.084 (0.031)	−0.097 (0.037)	19,927

Notes: The change in school effects, school entered minus school closed (full-sample estimates), by the closed school's last operating year. Each school is measured relative to the enrollment-weighted mean of the municipality's surviving schools. Pupil records: the year's records of displaced children with both schools' estimates. Standard errors in parentheses, clustered on the closed school's municipality. Year rows count the 19,982 records with both schools' estimates; the pooled row is Table 4's 19,927.


C.9 The Spread across Children, Net of Noise

The group means of Table 4 sit inside a wide spread across children (Table C7). Net of estimation noise, the SD of the change in school effects across displaced children is 0.42, with a 95 percent interval from 0.33 to 0.52. Relative to the average surviving school of the municipality, the SD of the school entered is 0.26 [0.19, 0.32] and that of the school closed 0.33 [0.21, 0.42]. The calculation keeps the schools that meet a coverage floor of 70 percent.

C.10 School Entered against School Closed

Figure C3 plots the entered-school term against the closed-school term for each closed school and wave, with the two terms estimated on separate halves of the school-years so that they share no estimation noise. Across 233 of the 312 closed school-waves, the slope is -0.04 , with a 95 percent interval from -0.41 to 0.17 . The school entered, relative to the municipality's average surviving school, does not rise detectably with the school closed.

Figure C3. The school entered against the school closed, split-sample bins



Notes: Closed schools of 2013–17 with estimates on both halves of the school-years (233 school-waves, 75 percent). School effects, in test-score SD, relative to the enrollment-weighted average surviving school of the municipality. Closed schools are put in five bins by their effect on one half of the school-years; each circle is the bin's mean for the school closed and for the school its children entered on the other half, the two directions averaged. Lines: 95 percent intervals of the school entered, a bootstrap over municipalities. Dashed: the 45-degree line, no change. Split-sample slope of the school entered on the school closed: all closed schools -0.04 [-0.41 , 0.17], rural -0.02 [-0.50 , 0.34], urban -0.10 [-0.68 , 0.34]. The comparison holds only across halves and only if closures were not chosen on the scores behind the estimates.

Table C7. The spread of school effects across displaced children, net of estimation noise

<i>Panel A. Change, school entered minus school closed</i>			
	Raw SD (1)	Net SD (2)	Coverage, percent (3)
All displaced children	0.45	0.42 [0.33, 0.52]	73
Rural poor	0.50	0.45 [0.31, 0.58]	73
Rural non-poor	0.45	0.44 [0.32, 0.56]	71
Urban poor	0.41	0.38 [0.30, 0.48]	74
Urban non-poor	0.42	0.42 [0.31, 0.54]	74
Children	14,538		
<i>Panel B. School closed, relative to the average surviving school</i>			
	Raw SD (4)	Net SD (5)	Coverage, percent (6)
All displaced children	0.40	0.33 [0.21, 0.42]	80
Rural poor	0.43	0.35 [0.18, 0.46]	79
Rural non-poor	0.43	0.38 [0.22, 0.51]	77
Urban poor	0.35	0.28 [0.20, 0.37]	81
Urban non-poor	0.34	0.31 [0.22, 0.40]	80
Children	15,855		
<i>Panel C. School entered, relative to the same average</i>			
	Raw SD (7)	Net SD (8)	Coverage, percent (9)
All displaced children	0.32	0.26 [0.19, 0.32]	89
Rural poor	0.34	0.27 [0.17, 0.34]	90
Rural non-poor	0.33	0.24 [0.15, 0.31]	90
Urban poor	0.31	0.26 [0.19, 0.34]	89
Urban non-poor	0.28	0.24 [0.14, 0.31]	90
Children	17,788		

Notes: Children displaced by closures of 2013–17 with a recorded next school, in the weights of Table 4. Raw SD: across children, of the full-sample estimate. Net SD: the square root of the covariance between two estimates of the same quantity, each from one half of the school-years with displaced pupils left out. Coverage: percent of the group's weight with both halves for every school involved. In brackets: 95 percent intervals from a bootstrap over municipalities.

D The Choice Model

This appendix gives the details behind Sections 6.2 and 7.3: the moments, estimator and fit of the choice model (Appendix D.1 and D.2), the richer specification with the responses to effectiveness (Appendix D.3), the check of the model's placements on the realized closures (Appendix D.4), the access loss under other assumptions (Appendix D.5), and the populations behind the model's numbers (Appendix D.6). The model is equation (2); its parameters are $(\tilde{\lambda}_0, \tilde{\lambda}_R, \tilde{\lambda}_P, \pi, \rho)$, and the school constants are recovered from enrollment shares at each candidate value.

D.1 Moments, Estimation and Fit

The rank moments describe where the children of a closed school went. For a family of type z in market m , removing the closed school o alone gives

$$K_{moj}(z; \theta) = s_{mj}(z; \mathcal{J}_m \setminus \{o\}, \theta) - s_{mj}(z; \mathcal{J}_m, \theta), \quad (\text{D.1})$$

which, under Assumption 2, is the probability that the family attended o and, once o closes, enters j . Let $b(o, j)$ be the distance rank of the surviving school j from o , and $\mathcal{B}_{mo}$ the surviving schools eligible as destinations, which exclude other schools closing in the same market. For group g , with poverty status p_g , let $\gamma_g(o)$ be the rural share of the closed school's surroundings, $A^R(o)$, for the rural groups and $1 - A^R(o)$ for the urban groups, and let $a_g(o, z) = \gamma_g(o) \mathbf{1}\{P = p_g\}$. With origin weights $\omega_{mo} = n_{mo}/S_{mo}$, the ratio of the origin's displaced count to its enrollment share, the model's share of group g entering a school of rank b is

$$M_{gb}(\theta) = \frac{\sum_{m,o} \omega_{mo} \int a_g(o, z) \sum_{j \in \mathcal{B}_{mo}} \mathbf{1}\{b(o, j) = b\} K_{moj}(z; \theta) dF_m(z)}{\sum_{m,o} \omega_{mo} \int a_g(o, z) \sum_{j \in \mathcal{B}_{mo}} K_{moj}(z; \theta) dF_m(z)}. \quad (\text{D.2})$$

Attendance at the closed school is itself a choice: in equation (D.1) a family's weight among a closed school's children is its probability of having attended that school, so each closed school's families are not simply the residents of its block groups. The ranks are 2, 3, 4 to 5, and 6 or more, in each of the four groups; the share at rank one is implied. The recorded shares apply the same restrictions to the re-enrolled children. The rank moments remove one closed school at a time, while the exit moment and the access loss remove all of a market's closed schools together. In markets with several closures the two need not give the same destinations, and the rank moments maintain the single-closure approximation.

The composition moment is the income-poor share of public-school students,

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

with M_m the market's population. Its target, from the survey's microdata for 2013–17, is 0.6706 (0.0036). Income poverty here differs from the family-poverty record that defines the groups.

The exit moment sets the nest parameter. Let C_m be the closing schools of market m and ω_m^E its weight in the exit sample. The model's exit per child of a closing school is

$$A(\theta) = \frac{\sum_m \omega_m^E \int [s_{m0}(z; \mathcal{J}_m \setminus C_m) - s_{m0}(z; \mathcal{J}_m)] dF_m(z)}{\sum_m \omega_m^E \int \sum_{o \in C_m} s_{mo}(z; \mathcal{J}_m) dF_m(z)}. \quad (\text{D.4})$$

Its target $\hat{\tau}$ is the exit effect in the year after a closure from the matched event study of Appendix A.6: 8.5 percentage points (1.5), of which 3.2 (0.9) come from leaving the public records and 5.3 (1.2) from entering another municipality's public school. The target is an excess-exit effect, not the exit rate. It is an estimating input to the choice model, distinct from the broader enrollment and achievement results of the companion study (Bobonis et al., 2026). A market whose public schools all close has a model exit of one whatever the parameters, so it carries no information on ρ .

Given ρ , the four other parameters minimize a weighted distance between the 17 moments and their targets. The rank discrepancies are weighted by the inverse of their bootstrap variances, clustered by municipality, and the composition discrepancy by 100 times its inverse variance. The nest parameter then solves $A(\hat{\vartheta}(\rho), \rho) = \hat{\tau}$, with the four parameters re-estimated at each candidate ρ . Writing $g(\vartheta, \rho)$ for the 17 discrepancies and G_ϑ for their derivative, the standard errors come from the stacked equations

$$h(\theta) = \begin{pmatrix} G'_\vartheta W g(\vartheta, \rho) \\ A(\vartheta, \rho) - \hat{\tau} \end{pmatrix}, \quad (\text{D.5})$$

with the full derivative $\partial h / \partial \theta'$, which keeps the terms that multiply nonzero discrepancies (Newey, 1984; Hall and Inoue, 2003). The calculation sets the covariance between the rank targets, the composition target and the exit target to zero; the rank and event-study samples are related, so this is an approximation. The population weights, the market sizes, the distances, the distance power and the weighting matrix are held fixed. The nest parameter is 0.847 (0.031), and inverting the exit equation gives a confidence set from 0.783 to 0.904.

Table D1 sets the 16 rank shares beside the model's; Section 7.3 reads the fit. The criterion is 26.90 for 17 moments and four parameters given ρ . The first-order test of the over-identifying restrictions ($p = 0.038$) treats the covariance of the targets as known, which makes the p -value too small. A one-percentage-point change in any second-nearest share moves the ratio of the rural-poor to the urban non-poor access loss by less than 0.01 to first order.

D.2 Rank of the School Entered

Figure D1 compares, by group, the recorded distance rank of the school a re-enrolled child entered with the rank the model predicts. It shows the pattern of the moments: in every group the model puts too many children in the nearest surviving school and too few in the second-nearest.

