

Student Flow Modeling for School Decongestion via Stochastic Gravity Estimation and Constrained Spatial Allocation

Sebastian Felipe R. Bundoc*
Paula Joy B. Martinez*
ecair.sbundoc@deped.gov.ph
ecair.pmartinez@deped.gov.ph
Center for AI Research
Department of Education
Philippines, PH

Sebastian C. Ibañez
Center for AI Research
Department of Education
Philippines, PH
ecair.sibanez@deped.gov.ph

Erika Fille T. Legara
Center for AI Research
Department of Education
Philippines, PH
ecair.elegara@deped.gov.ph

Abstract

School congestion, where student enrollment exceeds school capacity, remains a major challenge in low- and middle-income countries. It highly impacts learning outcomes and deepens inequities in education. While education subsidy programs that transfer students from public to private schools offer a mechanism to alleviate congestion without capital-intensive construction, they often underperform due to fragmented data systems that hinder effective implementation. The Philippine Educational Service Contracting program, one of the world’s largest educational subsidy programs, exemplifies these challenges, falling short of its goal to decongest public schools. This data fragmentation prevents the science-based and data-driven analyses needed to understand what shapes student enrollment flows, particularly how families respond to economic incentives and spatial constraints. In this work, we introduce a computational framework for modeling student flow patterns and simulating policy scenarios. By synthesizing heterogeneous government data across nearly 3,000 institutions, including school geolocations, infrastructure records, enrollment statistics, school tuition fees, student flows, road networks, subsidy amounts and slots, we employ a stochastic gravity model estimated via negative binomial regression to derive behavioral elasticities for distance, net tuition cost, and socioeconomic determinants. These elasticities inform a doubly constrained spatial allocation mechanism that simulates student redistribution under varying subsidy amounts while respecting both origin candidate pools and destination slot capacities. We find that geographic proximity constrains school choice four times more strongly than tuition cost, and that slot capacity—not subsidy amounts—is the binding constraint. Our work demonstrates that subsidy programs alone cannot resolve systemic overcrowding, and computational modeling can empower education policymakers to make equitable, data-driven decisions by revealing the structural

constraints that shape effective resource allocation, even when resources are limited.

CCS Concepts

• **Computing methodologies** → **Discrete-event simulation**;
• **Supervised learning by regression**; • **Applied computing** → **Education**.

Keywords

school choice, stochastic gravity modeling, constrained spatial allocation, policy simulation

ACM Reference Format:

Sebastian Felipe R. Bundoc, Paula Joy B. Martinez, Sebastian C. Ibañez, and Erika Fille T. Legara. 2026. Student Flow Modeling for School Decongestion via Stochastic Gravity Estimation and Constrained Spatial Allocation. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '26)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 Introduction

School congestion, defined as enrollment exceeding a school’s physical and pedagogical capacity, remains a persistent barrier to global education quality [5, 8]. In Lower- and Middle-Income Countries (LMICs), this challenge is particularly acute, with student-teacher ratios often doubling global averages [27]. In the Philippines, over 5.1 million students are enrolled in overcrowded classrooms [23]. To address this, the Department of Education utilizes the Educational Service Contracting (ESC) program—one of the world’s largest public-private partnerships in education [29]—to provide tuition subsidies that incentivize the transition of Grade 7 students from overextended public schools to participating private institutions [9] nationwide.

Despite the program’s scale, the Department’s internal capacity to measure its impact on systemic decongestion has historically been limited. To resolve this, our team, together with the Government Assistance and Subsidies Services Office, was tasked with transforming how the Department understands and relieves public school overcrowding. This paper documents the computational policy framework that we developed to address the systemic barriers the Department of Education encounters. Through our experience in administering the ESC, it became evident that our primary obstacles were fragmented information systems and analytical frameworks that failed to account for necessary complexity. First, our critical administrative data on school capacity, spatial accessibility,

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
KDD '26, Jeju, Korea

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2026/08
<https://doi.org/XXXXXXXX.XXXXXXX>

and student mobility were siloed across disparate bureaus and other national government agencies, preventing a unified view of the system. Second, traditional policy analysis relied on linear, aggregate models that failed to capture the nonlinear interactions between student flows, spatial constraints, and economic incentives.

Our computational framework integrates heterogeneous administrative data from nearly 8,000 schools with open-source geospatial data and government data from other national government agencies to construct a comprehensive representation of student flows in the Philippine basic education system. Our contribution lies in using a stochastic gravity-based model to implement counterfactual decongestion and perform policy simulations. By framing student mobility within a gravity approach, we characterize the interaction between origins and destinations as a function of mass (school capacity and candidate beneficiaries) and friction (e.g., road distance and net tuition costs). These parameters are then integrated into a doubly constrained spatial allocation mechanism that simulates how student flows shift across the network when we adjust subsidy amounts.

Our work makes three primary contributions to the science of educational resource allocation. First, we demonstrate how harmonizing fragmented administrative datasets with open-source data and records from other national government agencies enables system-wide visibility into student flow networks that were previously opaque to both policymakers and researchers. This yields the first origin-destination flow graph of the Philippine basic education system at national scale. Second, through a gravity-based spatial interaction model, we establish an empirical characterization of school choice behavior among subsidy-eligible families. Geographic accessibility dominates economic cost by a factor of four in determining enrollment flows, a finding that challenges the prevailing policy assumption that subsidy increases are the primary lever for driving private school participation. Third, by coupling this behavioral estimation with a doubly constrained stochastic allocation framework, we uncover structural properties of the system that aggregate statistics alone cannot reveal. We show that ESC enrollment is capacity-bound rather than price-bound, and that over half of the system's untapped decongestion potential lies in latent origin-destination pathways not yet observed in historical data. We apply our computational framework to the three most populous regions of the Philippines, the National Capital Region, Region III, and Region IV-A, which together account for the vast majority of the country's school congestion.

Beyond these empirical discoveries, this work demonstrates how computational methods from spatial science and AI can make the structure of a national educational system accessible to scientific inquiry. Where the Department of Education previously relied on descriptive reporting of historical enrollment, our framework enables prospective simulation of counterfactual policy scenarios. This reveals not only how students would redistribute under alternative subsidy regimes but, more fundamentally, identifies the binding constraints on the system's decongestion potential: physical school capacity rather than tuition cost. While developed for the Philippine basic education context, the methodology generalizes to spatial allocation problems in which behavioral models of individual choice must be reconciled with physical capacity constraints

across a network of facilities, including healthcare access, public transit planning, and social service delivery.

2 Related Literature

2.1 Philippine Educational Service Contracting

The ESC program was established in 1998 to reduce public school congestion by providing tuition subsidies for elementary graduates to enroll in partner private secondary schools [9]. It has grown exponentially over the years, serving nearly one million students through collaboration with over 3,000 private schools by the 2017–2018 school year [20]. However, despite the program's substantial scale, fundamental weaknesses in its design and governance are well-documented, particularly with respect to its two levers of control: subsidy slots and subsidy amount. A World Bank [29] evaluation discovered that the current distribution of grantees by region is not proportional with the estimated number of excess students in each region, resulting in some regions receiving disproportionately larger shares while others were underserved. A local study also found issues with the distribution mechanism at the school level [20]. With respect to subsidy amounts, [7] revealed significant inadequacies in current grant amounts, which have remained unchanged since 2017 and are merely based on the geographic location of the participating private school—those located in the capital region receive the highest amount, while schools established in more provincial regions are granted the lowest. While existing literature identifies the misallocation of slots and the inadequacy of subsidies, it often treats these factors as isolated administrative issues. We argue that these two levers, subsidy slots and amounts, must be modeled as interdependent components of a redistribution mechanism that analyzes student flow patterns, identifies underutilized private sector capacity, and systematically reallocates students from congested public schools to available private seats.

2.2 Modeling Student Flow

We model student flows using a framework analogous to Newton's law of gravitation. The gravity model posits that flows between origin i and destination j are proportional to origin "mass" and destination "mass" and inversely proportional to the distance between their centers of mass. This framework has been widely used across diverse domains such as in transportation and urban planning [17, 22, 25, 26, 31], as well as in education, particularly in understanding student migration patterns [4, 6, 15, 18, 21]. However, unlike physical gravity, which holds exactly, economic gravity models are approximations requiring stochastic specifications [24]. Individual decisions depend on many unmeasured factors, including personal preferences, family connections, information availability, and random events that create deviations from theoretical predictions. The stochastic gravity equation takes the form $F_{ij} = \alpha_0 O_i^{\alpha_1} D_j^{\alpha_2} d_{ij}^{\alpha_3} \eta_{ij}$, where η_{ij} is a multiplicative error term with $E(\eta_{ij} | \text{regressors}) = 1$, assumed to be statistically independent of the regressors.

Traditionally, gravity models were estimated by ordinary least squares on log-transformed flows. However, Silva and Tenreiro [24] demonstrate that heteroskedasticity renders OLS estimates biased and inconsistent, even after controlling for fixed effects. They proposed the Poisson Pseudo-Maximum Likelihood estimation, which

handles heteroskedasticity consistently and accommodates zero flows. This approach has since become standard practice in spatial interaction modeling [16].

Because student enrollment flows emerge from individual school choice decisions, the gravity model can also be viewed through the lens of microeconomic choice behavior. McFadden [19] developed the seminal multinomial logit model based on random utility maximization, originally applied to travel behavior analysis but since widely adopted in education contexts [1, 3, 10, 12]. Under this framework, the aggregate student counts we observe are the realized outcomes of individual utility-maximizing choices. Following the mathematical framework established by [11], when the number of alternatives in a choice set is large, as is the case with thousands of origin-destination pairs in our dataset, a Poisson-based regression becomes mathematically equivalent to a conditional logit model. This equivalence allows us to treat the gravity equation as a computationally tractable method to analyze individual choice behavior.

2.3 Constrained Spatial Allocation

The mathematical foundation for redistributing students across a network originates in the transportation problem of linear programming [13, 14], which optimizes resource flows from supply points (origin schools) to demand points (destination schools) while minimizing a cost function, typically transportation costs. Modern extensions of this logic combine machine learning-based discrete choice models with stochastic optimization to predict student enrollment under different zoning plans [10], while others emphasize that reserving seats for specific assignment rules can reduce strategic misrepresentation and improve equity [30]. Our algorithm follows this logic by treating the candidate pool at origin schools as the volume to be allocated and the unutilized slots at destination schools as the system's capacity constraints. We extend the classical framework by using the OD-specific behavioral elasticities derived from the gravity model to weight these allocations, rather than assuming fixed or uniform preferences across all school pairs.

This redistribution mechanism aligns further with Wilson [28]'s framework for doubly constrained spatial interaction. Unlike production-constrained models that only fix total movers from an origin, doubly constrained models introduce balancing factors to ensure that destination totals do not exceed known capacities, represented in this paper as the available subsidy slots. We then incorporate lexicographical searches to identify new, viable origin-destination pairs. This approach is grounded in the literature on non-compensatory decision-making and student-optimal matching theories, which suggests that households do not weigh all school attributes simultaneously but rather filter options through a hierarchy of priorities, where distance is often the highest [3].

3 Methodology

3.1 Data Sources

This study was conducted through an internal collaboration between the Center for AI Research and the Government Assistance and Subsidies Service Office of the Department of Education in the Philippines. The datasets used span school years 2023-2024 and

2024-2025, and were sourced from multiple information systems within the Department.

3.1.1 Public and private school coordinates. School coordinates enable spatial analysis to identify nearby private school alternatives when redirecting incoming Grade 7 students away from overcrowded public high schools. Public school coordinates are obtained from the national school registry, while private school coordinates are compiled from the administrative records of the country's 18 regional field offices. Preprocessing standardizes coordinate formats across sources and filters out records lacking valid coordinates or school identifiers, yielding 47,821 public and 11,831 private school locations for geospatial matching.

3.1.2 Inter-school road network distances. Road network distances between schools are computed using the Open Source Routing Machine, a routing engine built on OpenStreetMap data. OSRM's Multi-Level Dijkstra algorithm enables efficient computation of pairwise distances for all 13,106 schools within the three largest regions in the Philippines—the National Capital Region and two adjacent regions—yielding a distance matrix of over 171 million school pairs. Unlike straight-line distances, road network distances account for actual travel paths, providing more realistic estimates of geographic accessibility for student redirection.

3.1.3 School-level enrollment. School-level enrollment data provides aggregate student counts for each school, disaggregated by grade level. Sourced from the national enrollment database and spanning three school years, the dataset covers approximately 60,000 public and private schools annually. High school enrollment figures are particularly critical: when compared against seating capacity, they reveal which public schools are operating beyond capacity and which private schools have room to absorb additional students.