Children of a closed school concentrate at one receiving school more than the model's probabilities

Table D1. Rank of the school entered among the surviving schools: data and model

	Rural poor		Rural non-poor		Urban poor		Urban non-poor	
	Data (1)	Model (2)	Data (3)	Model (4)	Data (5)	Model (6)	Data (7)	Model (8)
Nearest surviving school	44.6 (2.4)	49.7	37.2 (3.7)	46.9	38.7 (2.9)	42.2	37.9 (3.2)	40.4
Targeted								
Second nearest	26.9 (2.1)	21.4	28.4 (3.4)	20.5	21.8 (1.6)	19.5	22.1 (2.1)	19.1
Third	9.8 (1.0)	11.1	9.2 (1.5)	12.0	12.5 (1.6)	10.8	11.7 (1.5)	11.2
Fourth or fifth	11.7 (1.3)	10.9	16.3 (2.3)	12.3	12.1 (1.5)	12.3	11.5 (2.4)	12.7
Sixth or farther	7.0 (0.9)	6.9	8.8 (1.6)	8.3	14.9 (1.9)	15.2	16.7 (2.8)	16.5

Notes: Percent of re-enrolled displaced children whose school entered has each rank in distance from the closed school among the municipality's surviving schools serving the grade. Data: recorded shares, standard errors in parentheses, clustered by municipality, those of the estimation's targets (for the nearest school, Tables 2 and 5 print the pupil records' own, at most 0.1 point apart); Model: the choice model's prediction at the estimates. The sixteen shares below the first row are targeted moments and the first row is their complement; a composition moment completes the seventeen. First-order test of the over-identifying restrictions, target covariance treated as known: $p = 0.038$. Groups by the closed school's rural share and the child's poverty record.


imply. For each closed school, grade and year, I take the surviving school that the other grades of that closed school in the same grade span enter most in excess of the model. Among the 14,034 displaced children who re-enrolled inside the closed school's set of alternatives and for whom that school is defined (62 percent of the displaced children who re-enrolled in the choice model's extract), 46 percent entered it; the model's probability is 21 percent, and the excess appears in all four groups. Assignment to a receiving school would produce this pattern. So would bus routes or siblings in other grades, an error in where the model places families, or an error in the cost per km; the records do not separate them. Re-estimating with that school removed from both the recorded and the model's destinations, with the nest parameter held, leaves 57 percent of the sample and gives a ratio of the rural-poor to the urban non-poor access loss of 2.04, against 1.90 on the same 805 markets (Appendix D.5). This concentration diagnostic and the masked re-estimation were added after the model's misfit was seen.

D.3 Variants and the Responses to Effectiveness

Write the within-nest response to effectiveness as $\beta_0 + b_R R_i + b_P P_i$: a response β_0 common to all families, and rural and poor differences b_R and b_P . Table D2 sets the coefficients and the access ratio of the specification of Table 5 beside two variants that free these responses.

The common response β_0 is not a separate parameter. It is part of the school constants δ_{mj} , which are recovered from enrollment shares and also hold every other attribute that all families value alike. A common coefficient on effectiveness could be separated from the constants only by a further restriction on how effectiveness moves with those other attributes. A projection of the constants on effectiveness, or on

Figure D1. Distance rank of the school entered: recorded and predicted



Notes: The 13,058 records of children displaced by closures of 2013–17 who re-enrolled in a surviving public school of their municipality and have a value in the choice model. Rows: the rank of the school entered among the municipality’s surviving schools serving the grade, by distance from the closed school. Filled: the recorded share; hollow: the choice model’s predicted share for the same records; both weighted by records, so they differ from the targets of Table D1, which weight by children. Lines: 95 percent intervals, standard errors clustered by municipality. Circles: poor children; squares: non-poor children.

school-average value-added, would mix the response with omitted attributes and with noise in the regressor, and that noise is large: the reliability of school-average value-added is 0.087 (Table 3). In the school-year split, the lower confidence limit for the within-market variance of school-average value-added net of noise is not positive, so the bound on a common coefficient that this variance would give is unbounded. The table therefore reports no estimate of β_0 .

The two differences are estimated, and neither is informative. Estimated in a second step, given the first step’s coefficients, the poor difference is $b_P = -0.063$ (0.398), with a 95 percent interval from -0.84 to 0.72 . The rural difference is $b_R = -3.85$ (4.52), and its profile confidence set does not close at the lower end of the searched range, -10 , and has an upper end of 3.0 . A test that both differences are zero gives a score of 1.06 and $p = 0.59$. Estimated jointly with the coefficients of the cost per km and attachment, with b_R held at zero, the poor difference is -0.061 (0.345), with an interval from -0.74 to 0.62 . Iterating the two steps to a common solution gives $b_R = -4.41$ and $b_P = -0.06$; the curvature check fails at that solution, so it carries no interval.

Four facts support the restriction $b_R = b_P = 0$, which was chosen after the results were seen. First, the rural difference cannot be distinguished from the rural difference in the cost per km, so its sign is not identified. Second, the poor difference is small and imprecise, and among displaced children the effectiveness of the school entered barely differs between poor and non-poor children (Section 6.2). Third, the test that both differences are zero does not reject. Fourth, under the restriction no estimate of a school effect enters the access numbers, so the noise in the school effects (Section 6.1) does not reach them. The restriction matters little for access (Appendix D.5). Appendix E.2 reports where the model places children when the rural difference is not zero.

Table D2. The choice model under three treatments of the responses to effectiveness

	Response held		
	Restricted (estimate)	at a first-pass value	Both steps iterated
	(1)	(2)	(3)
Coefficients			
Distance, urban non-poor family	−2.785 (0.577)	−2.697 (0.505)	−3.053 (1.454)
Distance, rural family, difference	1.219 (0.637)	1.146 (0.569)	1.251 (1.835)
Distance, poor family, difference	−0.202 (0.094)	−0.201 (0.093)	−0.206 (0.089)
Nest parameter	0.847 (0.031)	0.846 (0.031)	0.859 (0.034)
Income-poor family: attachment	1.411 (0.062)	1.409 (0.061)	1.413 (0.074)
Implied			
Rural over urban cost of distance (non-poor)	0.56 (0.15)	0.58 (0.15)	0.59 (0.42)
Access-loss ratio, rural poor over urban non-poor	1.84 (0.16)	1.81	1.78 (0.35)

Notes: Nested logit of the school a displaced child entered (Table 5). (1): responses to a school’s effectiveness restricted not to differ by family type (the estimates of the paper). (2): the responses held at the first pass of the two-step estimation. (3): both steps iterated to a fixed point, where the rural and poor responses are −4.41 and −0.06. Standard errors in parentheses, first-order; in (3) the two first-order formulas differ by more than a quarter for the distance coefficients and the ratios, and no interval is formed; the ratio in (2) is a point. Access-loss ratio: the 28,024 displaced children of Table 6.

D.4 Checks before Placing Children

Table D3 reports in full the check of Section 6.2. The record-level error is the school effect, or the effectiveness, of the school a child entered in the records minus that of the school the model places the child in, averaged over the model’s probabilities. Both use cross-fitted estimates: a school’s effect in the year compared comes from the half of the school-years that does not contain that school-year. The error is measured on 13,058 re-enrolled records in 73 municipalities and regressed, without weights, on the closed school’s size band, its km to the nearest surviving school and the log number of surviving schools. For each list of Section 8, the predicted change in the error is the difference between that list’s mean features and the realized closures’, per displaced child, times the coefficients. Its standard error is the larger of a municipality-clustered one and a municipality bootstrap one, with the choice-model estimates held fixed.

The largest predicted change across the five lists, for all displaced children and for rural-poor children, is 0.007 SD. The thresholds were set after the model’s error on the realized closures was known: with a largest change of at most 0.03 SD, differences between lists in where the model places children are read against a tolerance, and above 0.06 SD none would be read. Each list’s tolerance is the larger of 0.02 SD and the largest predicted change plus two standard errors, over the features of the check and over the same features with the closed school’s average test score and enrollment added. For rural-poor children’s school effects the tolerances run from 0.020 to 0.033 SD; for effectiveness they are 0.020 SD. The features do not jointly

predict the error ($p = 0.70$ for school effects and 0.89 for effectiveness). The table also reports the error itself, for all displaced children and for rural-poor children: in effectiveness it is -0.022 (0.015) and -0.031 (0.018). Two mechanical rules that place each child at the nearest surviving school or by the recorded rank shares give errors in effectiveness of -0.053 (0.027) and -0.025 (0.014) for all displaced children; they are comparisons and play no part in the check.

Table D3. Check of the choice model's placements on the realized closures

<i>Panel A. The model's error on the realized closures, SD</i>				
	School effects		Effectiveness	
	All (1)	Rural poor (2)	All (3)	Rural poor (4)
School entered, recorded minus model	-0.026 (0.020)	-0.041 (0.024)	-0.022 (0.015)	-0.031 (0.018)
<i>Panel B. Predicted change in the error under each list, SD</i>				
	School effects		Effectiveness	
	All (5)	Rural poor (6)	All (7)	Rural poor (8)
Smallest first	-0.002 (0.003)	-0.003 (0.005)	0.002 (0.002)	0.002 (0.003)
Smallest first, alternative within 4.8 km	0.002 (0.002)	0.003 (0.003)	0.000 (0.001)	-0.001 (0.002)
Lowest average test scores first †	-0.002 (0.003)	-0.002 (0.004)	0.000 (0.002)	0.000 (0.002)
Lowest scores first, within 4.8 km †	0.001 (0.003)	0.003 (0.003)	-0.001 (0.002)	-0.001 (0.002)
Nearest alternative first	0.006 (0.008)	0.007 (0.009)	-0.002 (0.005)	-0.001 (0.006)
<i>Panel C. Tolerance for the model's error, by list, SD</i>				
	School effects		Effectiveness	
	All (9)	Rural poor (10)	All (11)	Rural poor (12)
Smallest first	0.026	0.026	0.020	0.020
Smallest first, alternative within 4.8 km	0.027	0.027	0.020	0.020
Lowest average test scores first †	0.027	0.027	0.020	0.020
Lowest scores first, within 4.8 km †	0.020	0.020	0.020	0.020
Nearest alternative first	0.027	0.033	0.020	0.020

Notes: The 13,058 re-enrolled records of children displaced by closures of 2013–17 with a model value. Panel A: the effect of the school recorded, less that of the school the model places a child in. Panel B: the model's error is regressed on the closed school's size band, its km to the nearest surviving school and the number of surviving schools; applied to each list's closures, the features predict the change shown. The largest change in absolute value is 0.007 SD, under the threshold of 0.03 set for this check, which was set after the model's error on the realized closures was known. Panel C: the larger of 0.02 and the predicted change plus two standard errors over three versions of the regression. Standard errors in parentheses, the larger of clustered and bootstrap, by municipality. †: as in Table 7.