3.1.4 Public and private school seats data. Seating capacity data quantifies the maximum number of students each school can physically accommodate. Public school capacity is derived from official seat-learner ratio records, while private school capacity is compiled from classroom furniture inventories. Schools where high school enrollment exceeds seating capacity are flagged as congested, with the excess enrollment termed "aisle learners" or students who lack dedicated seats. This metric directly informs the choice and redirection models by quantifying overcrowding at each public school and the absorptive capacity at nearby private schools.

3.1.5 ESC beneficiary records. Beneficiary records from the Education Service Contracting (ESC) program identify students currently receiving tuition subsidies at participating private schools. Spanning three school years and totaling approximately 2.7 million individual records, this dataset enables the model to distinguish between existing beneficiaries and non-beneficiaries within observed student flows. Since beneficiaries are already enrolled in private schools through the subsidy program, only non-beneficiaries, or students flowing toward congested public schools, are candidates for redirection to private alternatives with unutilized subsidy slots.

3.1.6 ESC program data. Program-level data describes the resources and constraints governing the ESC subsidy. Tuition fee records indicate what participating private schools charge, while mandated

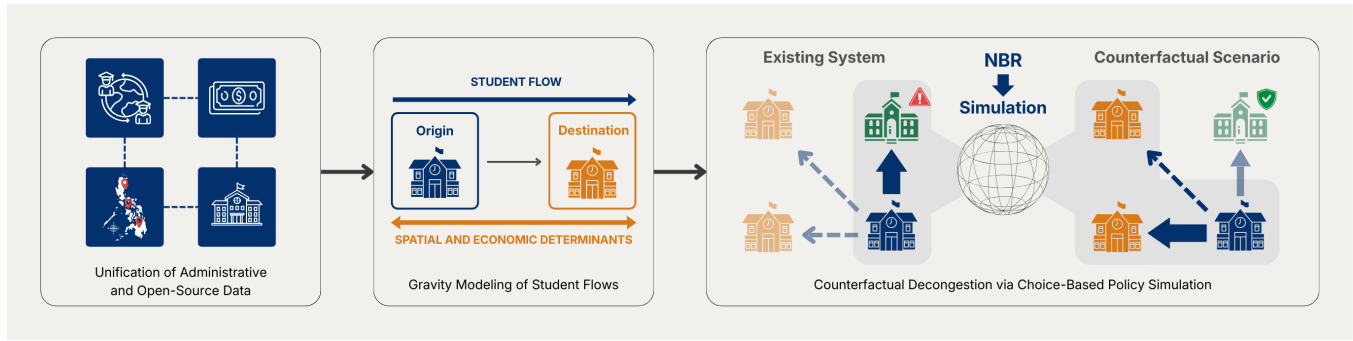


Figure 1: Overview of the computational policy framework for student redistribution. Institutional data from multiple national government agencies and open-source geospatial datasets are unified to represent the national school network. A stochastic gravity model estimated via negative binomial regression derives behavioral elasticities for distance, net tuition cost, and other determinants. Predicted flows are then allocated through a doubly constrained mechanism that iteratively assigns students until candidate pools or slot capacities are exhausted, enabling counterfactual evaluation of subsidy and slot allocation policies.

subsidy amounts specify the government contribution per beneficiary. The gap between the two represents the out-of-pocket burden on families. Slot allocation data specifies how many subsidized seats each school is authorized to fill. Crucially, unutilized slot counts from official program records indicate remaining capacity at each school, serving as the binding constraint in the redistribution algorithm. Finally, certification ratings from the national quality assurance system provide a measure of school performance, enabling the model to prioritize higher-quality alternatives when multiple redirection options exist. These institutional quality assurance ratings were translated into hierarchical ordinal values to enable quantitative comparison of school quality during choice modeling and redirection prioritization.

3.1.7 Municipality Income Data. To account for the socioeconomic determinants of student flow, we integrated municipal-level income data from the Cities and Municipalities Competitiveness Index. Published annually by the Department of Trade and Industry, this dataset provides a standardized framework for evaluating the financial capacity and economic dynamism of Local Government Units (LGUs). We specifically utilized LGU income as a proxy for the socioeconomic status of both origin and destination municipalities. Origin-side income data allows the model to identify the financial capacity of a student’s home municipality, while destination-side income serves as a proxy for the localized cost of living.

The student flow dataset is the primary analytical output of this study, synthesized by merging the administrative records described above. By joining enrollment databases, beneficiary records, and school registries, we constructed a dataset of student transitions from Grade 6 to Grade 7. These origin-destination (OD) records, spanning three school years and totaling approximately 2.7 million individual transitions, form the empirical foundation for our gravity modeling and redistribution simulation.

For each transition, we integrated the road network distance between the origin and destination, the net tuition cost (calculated as the difference between private school fees and the ESC subsidy), and the socioeconomic status of the origin municipality. This unified

view enables the model to differentiate between spatial interactions—geographic movements driven by proximity and historical feeder patterns—and subsidy-responsive flows incentivized by the ESC program.

3.2 The Empirical Model of Student Flows

We employ a stochastic gravity model to analyze student flows across the three largest regions in the Philippines. Individuals are assumed to be rational agents who maximize their utility by evaluating the benefits of a destination against costs, such as travel costs implied through distance. We describe student flows F_{ij} by a general gravity model where the movement from origin i to destination j is a function of origin-side socioeconomic characteristics, destination-side attractiveness, and the spatial separation between them. Traditionally, gravity models are estimated by ordinary least squares (OLS) on log-transformed flows. However, the implications of Jensen’s inequality render OLS estimates biased and inconsistent in the presence of heteroskedasticity, even after controlling for fixed effects [16, 24]. As $E(\ln y) \neq \ln E(y)$, log-linearization of the gravity equation leads to systematic errors in estimating true elasticities. While Poisson models are known to solve this, they only provide correct estimations if the equidispersion theorem applies. Because student transitions exhibit a power law distribution, where a small number of OD pairs account for the majority of flows, the resulting overdispersion violates the equidispersion requirement of a Poisson model, necessitating the use of a Negative Binomial Regression (NBR). A detailed comparison justifying the transition from OLS and Poisson to NBR is provided in Appendix A.

We formalize the stochastic gravity model as:

$$F_{i,j} = \alpha_0 O_i^{\alpha_1} D_j^{\alpha_2} d_{i,j}^{\alpha_3} c_{i,j}^{\alpha_4} \eta_{i,j} \quad (1)$$

$$E(F_{i,j}|O_i, D_j, d_{i,j}, c_{i,j}) = \alpha_0 O_i^{\alpha_1} D_j^{\alpha_2} d_{i,j}^{\alpha_3} c_{i,j}^{\alpha_4} \quad (2)$$

where $F_{i,j}$ is the flow of ESC beneficiaries from origin i to school j , O_i captures origin school characteristics, D_j represents destination school attractiveness, $d_{i,j}$ is road network distance (km), $c_{i,j}$ is net

enrollment cost (tuition minus ESC subsidy, in thousands of pesos), and $\eta_{i,j}$ is a multiplicative error term with $E(\eta_{i,j}|\text{regressors}) = 1$.

Following [2], we also use fixed effects to control for all time-invariant origin and destination factors. Our final equation utilizes a log-link function to maintain the multiplicative nature of the gravity framework while avoiding the biases inherent in manual log-linearization:

$$\begin{aligned} \ln E(F_{i,j}) = & \beta_0 + \gamma_i + \delta_j \\ & + \alpha_3 \ln d_{i,j} + \alpha_4 \ln c_{i,j} \\ & + \beta_1 \ln \text{OriginIncome}_i + \beta_2 \ln \text{DestIncome}_j \\ & + \beta_3 \text{Rating}_j \end{aligned} \quad (3)$$

Where β_0 is the intercept, γ_i and δ_j represent regional fixed effects, and $\ln d_{i,j}$ and $\ln c_{i,j}$ are the log-transformed distance and net cost variables used to estimate the elasticities α_3 and α_4 . The iterative process for selecting these parameters is detailed in Appendix B. Furthermore, the dispersion parameter ($\alpha = 0.3948$) confirms the NBR as the robust choice for this dataset over the restrictive Poisson alternative. To account for the potential intra-cluster correlation of student flows originating from the same school, we employ clustered robust standard errors at the origin school level. This ensures that our inference remains valid in the presence of unobserved local factors that may drive correlated enrollment behaviors among students within the same geographic catchment area.

3.3 Candidate Beneficiary Pool and Policy Simulation

To assess the efficacy of varying policy interventions, we apply the estimated model parameters to simulate the behavioral response of students to changes in the subsidy amount. The subsidy is treated as a direct reduction in the economic constraint c_{ij} , allowing us to model the sensitivity of non-beneficiary flows and predict the resulting student redistribution under different levels of financial support.

A critical component of this simulation is the definition of the candidate beneficiary pool (E_i) for each origin school i . This is the untapped cohort of potential subsidy beneficiaries. While the model is trained on historical beneficiary flows, the policy objective is to project the movement of students who have not yet opted into the ESC program. We define this candidate pool as:

$$E_i = \text{Enrollment}_i - \sum_j \text{Beneficiaries}_{ij} \quad (4)$$

Enrollment_i is the total Grade 6 student population at origin school i , and the summation represents the existing flow of students into the private sector through the subsidy program. This pool ensures that the simulation respects the physical population limits and avoids overestimating decongestion by constraining predicted flows such that $\sum_j \hat{y}_{ij} \leq E_i$. To capture demand beyond historical observations, the simulation expands the set of possible school choices (OD pairs) for each student. This augmented set includes not only established pathways but also counterfactual pairings identified through a lexicographical search of the school network.

3.4 Counterfactual Decongestion via Choice-Based Policy Simulation

The augmented candidate beneficiary pool from Section 3.3 provides, for each origin-destination pair (i, j) , a scenario-specific predicted flow $\hat{y}_{ij}^{(s)}$ representing the expected number of students from origin i who would enroll at ESC school j under policy scenario s . These predictions are unconstrained: they reflect the behavioral response to changes in net cost c_{ij} but do not account for finite school capacity or the competition between ESC destinations for the same origin pool. We now introduce a simulation framework that imposes both supply- and demand-side constraints to produce realistic enrollment projections, and then attributes the resulting flows to public school decongestion.

Constrained Allocation under Coupled Depletion. Two capacity constraints bind simultaneously in the ESC system. The unconstrained predictions from Section 3.3 must satisfy a doubly constrained structure: on the demand side, each origin i has a finite candidate pool E_i (Section 3.3), ensuring that $\sum_j a_{ij} \leq E_i$; on the supply side, each ESC school j has a finite number of contracted slots K_j , ensuring that $\sum_i a_{ij} \leq K_j$. Because multiple ESC schools draw from overlapping origin catchment areas, accepting students along one path reduces availability for others—a coupled depletion structure that cannot be resolved by independent per-school calculations. We formalize this as a stochastic allocation process. For a given scenario s and random seed r , let $\pi^{(r)}$ be a uniformly random permutation of all origin-destination pairs in the candidate beneficiary pool \mathcal{P} . We process pairs sequentially in the order $\pi^{(r)}$, maintaining residual capacities e_i (initialized to E_i) and k_j (initialized to K_j). For each pair (i, j) , the accepted flow is:

$$a_{ij}^{(s,r)} = \min(e_i, k_j, \hat{y}_{ij}^{(s)}) \quad (5)$$

after which both residual capacities are updated: $e_i \leftarrow e_i - a_{ij}^{(s,r)}$ and $k_j \leftarrow k_j - a_{ij}^{(s,r)}$. The iteration terminates early when either total remaining candidates or total remaining slots reaches zero. The random permutation $\pi^{(r)}$ determines which paths are prioritized when capacity is scarce; by averaging over $R = 100$ independent seeds ($r = 0, \dots, 99$), we obtain Monte Carlo means and 95% confidence intervals that capture the allocation uncertainty due to processing order. The total simulated enrollment at ESC school j under scenario s and seed r is $Y_j^{(s,r)} = \sum_i a_{ij}^{(s,r)}$.

Decongestion Attribution. Not all ESC enrollment contributes to public school decongestion—only students diverted from origins that also feed congested public schools. Let \mathcal{G} denote the set of congested public junior high schools and let $\mathcal{O}^* = \{i : \exists g \in \mathcal{G} \text{ s.t. } F_{ig} > 0\}$ be the set of *congested-feeding origins*, i.e., elementary schools whose graduates are observed to flow to at least one school in \mathcal{G} . For each ESC school j , we define the congested fraction:

$$\phi_j = \frac{\sum_{i \in \mathcal{O}^*} \hat{y}_{ij}}{\sum_i \hat{y}_{ij}} \quad (6)$$

representing the share of j 's predicted demand originating from congested-feeding schools, computed from unconstrained model predictions as a pre-simulation approximation. We then define two complementary decongestion measures for each ESC school j .