D.5 Access Loss under Other Assumptions

Table D4 repeats the access loss of Table 6 under other assumptions, on the same 28,024 children. Across the specification of Table 5 and four variants (the joint estimation with b_R held at zero or at either end of its interval, and the two steps iterated), the ratio of the rural-poor to the urban non-poor loss runs from 1.78 to 1.85. Holding the first step fixed, a separate sensitivity exercise evaluates eight combinations of the response differences, using the assumption-dependent rural range $[-4.34, 2.35]$ and the recorded endpoints for the poor response. The ratio runs from 1.74 to 1.84 at these points; the rural sensitivity range is not the profile confidence set above. Over the 15 admissible points of the confidence sets for the rural difference in the cost per km, it runs from 1.4 to 2.1, and the difference between the two groups' losses is at least 0.47 km-equivalents. The ratio is above one at every admissible point evaluated. These are ranges over evaluated points, not confidence intervals. With the two steps iterated the ratio is 1.78, but the curvature check fails there, so that point carries no interval.

Counting only the 26,210 displaced children in the 805 markets where displaced children's next school is recorded, the ratio is 1.90 (0.18); 1.83 to 1.91 across the variants, 1.79 to 1.90 at the ends of the intervals, and 1.4 to 2.2 over the confidence sets. The two counts agree exactly on the markets both include; the other 80 markets raise urban children's loss per child. On these 805 markets the ratio in units of within-nest utility is 1.18 at the estimates and runs from 0.78 to 1.71 over the points of the confidence set for the rural difference evaluated on this base, so the ordering in utility is not determined, while the ratio in km-equivalents stays above one. Each displaced family's loss is at least one unit of within-nest utility whatever the map; at the estimates the closed schools' lower school constants offset most of the part that does not depend on distance, so the loss in km-equivalents is mostly its distance part: 1.86 km-equivalents for rural-poor children against a total of 1.83, and 0.90 against 0.96 for urban non-poor children, a ratio of 2.05 in the distance parts.

The standard errors of Table 6 come from the delta method in the five coefficients. Along a direction of the coefficients' covariance, the ratio's response at one and at 1.96 standard errors stays within the tolerance of its linear approximation, so the local calculation is adequate at the main estimates.

D.6 Populations

Three populations carry the numbers of Sections 7 and 8, and Table D5 describes their counts. The 34,975 displaced children of Table 1 include the 938 children of the 19 closures held in every list of Section 8. The other 34,037, of whom 13,618 are rural poor and 20,419 are other children as fractional sums of the group shares, are the children of the 315 closures chosen again, and the change in school effects before children re-sort counts them. The choice model's 28,024 displaced children are its families in the 885 markets in which a closed school was among the choices. The difference of 6,013 between 34,037 and 28,024 splits into 781 children in school-grade pairs with no market in the choice model and 5,232 by which the enrollment records exceed the enrollment the model uses.

Rural-poor families are 50.0 percent of the model's displaced children (14,026 of 28,024), against 39.7

Table D4. Access loss under other assumptions*Panel A. As estimated*

	Km-equivalents per displaced child				Rural poor and urban non-poor	
	Urban non-poor (1)	Urban poor (2)	Rural non-poor (3)	Rural poor (4)	Ratio (5)	Difference (6)
28,024 children of 885 markets	1.00 (0.09)	0.97 (0.08)	1.91 (0.15)	1.84 (0.11)	1.84 (0.16)	0.84 (0.12)
26,210 children of 805 markets	0.96 (0.09)	0.93 (0.08)	1.90 (0.15)	1.83 (0.11)	1.90 (0.18)	0.87 (0.12)

Panel B. The two steps of the choice model iterated, 28,024 children

	Km-equivalents per displaced child				Rural poor and urban non-poor	
	Urban non-poor (7)	Urban poor (8)	Rural non-poor (9)	Rural poor (10)	Ratio (11)	Difference (12)
At the joint root, point only	0.98	0.95	1.80	1.74	1.78	0.76

Panel C. Ranges over evaluated points

	Ratio, 28,024		Difference, 28,024		Ratio, 26,210	
	Lowest (13)	Highest (14)	Lowest (15)	Highest (16)	Lowest (17)	Highest (18)
The specification and four variants (5 points)	1.78	1.85	0.76	0.84	1.83	1.91
Ends of the intervals of the responses (8 points)	1.74	1.84	0.78	0.84	1.79	1.90
Confidence sets for the rural cost per km (15 points)	1.42	2.12	0.47	0.99	1.44	2.22

Notes: Closures of 2013–17; the access loss per displaced child of Table 6, by group, standard errors in parentheses by the delta method in the five coefficients. The 805 markets are those where displaced children's next school is recorded; the two counts agree on the markets both include. B: the joint estimation's curvature check fails at its root, so that point carries no standard error. C: the variants hold the rural difference in the response to effectiveness at zero or at either end of its interval; the ends set both differences in the response at the ends of their intervals with the first step fixed. Ranges over evaluated points, not confidence intervals. Ratio and difference: rural poor relative to urban non-poor.

percent of the 34,975 in Table 1. The model types families by their block group's rural indicator and the family-poverty record, while Table 1 types children by the closed school's rural and poor shares (Appendix A); the difference between the two shares is not decomposed here.

Table D5. Three counts of displaced children

	Rural poor (1)	Rural non-poor (2)	Urban poor (3)	Urban non-poor (4)	All (5)
Children displaced, all closures	13,900	3,051	15,379	2,646	34,975
Children of the 315 list closures	13,618	2,964	14,910	2,545	34,037
Children in the choice model's markets	14,026	2,510	9,882	1,606	28,024

Notes: Closures of 2013–17. The first two rows differ by the 938 children of the 19 closures held in every list. The last two differ by 6,013 children: 781 in school-grade pairs of a closed school with no market or school in the choice model, and 5,232 in pairs where the enrollment records exceed the enrollment of the choice model's school data. Groups: rows 1 and 2 by the closed school's rural share and the family-poverty record; row 3 by the block group's rural indicator and the family-poverty record.

E Counterfactuals: Definitions and Further Results

This appendix supports Section 8. It defines the lists and what every comparison holds fixed (Appendix E.1), reports further results at the Department's counts (Appendix E.2), for closing fewer schools (Appendix E.3) and with the closures held fixed (Appendix E.4), and sets out how the standard errors and the statements of direction are computed (Appendix E.5).

E.1 Definitions, What Is Held Fixed, Seats and Coverage

Appendix D.6 describes the three populations of the comparisons. Distances use the 34,975 displaced children, including the 938 children of the held closures; the change in school effects before children re-sort uses the 34,037 children of the 315 closures chosen again; and the model-weighted access calculation represents 28,024 displaced children. Placed school effects instead use enrollment-weighted origins, with a covered weight of 32,398 under the realized closures. The account of savings uses each closed school's total enrollment in its last operating year: 46,348 pupils in the 315 closures, or 147.1 per closure.

Table E1 defines each list. Smallest first closes, in each municipality, year and size band, the smallest comparable schools. Smallest first with an alternative within 4.8 km does the same among the schools that leave another school serving their grades open within 4.8 km. Where too few schools qualify, it completes the count from the other schools, so the condition is a priority, not a guarantee: under the lists that use it, about a tenth of displaced rural-poor children still have no surviving school within 4.8 km. Lowest average test scores first closes the comparable schools with the lowest average test score in the year before, ranked on one half of the school-years. A school with fewer than 10 scored pupils on that half takes the median of its municipality, year and size band, with ties broken by enrollment, and in the first closure year, when no school has a score, the rule closes by size. The same list with an alternative within 4.8 km adds the distance priority. Nearest alternative first closes, one school at a time, the comparable school whose displaced children would travel the fewest child-km to their nearest surviving school, and recomputes this after each closure.