The *total predicted decongestion role* captures the full congestion-relevant simulated flow:

$$D_j^{\text{total}(s,r)} = Y_j^{(s,r)} \cdot \phi_j \quad (7)$$

The *marginal decongestion potential* isolates untapped capacity beyond current observed enrollment $b_j = \sum_i F_{ij}$:

$$D_j^{\text{marg}(s,r)} = \max(0, Y_j^{(s,r)} - b_j) \cdot \phi_j \quad (8)$$

The total measure quantifies the ESC system’s structural role in absorbing students who would otherwise attend congested public schools; the marginal measure isolates the additional relief achievable if all contracted slots were fully utilized.

Existing and Hypothetical Path Decomposition. The augmented candidate beneficiary pool (Section 3.3) includes both historically observed pathways and counterfactual pairings identified through a lexicographical search of the school network. We classify each pair $(i, j) \in \mathcal{P}$ as *existing* if at least one beneficiary was observed on that route ($F_{ij} > 0$), and *hypothetical* otherwise. Let $\mathcal{E} \subset \mathcal{P}$ denote existing paths and $\mathcal{H} = \mathcal{P} \setminus \mathcal{E}$ the hypothetical paths. We decompose each school’s simulated enrollment accordingly:

$$Y_j^{(s,r)} = \underbrace{\sum_{i:(i,j) \in \mathcal{E}} a_{ij}^{(s,r)}}_{\text{existing}} + \underbrace{\sum_{i:(i,j) \in \mathcal{H}} a_{ij}^{(s,r)}}_{\text{hypothetical}} \quad (9)$$

This decomposition carries through to both decongestion measures and distinguishes gains achievable by scaling up current enrollment patterns from those requiring the activation of new routes not yet observed in practice.

Counterfactual Scenarios. We evaluate five counterfactual scenarios corresponding to net cost reductions of {1K, 5K, 10K, 15K, 20K} Philippine pesos per student, which modify the cost term c_{ij} in the gravity model (Section 3.2) and thereby adjust the predicted flows $\hat{y}_{ij}^{(s)}$. The slot capacity K_j and candidate pools E_i are held constant across scenarios, isolating the demand-side effect of subsidy changes under fixed supply constraints. This design allows us to disentangle whether the binding constraint on ESC-mediated decongestion is the price of access or the physical supply of school places.

4 Results

4.1 Student Flow Modeling Results

Table 1 presents the estimation result of the gravity model using negative binomial regression. The model employs a log-log specification for distance and cost, allowing coefficients to be interpreted as elasticities.

Distance Elasticity. The distance decay effect is highly significant. This implies that a 10% increase in road distance between origin and destination schools results in approximately a 4.5% decrease in expected student flows. Equivalently, doubling the distance reduces expected enrollment by approximately 27% ($= 1 - 2^{-0.45}$). Distance has the strongest effect in the model, indicating that geographic accessibility is the primary barrier to private school enrollment among subsidy-eligible families.

Cost Elasticity. Net enrollment cost (tuition – subsidy) also exhibits a statistically significant negative effect. This indicates that a 10% reduction in out-of-pocket tuition cost increases expected

Table 1: Negative Binomial Regression Results

Variable	Coef.	Std. Error	z-value	p-value
Intercept	3.3944	0.130	26.121	0.000
ln(distance)	-0.4509	0.010	-46.488	0.000
ln(net cost)	-0.1004	0.013	-7.582	0.000
ESC School Rating	-0.0204	0.008	-2.503	0.012
ln(origin LGU income)	-0.0207	0.004	-5.053	0.000
ln(dest. LGU income)	-0.0489	0.004	-11.406	0.000
<i>Origin Region FE (Reference: NCR)</i>				
Region III	-0.0205	0.026	-0.781	0.435
Region IV-A	-0.0500	0.021	-2.438	0.015
<i>Destination Region FE (Reference: NCR)</i>				
Region III	-0.0237	0.026	-0.911	0.362
Region IV-A	0.0178	0.019	0.921	0.357
Dispersion (α)	0.3925	0.011	35.737	0.000
<i>Observations: 29,224</i>				
<i>Pseudo R² (McFadden): 0.08625</i>				
<i>Log-Likelihood: -53,592</i>				

Note: Standard errors are clustered at the origin school level. Distance and cost are log-transformed, representing elasticities. All income variables are LGU-level revenues.

student flows by approximately 1.1%. Doubling the net cost would decrease enrollment by roughly 7.4% ($= 0.5^{-0.11} - 1$). Because the magnitude of this cost elasticity is notably smaller than the distance elasticity (by nearly a factor of 4), it can be said that subsidies alone have limited effectiveness in driving enrollment shifts if students lack access to nearby private schools. This suggests that geographic expansion of ESC-participating schools may be more impactful than purely increasing subsidy amounts for geographically constrained families.

Destination School Quality. The ESC rating coefficient indicates that higher-rated schools receive fewer ESC beneficiaries. A post-hoc analysis revealed that higher-rated schools charge premium tuition that exceeds ESC subsidy levels, making them financially inaccessible despite higher ratings. Alternatively, top-rated schools may also be capacity-constrained and selective, limiting ESC beneficiary intake. A third possibility is that the ESC rating, which is based on DepEd compliance and administrative criteria, may not align with quality dimensions families prioritize, such as proximity, transportation availability, or peer composition. The small magnitude ($\beta = -0.02$) suggests this effect, while statistically significant, has limited practical importance compared to distance and cost factors.

Origin LGU Income Effects. The scaled origin income coefficient is negative and highly significant ($\beta = -0.0524$, $p < 0.001$), indicating that students from higher-income municipalities are less likely to utilize the subsidy program. This could indicate that the ESC program is successfully reaching its intended beneficiaries, the students from lower-income communities where public school quality gaps are most pronounced. The magnitude of this income effect is comparable to the cost elasticity, suggesting that origin municipality wealth plays a substantive role in enrollment decisions.

Destination LGU Income Effects. The negative coefficient for destination LGU income suggests that students avoid schools in the wealthiest municipalities. This implies that hidden costs, such as daily meals, school supplies, and other miscellaneous costs act as a barrier even when tuition is subsidized. Students from modest backgrounds may avoid schools in wealthy areas because of the financial burden of higher cost-of-living standards.

Regional Fixed Effects. The regional fixed effects reveal substantial spatial heterogeneity in subsidy participation patterns. On the demand side, students from Region III and Region IV-A exhibit lower participation per origin-destination pair compared to NCR students, controlling for distance, cost, and income. However, in absolute terms, Region IV-A generates the largest share of ESC beneficiaries (44.5% of total flows), reflecting its larger population base.

On the supply side, destination schools in Region III and Region IV-A enroll more ESC students per school than NCR schools, all else equal. This pattern manifests in aggregate flows, where NCR experiences a net outflow of students while Region IV-A receives a net inflow. This demand-supply mismatch suggests that NCR private schools face binding capacity constraints or strategically limit ESC participation despite higher per-capita demand from NCR families, forcing students to seek ESC slots in neighboring regions at additional cost.

Overdispersion The estimated dispersion parameter ($\alpha = 0.3948$, $p < 0.001$) confirms substantial overdispersion in the data, with variance significantly exceeding the mean. This validates the use of negative binomial regression over standard Poisson, ensuring appropriate handling of flow heterogeneity.

4.2 Policy Simulation Results

Drawing on the empirical finding that geographic proximity serves as a significantly stronger driver of school choice than tuition cost, we evaluate the system-wide impact of these behavioral elasticities by integrating our stochastic gravity model predictions into a constrained spatial allocation framework, and measure the resulting decongestion under five subsidy scenarios.

Table 2 reports the Monte Carlo simulation results across all scenarios, with observed ESC beneficiary enrollment as a baseline. Two comparisons are presented: the gap between observed and simulated enrollment (reflecting untapped slot capacity), and the change across subsidy scenarios (reflecting the demand-side effect of cost reductions).

The slot ceiling dominates price sensitivity.

The gap between observed enrollment and simulated enrollment is substantial. Under even the most conservative scenario (−1K), the model predicts approximately 26,000 additional students would fill ESC slots if all contracted capacity were utilized, a 34.7% increase over current enrollment. Yet varying the subsidy from a 1,000-peso to a 20,000-peso net cost reduction increases predicted flow by only 1,826 students, or 1.8%. This near-flat response is a direct consequence of the structural imbalance in the system: 562,958 candidates compete for 106,654 contracted slots, a demand-to-supply ratio of 5.3×. At this level of oversubscription, reducing price reshuffles which students fill slots but barely changes how many slots are filled.

Table 2: Constrained Simulation Results Across Subsidy Scenarios

Scenario	Net Cost Reduction	Predicted Flow $\bar{Y}^{(s)}$	Δ from Observed	Δ from −1K
Observed (b_j)	—	74,232	—	—
−1K	P1,000	99,992	+34.7%	—
−5K	P5,000	100,442	+35.3%	+0.4%
−10K	P10,000	100,944	+36.0%	+1.0%
−15K	P15,000	101,428	+36.6%	+1.4%
−20K	P20,000	101,818	+37.2%	+1.8%

Note: Predicted flow is the Monte Carlo mean over $R = 100$ iterations. System-level standard deviation is ~ 11 students ($< 0.01\%$ of mean), indicating stable convergence. Observed baseline is total ESC beneficiaries from student flow records. All scenarios use the same slot capacity K_j and candidate pool E_i .

Decongestion is geographically concentrated. We report the decongestion decomposition for the −1K scenario, which is representative of all scenarios given the minimal cross-scenario variation. Of the 99,992 predicted enrollees, 95,695 (95.7%) flow through congested-feeding origins—elementary schools whose graduates also attend congested public junior high schools. The marginal decongestion potential, $\sum_j D_j^{\text{marg}}$, totals 24,820 students, of which 98.1% originates from congested-feeding schools. In other words, virtually all of the ESC system’s untapped absorption capacity is geographically positioned to relieve the most overcrowded public schools. Among the 1,373 ESC schools in the system, 1,057 exhibit positive marginal capacity (predicted enrollment exceeding current enrollment), while 284 show negative marginals—of which 185 are structurally over-enrolled (current enrollment already exceeds contracted slots) and 99 reflect model under-prediction.

Over Half of Untapped Potential Lies in Unobserved Pathways. The existing-hypothetical decomposition (Equation 9) reveals that 53.9% of marginal decongestion flows through hypothetical paths—origin-destination pairs where no beneficiary is currently observed. The remaining 46.1% flows through existing paths that the model predicts could absorb additional students beyond current levels. This split suggests that realizing the full decongestion potential of the ESC system requires not only expanding enrollment along established routes but also activating new pathways between congested-feeding origins and nearby ESC schools that students have not yet considered.

Decongestion is meaningful but bounded. To contextualize the magnitude of the ESC system’s decongestion role, we compare it against the scale of public school congestion. A total of 416,458 students currently flow to the 1,254 congested public junior high schools, representing 23.3% of all Grade 7 transitions. The model’s total predicted congestion-relevant flow of 95,695 represents a substantial structural counterweight to this pressure. However, the marginal component—the 24,820 additional students absorbable through full slot utilization—accounts for approximately 6% of the total congested flow. Full utilization of existing ESC capacity would meaningfully reduce pressure on the most crowded public schools but would not, on its own, resolve the congestion problem.

These findings indicate that the primary policy lever for ESC-mediated decongestion is not the subsidy amount but the supply of school places. Expanding the number of contracted slots, or reallocating existing slots from surplus to deficit areas, would directly raise the slot ceiling that currently binds the system. However, these projections assume that all contracted capacity is operationally accessible, an assumption we revisit in the next section.

5 Limitations and Ethical Considerations

Our behavioral estimation assumes that historical enrollment patterns serve as a reliable proxy for future preferences. This assumption does not account for temporal shifts in school reputation or sudden socioeconomic shocks that may occur between academic years, which could significantly alter the predictive accuracy of the model. Furthermore, while our lexicographical redistribution algorithm mirrors the hierarchy of choice observed in our negative binomial results—prioritizing geographic proximity and net tuition cost—real-world household decision-making often involves qualitative factors not captured in administrative datasets. These factors include religious affiliation, specialized curricula, or family legacy, all of which may override the economic and spatial friction modeled in our gravity framework.

A critical operational limitation of the current redistribution algorithm is the assumption that private institutions can instantaneously absorb students up to their authorized capacity. In practice, schools may face internal bottlenecks, such as teacher shortages or classroom maintenance, that temporarily reduce their functional capacity below the officially recorded limit. Consequently, our computational framework should be treated as a decision support tool for human policymakers rather than an automated assignment system. We emphasize that these simulations must be integrated with rigorous auditing and verification of school facilities to ensure the integrity of simulated outcomes, allowing for manual overrides based on ground-level intelligence.