Each year's list starts from the schools open that year under the realized closures of earlier years. A list can therefore close a school the Department kept open, and it can close the same school in two years; the

Table E1. Closure lists at the Department's counts: rules, seats and children

	Rule (1)	Counts		Beyond 4.8 km, percent	
		Seats removed (2)	Children displaced (3)	Rural poor (4)	All other (5)
Realized list	The closures the Department chose.	35,455	34,975	19.6	7.5
Smallest first	The smallest comparable schools close first (enrollment the year before).	29,272	33,242	22.2	6.4
Smallest first, alternative within 4.8 km	Smallest first among schools that leave a school serving each of their grades open within 4.8 km; where too few qualify, the count is completed from the other schools, so the condition is a priority, not a guarantee.	30,134	33,698	9.1	3.3
Smallest first, alternative within 3.2 km	The same with 3.2 km.	30,907	34,750	9.3	3.1
Smallest first, alternative within 4.8 km, growth condition †	The same with 4.8 km; a school whose test-score growth is in the top third of its municipality is not closed.	30,400	34,000	10.3	3.5
Lowest average test scores first †	The lowest average test score the year before closes first; a school with fewer than 10 scored pupils takes the median of its municipality, year and size band.	35,007	36,850	21.2	7.3
Lowest average test scores first, alternative within 4.8 km †	The same, with the 4.8 km priority.	34,855	37,339	10.1	3.4
Nearest alternative first	The school whose displaced children would travel the fewest child-km to their nearest surviving school closes next; recomputed after each closure.	32,979	35,217	10.7	3.0

Notes: Closures of 2013–17. Each list re-chooses 315 of the 334 closures, the realized number by municipality, year and size band (thirds of the size of that year's closed schools), one year at a time; the 19 others close under every list. Comparable schools: open that year and no larger than that year's largest closed school. Seat proxy: each closed school's earlier maximum enrollment in grades 3 to 8; it is not capacity for the grades 1 to 7 children counted beside it. Children displaced: grades 1 to 7, the held closures' children included. Beyond 4.8 km: no surviving school serving the grade within 4.8 km. †: ranked on one half of the school-years, measured on the other; mean of the two directions. Summed over the rural-poor children each list displaces, km to the nearest surviving school changes against the realized list by –12.5 percent [–21.0, –2.5] under smallest first with an alternative within 4.8 km and by –32.4 percent [–41.2, –23.0] under nearest alternative first (95 percent bootstrap intervals).

second score list does so for 29 and 35 schools in its two directions. Counts refer to closures per year, not to schools. Random lists draw, in every municipality, year and size band, as many comparable schools as closed there. The random orders of Section 8.3 are a different object: they permute the order in which the 315 realized closures are spared. Neither assigns children at random.

Besides what Section 8.1 holds fixed, the option of leaving the municipality's public schools keeps its estimated value, and no surviving school is designated to receive a closed school's children (Section 2.2).

The lists impose no seat limit. For a descriptive check, I count each closed school's earlier maximum enrollment in grades 3 to 8 as a proxy for seats removed (Table 7, Panel A). Displaced children are counted in grades 1 to 7, so these totals do not establish that receiving schools can accommodate them. Model and recorded inflows above receiving schools' earlier maximum enrollment are not compared, because they may be counted on different bases.

For rural-poor children, under every list, the share the choice model places is within 5 percentage points of the realized closures' share on the same estimates. Of the 19,927 pupil records of Table 4, the model covers 18,834. The change before children re-sort has its own coverage rule: a closed school without a school effect, or a municipality without a surviving school that has one, contributes zero, while its children stay in the denominator. This coverage is not tabulated by list.

Table E2 lists the reported lists, orders and counts, along with additional documented variants of Sections 8.2 and 8.3, each marked as fixed before its results were seen or added after, with where it is reported. It also identifies the planned cross-fitted comparison at the farthest-first counts, which has not been run and for which no results are reported.

Table E2. Reported and planned lists, orders and counts

	Status	Reported in
Same counts as the Department (315 of the 334 closures)		
Realized list	The reference	Table 7; Table E1 to Table E5
1,000 random lists	Fixed before the results	Figure 7; Table E3
Smallest first	Fixed before the results	Table 7; Table E1, Table E4, Table E5
Smallest first, alternative within 4.8 km	Fixed before the results	Table 7; Table E1, Table E4, Table E5
Smallest first, alternative within 3.2 km	Fixed before the results	Table E1, Table E4
Smallest first, alternative within 1.6 km	Fixed before the results	Not reported
Smallest first, 4.8 km, growth condition	Fixed before the results	Table E1, Table E4
Lowest average test scores first	Added after the results were seen	Table 7; Table E1, Table E4, Table E5
Lowest average test scores first, alternative within 4.8 km	Added after the results were seen	Table 7; Table E1, Table E4, Table E5
Nearest alternative first	Fixed before the results	Table 7; Table E1, Table E4, Table E5
Nearest alternative first, counts by region only	Added after the results were seen	Not reported: other counts
Fewer closures (32, 47, 63 and 95 of the 315 kept open)		
Farthest from a surviving school first	Added after the results were seen	Figure 8; Table E6
Largest model access gain first	Added after the results were seen	Figure 8; Table E6
1,000 random orders	Added after the results were seen	Figure 8; Table E6
Lowest school effects first, cross-fitted	Added after the results were seen	Table E7
Highest average test scores first, cross-fitted	Added after the results were seen	Table E8
4.8 km, then highest average test scores, cross-fitted	Added after the results were seen	Table E8
Lowest average test scores with the 4.8 km guarantee, at farthest first's counts	Fixed before its run	Not yet run

Notes: Lists and orders of Section 8. Fixed before the results: the rule was written before any of its results was seen. Statements of direction are adjusted for multiple comparisons within each set of comparisons fixed before its results were seen, not across sets. Lists at the same counts are compared with the realized list; fewer closures keep some realized closures open.

E.2 Different Schools at the Department's Counts: Further Results

Table E3 gives the numbers behind Figure 7. Its percentiles and ranges are spreads across the 1,000 random lists at the full-sample estimates, not sampling intervals. On distance, the realized closures lie at percentiles 1.2, 2.2, 28.2 and 47.9 of the random lists for rural-poor, rural non-poor, urban-poor and urban non-poor children (Panel A). On school effects, the gap is the amount by which the change before children re-sort is less favorable for rural-poor children than for all other displaced children (Panel C). Per child the realized

closures displaced, the realized gap of 0.099 SD lies at the 46.7th percentile of the random lists, whose mean is 0.102. Per child each list displaces, it exceeds the random mean by 0.007, and the middle 95 percent of that difference runs from -0.042 to 0.058 . Measured on the two halves of the school-years, the realized gap lies at the 9.9th and the 59.5th percentiles.

The level of the change, unlike the gap, differs between the realized closures and random lists (Panels B and D). Per child the realized closures displaced, their change before children re-sort is more favorable than that of the average random list by 0.061 SD (0.024) for rural-poor children, 0.092 (0.033) for rural non-poor, 0.053 (0.017) for urban-poor and 0.049 (0.025) for urban non-poor children, and by 0.059 (0.017) for all displaced children. Section 7.2 shows that the level moves when the school effects are estimated without the closure-decision years while the gap does not, so I do not interpret these levels.

Table E4 reports, as numbers only, the choice model's access loss per child each list displaces, in km-equivalents with delta-method standard errors, beside the km to the nearest surviving school. It adds two variants of the rules: smallest first with an alternative within 3.2 km, and smallest first with an alternative within 4.8 km and a condition on test-score growth. None of these figures carries a statement of direction, because the values for three of the lists were known before any test could be fixed.

Table E5 reports further results for the children placed by the choice model, as numbers only. On effectiveness, the school effect plus school-average value-added, each list's contrast with the realized closures has the same sign as on school effects, for rural-poor children and for all others (Panels A and B). For the size and map lists, nearly all of each rural-poor contrast comes from the schools closed. For the two score lists, the schools entered contribute 0.04 to 0.06 SD and the schools closed between -0.01 and $+0.06$ SD, depending on the direction.

Three checks concern the placed contrasts (Panel C). First, school effects can absorb any effect of the move itself on displaced children's scores. Leave-out estimates, which read each school's effect from the half of the school-years that does not contain the closure year, give contrasts of the same sign for all five lists, although they cover only 70.5 to 80.0 percent of the children, against about 94 percent for the main estimates. Second, placing children with shares by family type or with shares pooled across types changes the contrasts by at most 0.007 SD. Third, letting rural families' response to effectiveness differ, at five values spanning the range its estimates allow, moves the realized rural-poor level from -0.058 to $+0.002$, and every contrast keeps its sign at both ends. At a nonzero rural difference families sort on school-average value-added, an estimate from the same test-score data, so these rows check signs only. The panel also gives the contrasts under the school effects estimated without the closure-decision years.

On the 18,834 pupil records the choice model covers, the simulated change is -0.056 SD (0.036) for rural-poor children minus all others, beside the recorded -0.074 (0.031). They are not differenced: that comparison is the check of Section 6.2 (Table D3). Expected exits, extrapolated from the event study, are 0.107 per displaced rural-poor child and 0.087 per displaced child under the realized closures. The list that sets the bound of 0.233 SD in Section 8.2 is lowest average test scores first. A contrast large enough to be

Table E3. The realized list among 1,000 random lists: percentiles and spreads

<i>Panel A. Access per child each list displaces</i>				
	Rural poor (1)	Rural non-poor (2)	Urban poor (3)	Urban non-poor (4)
Km to the nearest surviving school				
Realized list	3.43	3.37	1.91	2.12
Random lists, median	3.68	3.60	1.96	2.13
Realized list's percentile	1.2	2.2	28.2	47.9
Km-equivalents, choice model				
Realized list	1.84	1.91	0.97	1.00
Random lists, median	2.00	2.07	1.06	1.09
Realized list's percentile	1.3	1.2	3.0	3.3
<i>Panel B. Change in school effects before children re-sort, per child the realized list displaced, SD</i>				
	Rural poor (5)	Rural non-poor (6)	Urban poor (7)	Urban non-poor (8)
Realized list	−0.054	0.023	0.048	0.059
Random lists, median	−0.114	−0.069	−0.005	0.011
Realized minus random mean	0.061 (0.024)	0.092 (0.033)	0.053 (0.017)	0.049 (0.025)
<i>Panel C. The gap: the change for all other minus that for rural poor, SD</i>				
	Full sample		Halves, per realized child	
	Per realized child (9)	Per own child (10)	First half (11)	Second half (12)
Realized list	0.099	0.099	0.073	0.130
Random lists, mean	0.102	0.093	0.122	0.120
2.5th percentile	0.048	0.041	0.054	0.047
97.5th percentile	0.154	0.141	0.197	0.187
Realized list's percentile	46.7	59.8	9.9	59.5
Realized minus random mean	−0.003	0.007	−0.049	0.010
2.5th percentile	−0.054	−0.042	−0.124	−0.057
97.5th percentile	0.051	0.058	0.019	0.083
<i>Panel D. Change per child each list displaces: realized minus random mean, SD</i>				
	Rural poor (13)		All other (14)	
Realized minus random mean	0.051		0.057	
2.5th percentile	0.001		0.005	
97.5th percentile	0.099		0.108	

Notes: Closures of 2013–17. Random lists close comparable schools at the realized counts in every municipality, year and size band (Table E1); percentiles and ranges are spreads across the 1,000 lists at the full-sample estimates, not sampling intervals. A: per child each list displaces. B and C: the surviving schools minus the school closed; positive is more favorable. B: the last row's standard errors from the two halves of the school-years. C: the gap is how much less favorable the change is for rural-poor children than for all other displaced children; halves: school effects estimated on each half of the school-years. D: per child each list displaces.