Looking forward, a future direction for this research involves the dynamic calibration of subsidy slots. While this study treats the ESC slot ceiling as a fixed institutional constraint, future iterations could transition toward a responsive system that adjusts functional capacity based on real-time operational readiness and facility audit data. Additionally, the gravity model’s predictive engine could be used to recommend the optimal geographic redistribution of authorized slots, ensuring that private school capacity is strategically expanded or reallocated to high-congestion corridors where the behavioral pull is strongest.

This study was conducted under a formal institutional partnership with the Government Assistance and Subsidies Service Office of the Department of Education. To ensure strict compliance with the Philippine Data Privacy Act of 2012, all sensitive administrative records were de-identified and aggregated into origin-destination pairs prior to analysis. No personally identifiable information (PII) was accessed or stored during the modeling process, ensuring that the pursuit of system-wide visibility remains consistent with our mandate to protect student privacy.

6 Conclusion

This work establishes a computational policy framework for addressing school congestion in the Philippine basic education system through data-driven student flow modeling and constrained spatial allocation. By harmonizing fragmented administrative records across multiple government agencies with open-source geospatial data, we constructed the first comprehensive representation of nationwide student mobility patterns, enabling systematic evaluation of subsidy interventions that was previously operationally infeasible.

Three empirical patterns fundamentally reshape how the Department of Education should approach decongestion policy. First, behavioral elasticities reveal that geographic accessibility constrains choice more severely than tuition barriers. Second, supply constraints bind the system more tightly than demand-side incentives. Contracted slot capacity limits absorption to 106,654 students despite 562,958 eligible candidates, and increasing subsidy amounts from P1,000 to P20,000 yields marginal enrollment gains of only 1.8%, while full slot utilization could absorb 34.7% more students. Third, untapped decongestion potential concentrates in unactivated pathways: 53.9% of marginal capacity flows through origin-destination pairs where no beneficiary currently exists, suggesting that strategic outreach and slot reallocation could yield substantial relief without additional infrastructure expenditure.

These findings impose architectural constraints on policy design. Subsidy programs cannot overcome spatial mismatches between congested public schools and available private capacity. That 95.7% of the predicted enrollees flow through congested-feeding origins, yet the marginal impact is only 6%, confirms that demand-side interventions alone cannot resolve systemic overcrowding. Supply-side expansion through targeted slot increases in high-congestion areas represents the most effective intervention.

7 GenAI Disclosure

The authors used large language models to assist with grammar checking in the manuscript and code debugging. All AI-generated content was reviewed, verified, and revised by the authors. The authors take full responsibility for the accuracy and integrity of the research.

References

- [1] Atila Abdulkadiroğlu, Parag A. Pathak, Jonathan Schellenberg, and Christopher R. Walters. 2020. Do Parents Value School Effectiveness? *American Economic Review* 110, 5 (May 2020), 1502–1539. doi:10.1257/aer.20172040
- [2] James E Anderson and Eric van Wincoop. 2003. Gravity with Gravitas: A Solution to the Border Puzzle. *American Economic Review* 93, 1 (Feb. 2003), 170–192. doi:10.1257/000282803321455214
- [3] Amel Awadelkarim, Arjun Seshadri, Itai Ashlagi, Irene Lo, and Johan Ugander. 2023. Rank-heterogeneous Preference Models for School Choice. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*. ACM, 47–56. doi:10.1145/3580305.3599484
- [4] S. Bacci and B. Bertaccini. 2020. Assessment of the University Reputation Through the Analysis of the Student Mobility. *Social Indicators Research* 156, 2–3 (March 2020), 363–388. doi:10.1007/s11205-020-02322-x
- [5] Peter Barrett, Alberto Treves, Tigran Shmis, Diego Ambasz, and Maria Ustinova. 2018. *The Impact of School Infrastructure on Learning: A Synthesis of the Evidence*. World Bank. doi:10.1596/978-1-4648-1378-8
- [6] Massimiliano Bratti and Stefano Verzillo. 2019. The ‘gravity’ of quality: research quality and the attractiveness of universities in Italy. *Regional Studies* 53, 10 (Feb. 2019), 1385–1396. doi:10.1080/00343404.2019.1566701
- [7] Don Brodeth. 2025. *Strengthening Educational Service Contracting (ESC) Through Targeted Support and Adjusted Grant Structures*. Technical Report. Taft Consulting