Table E4. Closure lists at the Department's counts: access by list*Panel A. Access, per child each list displaces*

	Km-equivalents		Km to the nearest surviving school	
	Rural poor (1)	All other (2)	Rural poor (3)	All other (4)
Realized list	1.84 (0.11)	1.14 (0.07)	3.43	2.15
Smallest first	2.20 (0.09)	1.28 (0.06)	3.79	2.25
Smallest first, alternative within 4.8 km	1.79 (0.11)	1.18 (0.07)	3.11	2.07
Smallest first, alternative within 3.2 km	1.68 (0.12)	1.12 (0.07)	2.75	1.83
Smallest first, 4.8 km, growth condition †, first half	1.81 (0.11)	1.20 (0.07)	3.19	2.10
Smallest first, 4.8 km, growth condition †, second half	1.83 (0.11)	1.19 (0.07)	3.25	2.11
Lowest average test scores first †, first half	1.96 (0.10)	1.17 (0.07)	3.72	2.19
Lowest average test scores first †, second half	1.93 (0.11)	1.22 (0.07)	3.63	2.18
Lowest average test scores first, 4.8 km †, first half	1.71 (0.12)	1.12 (0.07)	3.08	1.99
Lowest average test scores first, 4.8 km †, second half	1.74 (0.12)	1.17 (0.07)	3.14	2.02
Nearest alternative first	1.65 (0.12)	1.01 (0.08)	2.67	1.50

Panel B. Change in school effects, list minus realized list, per child the realized list displaced, SD

	Before re-sorting		Placed by the model	
	Rural poor (5)	All other (6)	Rural poor (7)	All other (8)
Smallest first	−0.098 (0.039)	−0.089 (0.030)	−0.092 (0.034)	−0.057 (0.031)
Smallest first, alternative within 4.8 km	−0.081 (0.035)	−0.084 (0.029)	−0.076 (0.032)	−0.061 (0.031)
Smallest first, alternative within 3.2 km	−0.073 (0.040)	−0.083 (0.028)	−0.071 (0.037)	−0.074 (0.029)
Smallest first, 4.8 km, growth condition †	−0.071 (0.049)	−0.062 (0.040)	−0.044 (0.043)	−0.020 (0.044)
Lowest average test scores first †	0.048 (0.042)	0.094 (0.032)	0.061 (0.045)	0.119 (0.038)
Lowest average test scores first, 4.8 km †	0.057 (0.036)	0.104 (0.034)	0.067 (0.039)	0.129 (0.041)
Nearest alternative first	−0.054 (0.035)	−0.028 (0.030)	−0.049 (0.034)	−0.006 (0.033)

Notes: Closures of 2013–17; lists as in Table E1. Panel A: the choice model's access loss (Table 6), standard errors by the delta method; and the km from the closed school to the nearest surviving school serving the grade, without a standard error. †: chosen on one half of the school-years; each direction shown. Panel B: before children re-sort as in Table 7, Panel C (standard errors: the largest of four imputation rules); placed by the model, all the list's placed children, the † lists the mean of the two directions with the mean of their standard errors. The part due to the number of children placed is below 0.006 SD in every entry.

Table E5. Children placed by the choice model: further results*Panel A. Rural-poor children: list minus realized list, SD*

	Smallest first (1)	Smallest first, 4.8 km (2)	Lowest scores † (3)	Lowest scores, 4.8 km † (4)	Nearest alternative (5)
School effects	−0.093 (0.036)	−0.085 (0.034)	0.072 (0.052)	0.080 (0.047)	−0.062 (0.040)
Effectiveness	−0.055 (0.020)	−0.023 (0.018)	0.087 (0.027)	0.077 (0.026)	−0.052 (0.018)
School effects, from the schools closed	−0.110 (0.034)	−0.097 (0.032)	0.024 (0.042)	0.030 (0.038)	−0.077 (0.033)
School effects, from the schools entered	0.017 (0.014)	0.012 (0.015)	0.048 (0.022)	0.050 (0.023)	0.015 (0.018)
School effects, leave-out estimates	−0.144 (0.067)	−0.068 (0.068)	0.060 (0.057)	0.109 (0.053)	−0.052 (0.063)

Panel B. All other children: list minus realized list, SD

	Smallest first (6)	Smallest first, 4.8 km (7)	Lowest scores † (8)	Lowest scores, 4.8 km † (9)	Nearest alternative (10)
School effects	−0.068 (0.036)	−0.068 (0.034)	0.141 (0.046)	0.144 (0.048)	−0.008 (0.034)
Effectiveness	−0.058 (0.018)	−0.050 (0.018)	0.150 (0.026)	0.145 (0.026)	−0.006 (0.016)

Panel C. Checks, rural-poor children, school effects: list minus realized list

	Smallest first (11)	Smallest first, 4.8 km (12)	Lowest scores † (13)	Lowest scores, 4.8 km † (14)	Nearest alternative (15)
Children with leave-out estimates, percent	70.5	71.4	80.0	78.8	73.9
Shares pooled across family types	−0.087	−0.081	0.075	0.084	−0.058
Without the closure-decision years, closures of 2014–17	−0.054	−0.047	n.a.	n.a.	−0.035
Rural response to effectiveness: lowest	−0.105	−0.095	0.025	0.073	−0.066
Rural response to effectiveness: highest	−0.082	−0.079	0.117	0.089	−0.059
Smallest difference read alone	0.118	0.116	0.160	0.141	0.137
Expected exits per displaced rural-poor child	0.111	0.109	0.110	0.108	0.105

Notes: Closures of 2013–17; lists as in Table E1; children placed by the choice model given re-enrollment, school effects fixed; effectiveness adds school-average value-added. A and B: as Table 7, Panel D, delta-method standard errors; †: mean of the two directions, standard error the mean of the two. C: numbers only; the response to effectiveness at five values spanning its estimates (lowest and highest contrast, both directions for †); no version without the decision years for †, measured on halves. Realized closures: expected exits 0.107 per rural-poor child; on 18,834 pupil records, simulated rural poor minus all other −0.056 (0.036), recorded −0.074 (0.031), not differenced. No list's difference exceeds 0.233 SD (95 percent, simultaneous, allowing for the model's error, Table D3).

read on its own, given the model's error and sampling error, would need to exceed 0.116 to 0.160 SD in absolute value, depending on the list.

E.3 Closing Fewer Schools: Further Results

Table E6 reports both orders at 32, 47, 63 and 95 closures spared: savings forgone, access loss avoided, break-even values, the change in school effects before children re-sort and with children placed by the choice model, and the far tail. "Largest model access gain first" forgoes the same savings as "Farthest from a surviving school first" at the low end and more at the high end, because it spares larger schools. At 47, the closures that remain under the first order remove 29,459 seats, against 35,455 for the realized closures.

The break-even value is the savings forgone a year divided by the access loss avoided, for all children, in dollars per km-equivalent child-year. Treating the savings as fixed amounts, its 95 percent interval divides the savings by the two ends of the interval for the access loss avoided. Under the second order the values run from \$32 to \$49 at the low end and from \$185 to \$267 at the high end. They are lower than under the first because the second order is chosen to reduce the model's access loss, the denominator. No break-even value is computed for school effects, whose differences are never distinguishable from zero here, and no price is attached to access.

The two orders differ at the far tail and on the model's measure. "Farthest from a surviving school first" brings the share of rural-poor children with no surviving school within 4.8 km to 0.14 percent at 47 and keeps it below 0.2 percent at 63 and 95; under "Largest model access gain first" the share stays between 9.4 and 10.8 percent at the four counts. The second order avoids more of the model's access loss, 33.7 to 63.2 percent of rural-poor children's. Both orders avoid more access loss for all children than the 95th percentile of the 1,000 random orders, at every count.

The table reports the second order's differences in school effects from the realized closures as numbers only. For the change before children re-sort, the random orders' 5th to 95th percentiles include zero at every count. The path of "Farthest from a surviving school first" lies below the random orders' median at all four counts, with both halves of the school-years agreeing; at 47 it lies at the 0.4th percentile of the random orders. With children placed by the choice model, the random orders' rural-poor medians lie between 0.003 and 0.008.

Figure 8 draws simultaneous 95 percent bands, added after the results were seen: with probability at least 0.95 a path lies inside its band at all 31 counts drawn. For access the bands use the delta method in the five coefficients of the choice model, with critical values of 2.071 and 2.078 for rural-poor and other children under the first order and 2.187 and 2.172 under the second. For school effects they use the split-half noise model with the largest standard error and the largest critical value of the four imputation rules (Appendix E.5); the critical values are 2.653, 2.764, 2.520 and 2.646 in the same order, and zero lies inside every band. Figure E1 draws Panel C of Figure 8 with the bands of both orders.

A third order uses estimates the Department did not compute (Table E7). It spares first the closures

Table E6. Closing fewer schools at 32, 47, 63 and 95 closures spared


	Farthest first				Largest model access gain first			
	32 (1)	47 (2)	63 (3)	95 (4)	32 (5)	47 (6)	63 (7)	95 (8)
Savings forgone, thousand dollars a year								
Low end	390.3	573.3	768.4	1,158.7	390.3	573.3	768.4	1,158.7
High end	2,008.4	2,925.0	3,833.9	5,755.1	2,242.6	3,226.8	4,247.3	6,349.6
Access loss avoided, km-equivalent child-years								
Rural poor	5,394.7 (123.7)	7,835.7 (183.8)	9,693.5 (280.3)	12,742.9 (405.4)	8,692.3 (221.8)	11,136.8 (356.4)	13,241.3 (480.4)	16,283.7 (732.3)
All other	2,075.8 (76.0)	2,656.6 (99.8)	3,068.8 (127.9)	4,466.8 (201.2)	3,434.0 (76.8)	4,406.3 (121.1)	5,392.9 (176.0)	7,474.2 (310.1)
All children	7,470.5 (180.9)	10,492.4 (269.6)	12,762.2 (392.8)	17,209.7 (596.4)	12,126.3 (295.4)	15,543.2 (460.6)	18,634.2 (617.6)	23,757.9 (937.1)
Access loss avoided, percent of the group's realized loss								
Rural poor	20.9 (1.0)	30.4 (1.4)	37.6 (1.6)	49.5 (1.8)	33.7 (1.4)	43.2 (1.5)	51.4 (1.6)	63.2 (1.4)
All other	13.0 (1.0)	16.6 (1.2)	19.2 (1.4)	28.0 (1.9)	21.5 (1.4)	27.6 (1.4)	33.8 (1.5)	46.8 (1.5)
All children	17.9 (0.6)	25.1 (0.8)	30.6 (0.9)	41.2 (1.0)	29.1 (1.0)	37.3 (1.2)	44.7 (1.2)	56.9 (1.2)
Break-even value, all children, dollars a year per km-equivalent child-year								
Low end	52 [50, 55]	55 [52, 58]	60 [57, 64]	67 [63, 72]	32 [31, 34]	37 [35, 39]	41 [39, 44]	49 [45, 53]
High end	269 [257, 282]	279 [265, 294]	300 [283, 320]	334 [313, 359]	185 [177, 194]	208 [196, 220]	228 [214, 244]	267 [248, 290]
Change in school effects before children re-sort, SD								
Rural poor	-0.013 (0.009)	-0.024 (0.022)	-0.015 (0.023)	-0.005 (0.034)	-0.020 (0.030)	-0.021 (0.031)	-0.004 (0.034)	0.011 (0.034)
All other	-0.006 (0.004)	-0.010 (0.009)	-0.011 (0.009)	-0.010 (0.016)	-0.008 (0.015)	-0.005 (0.016)	0.007 (0.018)	0.005 (0.019)
Change in school effects, children placed by the choice model, SD								
Rural poor	-0.010 (0.011)	-0.016 (0.020)	-0.009 (0.022)	-0.005 (0.034)	-0.018 (0.028)	-0.014 (0.030)	-0.002 (0.034)	0.013 (0.035)
All other	-0.004 (0.004)	-0.006 (0.008)	-0.009 (0.008)	-0.007 (0.016)	-0.001 (0.015)	0.004 (0.016)	0.025 (0.021)	0.029 (0.022)
Children spared and the far tail								
Rural-poor children spared	2,002	2,964	3,884	5,651	2,780	3,711	4,741	6,424
All other children spared	1,057	1,570	1,934	3,260	2,244	3,310	4,453	7,034
Rural poor beyond 4.8 km, %	6.43	0.14	0.16	0.19	10.81	10.42	9.61	9.43
1,000 random orders at the same count (the same under both orders)								
Access, all: median	4,401.4	6,521.7	8,734.8	13,095.6	4,401.4	6,521.7	8,734.8	13,095.6
5th percentile	3,534.8	5,512.0	7,566.1	11,812.8	3,534.8	5,512.0	7,566.1	11,812.8
95th percentile	5,401.2	7,608.0	9,982.5	14,416.0	5,401.2	7,608.0	9,982.5	14,416.0
School effects, rural poor: median	0.006	0.009	0.013	0.018	0.006	0.009	0.013	0.018
Placed, rural poor: median	0.003	0.004	0.005	0.008	0.003	0.004	0.005	0.008

Notes: Closures of 2013–17 kept open, of 315, in two orders (Figure 8); farthest first is “Farthest from a surviving school first.” Savings: Table F1’s two ends. Access: the choice model’s measure, 28,024 children, against the realized loss of 41,722.8 (25,764.8 rural poor); standard errors in parentheses by the delta method. Break-even: savings over access avoided; its interval divides the savings by the ends of the access interval. School effects: the closures that remain minus the realized closures, per child the realized closures displaced (34,037), standard errors the largest of four imputation rules; placed by the model: given re-enrollment. Far tail: of displaced rural-poor children (19.6 percent under the realized closures). Random orders: a spread across orders, not sampling error. In brackets: 95 percent intervals, delta method.

Figure E1. Closing fewer schools: the change in school effects, with both orders' bands

Change in school effects before children re-sort

The closures that remain minus the realized closures, SD per realized displaced child; above 0 is more favorable



Notes: Panel C of Figure 8 with the simultaneous 95 percent bands of both orders. The closures that remain minus the realized closures, before children re-sort, SD per child the realized closures displaced; above zero is more favorable. Bars: pointwise 95 percent intervals at 32, 47, 63 and 95; standard errors: the largest of four imputation rules (Table E10). Gray: 5th to 95th percentile of 1,000 random orders.

whose estimated contribution to the change in school effects before children re-sort is least favorable on one half of the school-years, measures the result on the other half, and averages the two directions; closures without estimates on both halves come last. For rural-poor children its differences from the realized closures run from 0.049 (0.023) at 32 to 0.098 (0.028) at 95, where positive is more favorable; each standard error is the mean of the two directions' standard errors, a bound. Its access loss avoided is 73 to 81 percent of that of “Farthest from a surviving school first” for all children and 60 to 70 percent for rural-poor children. These figures carry no statement of direction.

Two further orders use what the Department held, past test scores and the map, and were added after the results were seen (Table E8). “Highest average test scores first (cross-fitted)” spares first the closures whose average test score the year before most exceeds the enrollment-weighted average of the surviving schools serving the same grades in the municipality, weighted by the children displaced. It ranks the 142 of 315 closures that have scores and school-effect estimates on both halves of the school-years, so that each can be measured on the half not used to rank it; the other 173 hold 55.4 percent of the displaced rural-poor children and come after all 142. “4.8 km, then highest average test scores (cross-fitted)” spares the 47 closures that

Table E7. Closing fewer schools in the order of school effects, cross-fitted

	Change before children re-sort		Minus the farthest-first order	
	Rural poor (1)	All other (2)	Rural poor (3)	All other (4)
32 closures kept open	0.049 (0.023)	0.042 (0.015)	0.054 (0.024)	0.046 (0.015)
47 closures kept open	0.067 (0.025)	0.047 (0.017)	0.089 (0.034)	0.059 (0.019)
63 closures kept open	0.079 (0.026)	0.048 (0.019)	0.096 (0.035)	0.061 (0.021)
95 closures kept open	0.098 (0.028)	0.054 (0.021)	0.102 (0.038)	0.063 (0.024)

Notes: Closures of 2013–17 kept open in the order of the closed schools’ estimated school effects, chosen on one half of the estimates and measured on the other, the two directions averaged; this order uses estimates the Department did not have. Change before children re-sort: the closures that remain minus the realized closures, per child the realized closures displaced; positive is more favorable. Standard errors in parentheses, clustered by municipality; for the average of the two directions, bounded by the average of their standard errors.

“Farthest from a surviving school first” spares up to its 4.8 km point, those that can be ranked first, and then continues by the same score. Each order is chosen on one half of the school-years and measured on the other.

Against “Farthest from a surviving school first” at the same count, the closures that remain under the first score order leave a higher access loss at all four counts: that order avoids 72 to 78 percent as much access loss for all children and 58 to 65 percent as much for rural-poor children. Its remaining closures give a more favorable change in school effects before children re-sort for rural-poor children at 47 and 63 closures spared, by 0.072 (0.020) and 0.065 (0.018) SD per child the realized closures displaced, and for other children at 32, 47 and 63. The second score order spares the same closures as the first distance order at 47. At 63, it avoids 518.5 (63.1) more km-equivalent child-years for all children, and its remaining closures give a more favorable change for both groups, by 0.040 (0.015) SD for rural-poor children and 0.031 (0.009) for others; its access loss avoided for rural-poor children is 3.7 percent below the first distance order’s, a point estimate with no test. At 95, neither its access difference for all children nor its difference in school effects for rural-poor children can be distinguished from zero, its remaining closures give other children a change more favorable by 0.046 SD (0.017), and its access loss avoided for rural-poor children is 9.2 percent below the first distance order’s, again a point estimate. At 32, both directions of the second score order spare the same closures, so the set there may be fixed by which closures can be ranked, not by their scores. These statements are adjusted for multiple comparisons within six sets of comparisons fixed before the results were seen, and each requires both directions to agree in sign. The two directions differ in size, and I do not interpret that difference.

Table E8. Closing fewer schools in orders that use test scores and the map

	Highest average test scores first				4.8 km, then highest average test scores			
	32 (1)	47 (2)	63 (3)	95 (4)	32 (5)	47 (6)	63 (7)	95 (8)
Savings forgone, thousand dollars a year								
Low end	390.3	573.3	768.4	1,158.7	390.3	573.3	768.4	1,158.7
High end	2,083.1	2,964.2	3,983.5	5,942.7	2,093.6	2,925.0	3,995.2	5,938.4
Access loss avoided, km-equivalent child-years								
Rural poor	3,512.0 (194.9)	4,651.6 (274.0)	5,630.7 (352.9)	8,225.6 (545.4)	6,174.6 (152.1)	7,835.7 (183.8)	9,338.9 (294.0)	11,575.8 (430.6)
All other	2,262.0 (152.5)	2,949.0 (187.3)	3,685.7 (236.5)	5,147.5 (346.9)	2,263.2 (86.0)	2,656.6 (99.8)	3,941.8 (134.8)	5,732.1 (241.0)
All children	5,774.0 (275.1)	7,600.6 (374.0)	9,316.4 (491.7)	13,373.1 (740.2)	8,437.8 (219.6)	10,492.4 (269.6)	13,280.7 (419.7)	17,307.8 (630.5)
All children, percent of the realized loss	13.8 (0.2)	18.2 (0.2)	22.3 (0.1)	32.1 (0.1)	20.2 (0.7)	25.1 (0.8)	31.8 (0.8)	41.5 (0.9)
All children, minus farthest first	-1,696.5 (145.7)	-2,891.8 (149.3)	-3,445.9 (155.8)	-3,836.6 (229.9)	967.3 (40.5)	0.0 —	518.5 (63.1)	98.1 (139.2)
Break-even value, all children, dollars a year per km-equivalent child-year								
Low end	68 [62, 75]	75 [69, 83]	82 [75, 92]	87 [78, 97]	46 [44, 49]	55 [52, 58]	58 [54, 62]	67 [62, 72]
High end	361 [330, 398]	390 [356, 432]	428 [387, 477]	444 [401, 498]	248 [236, 261]	279 [265, 294]	301 [283, 321]	343 [320, 369]
Change in school effects before children re-sort, SD								
Rural poor	0.036 (0.026)	0.050 (0.027)	0.049 (0.027)	0.045 (0.029)	-0.032 (0.027)	-0.023 (0.028)	0.024 (0.028)	0.043 (0.030)
All other	0.033 (0.012)	0.042 (0.014)	0.041 (0.016)	0.027 (0.018)	-0.011 (0.010)	-0.012 (0.010)	0.018 (0.013)	0.037 (0.017)
Rural poor, minus farthest first	0.040 (0.028)	0.072 (0.020)	0.065 (0.018)	0.048 (0.027)	-0.027 (0.025)	0.000 —	0.040 (0.015)	0.047 (0.025)
All other, minus farthest first	0.037 (0.013)	0.055 (0.014)	0.053 (0.015)	0.036 (0.019)	-0.006 (0.009)	0.000 —	0.031 (0.009)	0.046 (0.017)

Notes: Closures of 2013–17 kept open, of 315; both orders were added after the results were seen. The first spares first the closures whose average test score the year before most exceeds that of the municipality's surviving schools, among the 142 with scores and school effects on both halves of the school-years; the other 173 (55.4 percent of displaced rural-poor children) follow. The second spares farthest first's 47 (ranked ones first), then continues by score: at 47 the two coincide. Each order is chosen on one half and measured on the other; mean of the two directions, standard errors bounded by the mean of the two. Measures as in Table E6; differences from farthest first on the same half. In brackets: 95 percent intervals, delta method.

E.4 Holding the Closures Fixed: Further Results

Table E9 reports the reach exercise of Section 8.4 and the information scenario below.

The reach exercise assigns each displaced child's pupil record to the most effective surviving school of the municipality serving the grade that has a seat and lies within a stated distance of the closed school. Distance from the closed school uses the pupil records' own km between schools where they exist and the straight line times the road ratio of the school's class otherwise. A seat is open when a school's enrollment is below its maximum recorded enrollment up to the closure year, with a floor that admits the realized allocation; this is an accounting rule, not a measure of classrooms. Effectiveness is picked on one half of the school-years and scored on the other, both ways round, and every other child's school is held fixed. The exercise uses 12,564 pupil records, with a rural-poor weight of 5,369, and 95 percent intervals from

Table E9. Holding the closures fixed: reach, information and teachers*Panel A. Gains, SD, both groups*

	Rural poor (1)	Urban non-poor (2)
Reach: no farther than the school entered	0.048 [0.028, 0.073]	0.067 [0.031, 0.102]
Reach: within one more km	0.064 [0.040, 0.090]	0.099 [0.057, 0.141]
Information on effectiveness, 1.1 km per SD	0.047 [0.037, 0.058]	0.062 [0.041, 0.087]

Panel B. Rural-poor children only

	Rural poor (3)
Gain in effectiveness, anywhere in the municipality, SD	0.204 [0.146, 0.269]
Change in km, within one more km	-0.24 [-0.33, -0.16]
Change in km, anywhere in the municipality	3.1 [2.4, 3.9]
Change in km, information	0.004 [-0.025, 0.036]

Panel C. Point estimates only, both groups

	Rural poor (4)	Urban non-poor (5)
Change in km, no farther than the school entered	-0.31	-0.37
Information: placed above earlier maximum enrollment, percent	78	70
Teachers reassigned: value-added received, SD	0.0007	0.0010

Panel D. No interval computed, urban non-poor children

	Urban non-poor (6)
Gain in effectiveness, anywhere in the municipality, SD	0.224
Change in km, within one more km	-0.20
Change in km, anywhere in the municipality	2.87

Notes: Closures of 2013–17, closures and other children held fixed. Reach and information: 12,564 re-enrolled pupil records; effectiveness (school effect plus school-average value-added) chosen on one half of the school-years and measured on the other; intervals from a bootstrap over municipalities. Reach: the most effective surviving school with a seat within the stated distance of the closed school. Information: the choice model's placements with an assumed response to a report card, full take-up, no seat limit. Change in km: to the school entered under the exercise less to the school entered in the records. Teachers: 626 posts reassigned within municipality, subject and grade band; 201 change teacher. No sampling interval is reported for teacher reassignment. In brackets: 95 percent intervals for reach and information.

a bootstrap over municipalities. Table E9 gives intervals for both groups' gains no farther than and within one more km of the school entered, and for the information scenario (Panel A), and for rural-poor children's gain anywhere in the municipality and their changes in km (Panel B); the km at no extra distance and the seat shares (Panel C), and urban non-poor children's gain and km anywhere in the municipality and km within one more km (Panel D), have no interval computed. Rural non-poor and urban-poor children are not reported separately.

A teacher-reassignment exercise keeps the teacher pool and the posts fixed. Of the 2,145 teachers whose schools closed, it retains the 626 who taught a subject with a standardized test at exactly one surviving school the next year and have an estimate of value-added. Their posts are reassigned within municipality, subject profile and grade band, maximizing the value-added children receive with rural-poor children counted twice and value-added adjusted for sampling noise; 201 posts change teacher. The point-estimated increase in the value-added received is 0.0007 SD for rural-poor children and 0.0010 SD for urban non-poor children (Table E9); no sampling interval is reported for teacher reassignment. These gains are conditional on the fixed pool of teachers and posts. This design does not establish whether alternative assignments satisfy seniority rules or teachers' preferences.

In the realized placement, within municipality, year, subject and grade band, teachers with higher estimated value-added went to posts serving fewer rural-poor children (permutation $p = 0.005$). The difference is about 0.0005 SD, smaller than the 0.00066 SD that the test detects with 80 percent power in simulation. This diagnostic describes the realized assignment; it is distinct from the gain from reassigning teachers.

The information scenario adds to the choice model a response to a report card on each surviving school's estimated effectiveness, with an assumed weight of 1.1 km per SD; the weight is an input of the scenario, not an estimate. One half of the school-years supplies the card and the other evaluates the effectiveness of the resulting allocation, with the roles reversed and averaged. The scenario compares two allocations from the model on the same records and assumes full take-up, no seat limit and unchanged choices by other families. It uses an earlier specification of the choice model, with 805 markets and 98,069 families, and 12,564 of the 15,025 re-enrolled records, which carry 84.6 percent of the rural-poor weight. The point-estimated gain in effectiveness for rural-poor children is 0.047 SD [0.037, 0.058], with the scenario's modeled distance changing by 0.004 km [−0.025, 0.036]; for urban non-poor children the gain is 0.062 [0.041, 0.087]. Seats above the receiving schools' earlier maximum enrollment account for 78 percent of the reallocation of rural-poor children and 70 percent for urban non-poor children.

E.5 Inference

Standard errors of access come from the delta method in the five coefficients of the choice model, clustered by municipality. The derivatives are taken numerically at two step sizes, and a figure is reported only when the two give standard errors within 1 percent of each other. Enrollment shares, the set of closed schools, population weights, the reference distance and the distance power are held fixed.

Standard errors of the change in school effects before children re-sort come from the split of the school-years into halves, clustered by municipality. The half estimates leave out displaced pupils' records from their move year on; the full-sample estimates keep them. Among the schools that carry weight in these comparisons, 960 have estimates on both halves, 209 on one half (174 of them with scored years in both halves) and 19 on neither (Table E10). Schools without an estimate in both halves have no half-difference, and they carry 13 to 19.5 percent of the absolute weight, so their noise variance is imputed. The rule set before the results took it from schools with exactly two scored years and both halves, of which fewer than 10 exist. Because that base is so small, the rule was changed after the results were seen. Each interval now uses the largest of four standard errors. Two come from the rule set before the results and from the same rule applied to all schools with both halves. Two were added after the results were seen: they make the noise variance inversely proportional to the school's scored years or to its scored pupil records in the half it has, scaled within enrollment terciles from all schools with both halves. Among schools with both halves, the squared half-difference falls with a school's scored pupil records, but less than in proportion, and within enrollment terciles it varies little with scored years.

The rule binds in two places. In Figure 8, the largest standard error comes from the rule on scored pupil records at 15 of the 16 estimates, and more than half of the variance is imputed at 95 closures spared under "Farthest from a surviving school first" (51.0 percent for rural-poor children and 78.8 percent for others) and at 7 of the 8 estimates for "Largest model access gain first." No statement about these comparisons depends on the choice of rule. In Table 7, Panel C, the rule changed one statement: under the larger standard error, 0.039 against 0.035, the difference between smallest first and the realized closures for rural-poor children can no longer be distinguished from zero after adjustment.

Lists chosen on one half of the school-years are measured on the other, so their noise comes from one half only and its variance is doubled. The mean of the two directions takes the average of the two standard errors, a bound that holds whatever the correlation between directions; the value under independence is printed beside it.

For children placed by the choice model, a standard error combines the uncertainty in the model's coefficients, from the delta method with the school constants re-solved at each step, and the noise in the school effects, from the split halves as above. The calculation treats the two parts as independent under the main specification, which does not use school-effect estimates to estimate the choice coefficients, and the coefficient part is negligible, at most 0.002 of the variance for the rural-poor contrasts. Covariance between enrollment choices and test-score shocks is not estimated. A contrast D with the realized closures is tested with $p = 2[1 - \Phi((|D| - a)/SE)]$, and $p = 1$ when $|D| \leq a$, where a is the list's tolerance for the model's error (Table D3). A contrast that passed would also have needed the same sign on effectiveness, on the leave-out estimates (for the three lists not chosen on one half) and across the range of the rural difference in the response to effectiveness, the same sign in both directions for the score lists, and a placed share within 5 percentage points of the realized closures'. No contrast passed, so these conditions decided nothing.

Statements of direction are adjusted for multiple comparisons with Holm's step-down procedure at the 5

Table E10. Imputed noise and the standard errors of the change in school effects*Panel A. Schools carrying weight, and their share of the weight*

	Schools (1)	Rural poor, weight share (2)	All other, weight share (3)
An estimate in both halves	960	0.842	0.859
In one half only	209	0.142	0.112
In neither half	19	0.015	0.029

Panel B. Fewer closures: the standard error under each imputation rule, SD

	Two scored years (4)	All with both halves (5)	Scored years (6)	Scored pupil records (7)	Largest (8)
<i>Farthest from a surviving school first</i>					
32 kept open, rural poor	0.0070	0.0079	0.0088	0.0086	0.0088
32 kept open, all other	0.0035	0.0036	0.0038	0.0039	0.0039
47 kept open, rural poor	0.0198	0.0203	0.0210	0.0215	0.0215
47 kept open, all other	0.0069	0.0071	0.0073	0.0087	0.0087
63 kept open, rural poor	0.0211	0.0218	0.0228	0.0232	0.0232
63 kept open, all other	0.0071	0.0073	0.0075	0.0089	0.0089
95 kept open, rural poor	0.0237	0.0246	0.0264	0.0338	0.0338
95 kept open, all other	0.0075	0.0079	0.0089	0.0162	0.0162
<i>Largest model access gain first</i>					
32 kept open, rural poor	0.0205	0.0211	0.0227	0.0298	0.0298
32 kept open, all other	0.0078	0.0080	0.0086	0.0152	0.0152
47 kept open, rural poor	0.0218	0.0226	0.0243	0.0310	0.0310
47 kept open, all other	0.0094	0.0099	0.0105	0.0163	0.0163
63 kept open, rural poor	0.0244	0.0256	0.0276	0.0341	0.0341
63 kept open, all other	0.0105	0.0113	0.0126	0.0184	0.0184
95 kept open, rural poor	0.0238	0.0251	0.0273	0.0338	0.0338
95 kept open, all other	0.0110	0.0123	0.0141	0.0192	0.0192

Notes: Schools that carry weight in the comparisons of the closure lists and of closing fewer schools; weight shares at 47 closures kept open, farthest first. A school without an estimate in both halves has no half-difference, so its noise variance is imputed. Panel B: the standard error of the change in school effects before children re-sort, per child the realized closures displaced, under a noise variance taken from schools with exactly two scored years; the same over all schools with both halves; inversely proportional to the school's scored years; or to its scored pupil records (the last two added after the results were seen). The largest is used, here and in Table 7.

percent level: within a set of comparisons, the p -values are ordered from smallest to largest, and each in turn is compared with 0.05 divided by the number of hypotheses not yet rejected, stopping at the first that fails. Each set was fixed before that set's own results were seen: the five lists of Table 7 for rural-poor children, separately for the change before children re-sort and for children placed by the choice model; for closing fewer schools, the four counts for each group and object; and the six sets of the orders on test scores. The adjustment applies within each set, not across sets. Figures outside these sets, among them every result for all other children in Table 7, effectiveness, the model's access loss by list and the contrast between the two orders of Section 8.3, print as numbers only.

The intervals for children placed by the choice model are conditional on the delta-method approximation, the split-half noise calculation and the independence assumption above. They have not received the numerical simulation check of parameter uncertainty. The 0.233 SD simultaneous bound is conditional on these standard errors; it is not evidence that the lists have equivalent effects.

A simulation check of the standard errors for placed children draws the five coefficients of the choice model 200 times from their estimated distribution, re-solves the school constants at each draw, and pairs each draw with random signs on the half-differences of the school effects: one per municipality, and one per school estimated on one half only, scaled by its imputed variance. It checks the computation and the curvature in the coefficients, not the split-half estimate of the noise. The intervals printed in this paper have not been through it. A figure whose draws' standard deviation and standard error differed by more than a quarter would be printed without an interval, and the bound of 0.233 SD would then not be printed. No statement of direction can change, because no contrast passed.

F Fiscal Accounting

This appendix sets out the account of savings used in Section 8.3, the only exercise in the paper that uses dollars. It gives the components of a closure's savings, the prices behind them, the totals, and what the account leaves out.

F.1 What a Full Budget Comparison Would Record

For a set of closures C and an assignment α of children and staff, a full annual comparison of budgets would record

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

where F is facilities spending avoided, W is spending on staff posts avoided, K is spending on the transition, ΔT is the change in transport spending and E is the change in spending from children who leave the public schools. A positive E is an accounting saving, not a gain to anyone. Spending on the transition is dated to its year, and the value of closed buildings belongs to a capital account. The account of Section 8.3 measures parts of F and W only, as a range of assumptions; it is not an estimate of equation (F.1).

F.2 Savings per Closure: Components and Prices

Facilities spending comes from a regression on the spending reported for 127 surviving schools in the 2019–20 school year: \$18,657 (5,200) plus \$121.16 (9.64) per enrolled pupil, standard errors in parentheses, with an R^2 of 0.558. No school closed in 2013–17 appears among these schools, so applying the regression to the closures is an extrapolation. The low end keeps the intercept and subtracts an assumed closed-site deduction of \$6,460 a year, giving \$12,197 per closure. The savings ledger derives the deduction by applying a 73 percent account share to reported utility payments of about \$8,850 per unused school. It is a calibration rather than a measured avoidable cost for a school in this sample. The high end adds the slope times the school’s enrollment in its last operating year, \$121.16 per pupil. Whether a closure avoids this part of facilities spending is not known, and that is the reason for the range.

The high end also adds \$31,544 a year per closure for staff posts. In the Department’s staff records, the staff of 467 closed schools show 496 more exits in the year after closure than the exit rate of staff in schools that stayed open implies, or 1.06 posts per closure. Each post is priced at \$2,475 a month, the median monthly wage of staff in closing schools, times twelve, or \$29,700 a year, which gives \$31,544 per closure. The count comes from the closure waves the staff records observe, which include closures after 2017, and the records cannot tell whether a released post was later refilled; the high end assumes it was not. The wage comes from a payroll file that holds the workforce of 2024 back to 2010. It omits everyone who left before 2024, at least 28.6 percent of the 2013 workforce, so it prices posts at the wages of staff who stayed. The file supplies prices only; no payroll total is read from it.

F.3 Totals and What the Account Leaves Out

For the 315 closures chosen again in Section 8.2, the low-end amount sums to \$3,842,055 a year, the enrollment component to \$5,615,524 and posts to \$9,936,360, so the range runs from \$3,842,055 to \$19,393,939 a year (Table F1). At the high end a closure saves \$43,741 plus \$121.16 per pupil: \$60,582 at the median closed school of 139 pupils and \$61,568 at the mean of 147.1 pupils. Because the high end grows with enrollment, equal counts need not forgo equal savings: keeping open the first 47 closures of “Farthest from a surviving school first” forgoes \$573,259 to \$2,925,029 a year, and the first 47 of “Largest model access gain first,” which are larger schools, forgo the same low end and \$3,226,838 at the high end.

The account leaves out five items. Per-pupil spending outside facilities follows the child to the receiving school, so it is not saved. Spending that stops when children leave the public schools is not a saving from the closure. Transport is not identified in the records, so the savings are gross of any added transport cost, and such a cost would lower the break-even values of Section 8.3. Spending on the transition is not counted. One-time values of closed buildings belong to a capital account and are never added. The range is a range of assumptions, not a confidence interval and not a bound on what the closures saved, because the spending regression does not identify the cost each closure avoided.

Valuing the closures would require more than this account supplies. It would need how long the access

Table F1. Savings per year from a closure: components and totals

	Per closure	All 315 closures
Low end		
Facilities: intercept	12,197	3,842,055
less assumed closed-site deduction		
Added at the high end		
Facilities: 121.16 dollars a pupil	17,827	5,615,524
Staff posts	31,544	9,936,360
High end	61,568	19,393,939

Notes: Dollars a year. Low end: the facilities intercept (\$18,657) less an assumed closed-site deduction (\$6,460), from a regression on 127 surviving schools in 2019–20 applied to earlier closures. High end: adds the enrollment term at each closed school’s last-year enrollment and 1.06 posts per closure; the per-closure column uses the mean closed school (147.1 pupils; 46,348 in all). Posts savings assume released positions are not refilled. Excluded: transport, transition spending, per-pupil costs that follow the child, and capital. The two ends are a range of assumptions, not a confidence interval.

loss lasts and a price for it, and an interpretation and a price for school effects. Comparing lists at a common budget would also need the seats and costs of receiving schools, transport and feasible assignments. The paper supplies none of these, and it compares lists at the Department’s counts, not at equal savings.