- Group. Published 06/10/2025.
- [8] Jesús M. Carro and Pedro Gallardo. 2023. Effect of class size on student achievement in the COVID-19 “new normal”. *Bulletin of Economic Research* 76, 2 (Oct. 2023), 303–318. doi:10.1111/boer.12426
- [9] Department of Education. 2017. Guidelines on the Implementation of the Educational Service Contracting and Teachers’ Salary Subsidy Programs in Junior High School Effective School Year 2017–2018. https://peac.org.ph/wp-content/uploads/2017/09/DO_s2017_020.pdf. DepEd Order No. 20, s. 2017. Accessed: 2025-07-18.
- [10] Hongzhao Guan, Nabeel Gillani, Tyler Simko, Jasmine Mangat, and Pascal Van Hentenryck. 2025. Contextual Stochastic Optimization for School Desegregation Policymaking. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 27 (April 2025), 28024–28032. doi:10.1609/aaai.v39i27.35020
- [11] Paulo Guimarães, Octávio Figueirdo, and Douglas Woodward. 2003. A Tractable Approach to the Firm Location Decision Problem. *Review of Economics and Statistics* 85, 1 (Feb. 2003), 201–204. doi:10.1162/003465303762687811
- [12] Justine S. Hastings, Thomas J. Kane, and Douglas O. Staiger. 2008. Heterogeneous Preferences and the Efficacy of Public School Choice. <https://api.semanticscholar.org/CorpusID:17263864>
- [13] Frank L. Hitchcock. 1941. The Distribution of a Product from Several Sources to Numerous Localities. *Journal of Mathematics and Physics* 20, 1–4 (April 1941), 224–230. doi:10.1002/sapm1941201224
- [14] Tjalling C. Koopmans. 1949. Optimum Utilization of the Transportation System. *Econometrica* 17 (July 1949), 136. doi:10.2307/1907301
- [15] Zsolt Tibor Kosztyán, Vivien Valéria Csányi, Zsuzsanna Banász, Ákos Jakobi, Ildikó Neumanné-Virág, and András Telcs. 2021. The role of higher education in spatial mobility. *Applied Network Science* 6, 1 (Nov. 2021). doi:10.1007/s41109-021-00428-w
- [16] Mario Larch, Serge Shikher, and Yoto V. Yotov. 2025. Estimating Gravity Equations: Theory Implications, Econometric Developments, and Practical Recommendations. *Review of International Economics* 33, 5 (June 2025), 1066–1092. doi:10.1111/roie.12789
- [17] Joshua J. Lewer and Hendrik Van den Berg. 2008. A gravity model of immigration. *Economics Letters* 99, 1 (April 2008), 164–167. doi:10.1016/j.econlet.2007.06.019
- [18] Marie-Louise Litmeyer, Philipp Gareis, and Stefan Hennemann. 2023. Comparing student mobility pattern models. *European Journal of Geography* 14, 1 (March 2023), 21–34. doi:10.48088/ejg.m.lit.14.1.21.34
- [19] Daniel McFadden. 1974. The measurement of urban travel demand. *Journal of Public Economics* 3, 4 (Nov. 1974), 303–328. doi:10.1016/0047-2727(74)90003-6
- [20] Dina Joana Saltiva Ocampo and Kathrina Lorraine Lucasan. 2019. *Key Issues in Governance, Finance, School Improvement, and ICT in Basic Education*. Technical Report Discussion Paper Series No. 2019-13. Philippine Institute for Development Studies (PIDS). <https://pidswebs.pids.gov.ph/CDN/PUBLICATIONS/pidsdps1913.pdf> PIDS Discussion Paper Series.
- [21] Tatiane Pelegrini, Carla Sá, and Marco Túlio Aniceto França. 2022. Factors associated with the mobility of college students in Brazil: an analysis using a gravity model. *Higher Education* 85, 1 (March 2022), 203–223. doi:10.1007/s10734-022-00829-5
- [22] Damiaan Persyn and Wouter Torfs. 2015. A gravity equation for commuting with an application to estimating regional border effects in Belgium. *Journal of Economic Geography* 16, 1 (March 2015), 155–175. doi:10.1093/jeg/lbv003
- [23] Second Congressional Commission on Education. 2025. Fixing the Foundations: A Matter of National Survival. <https://edcom2.gov.ph/media/2025/01/EDCOM-2-Year-2-Report-Fixing-the-Foundations-2025.pdf>. Accessed: 2025-07-18.
- [24] J. M. C. Santos Silva and Silvana Tenreyro. 2006. The Log of Gravity. *The Review of Economics and Statistics* 88, 4 (Nov. 2006), 641–658. doi:10.1162/rest.88.4.641
- [25] Filippo Simini, Gianni Barlacchi, Massimiliano Luca, and Luca Pappalardo. 2021. A Deep Gravity model for mobility flows generation. *Nature Communications* 12, 1 (Nov. 2021). doi:10.1038/s41467-021-26752-4
- [26] Theodore Tsekeris and Antony Stathopoulos. 2006. Gravity models for dynamic transport planning: Development and implementation in urban networks. *Journal of Transport Geography* 14, 2 (2006), 152–160. doi:10.1016/j.jtrangeo.2005.06.009
- [27] UNESCO. 2023. World Teachers’ Day: UNESCO sounds the alarm on global teacher shortage crisis. <https://www.unesco.org/en/articles/world-teachers-day-unesco-sounds-alarm-global-teacher-shortage-crisis>. Accessed: 2025-07-18.
- [28] A G Wilson. 1971. A Family of Spatial Interaction Models, and Associated Developments. *Environment and Planning A: Economy and Space* 3, 1 (March 1971), 1–32. doi:10.1068/a030001
- [29] World Bank. 2011. *Philippines Private Provision, Public Purpose*. Washington, DC. doi:10.1596/27406
- [30] Zhiyi Xu and Robert G. Hammond. 2023. Designing school choice mechanisms: A structural model and demand estimation. *Economic Inquiry* 62, 2 (Dec. 2023), 505–524. doi:10.1111/ecin.13193
- [31] Lan Zhang and Kaijian Liu. 2025. Modeling Purposes of Public Transportation Trips for Human Need-Responsive Urban Mobility Efficiency. *IEEE Access* 13 (2025), 124413–124428. doi:10.1109/ACCESS.2025.3587994

A Model Comparison with OLS and Poisson

To further justify the selection of the Negative Binomial Regression framework, we compared its performance against two traditional benchmarks: an Ordinary Least Squares (OLS) model on log-transformed flows and a standard Poisson regression.

Table 3: Comparison of Count-Based Models

Metric	Poisson Regression	Negative Binomial
Log-Likelihood	−64,568.73	−53,653.45
AIC	129,155.47	107,326.91
Dispersion (α)	Fixed at 0	0.3948***

Note: *** indicates significance at the $p < 0.001$ level. The significantly higher log-likelihood and lower AIC for the NBR model indicate superior fit for the overdispersed flow data.

A.1 Jensen’s Inequality

Traditionally, gravity models were estimated using OLS:

$$\ln(F_{i,j}) = \beta_0 + \beta_1 \ln(d_{i,j}) + \dots + \epsilon_{i,j} \quad (10)$$

However, as established by [24], this approach is structurally flawed for spatial interaction data. First, the implications of Jensen’s inequality $E(\ln y) \neq \ln E(y)$ mean that OLS estimates are biased and inconsistent in the presence of heteroskedasticity, which is present in our data. Second, log-linearization requires the exclusion or ad hoc adjustment of zero-flow observations, which truncates the dataset and further biases elasticity estimates. Finally, OLS is a continuous estimator that fails to account for the discrete nature of student transitions. By treating student flows as non-negative integers rather than continuous values, the NBR framework more accurately reflects the underlying data-generating process of individual enrollment decisions. It is possible that the OLS model will yield a lower AIC (in fact, it does), but this is a mathematical artifact of the logarithmic transformation that reduces the scale of the dependent variable and its associated residuals. Consequently, the OLS likelihood is not directly comparable to count-based models (Poisson and NBR), and we reject it on theoretical grounds in favor of a framework that preserves the original scale of student flows.

A.2 Overdispersion in Poisson Models

While the Poisson model uses a log-link function to avoid the bias of OLS, it assumes equidispersion, where $Var(y) = E(y)$. In the context of the Philippine basic education system, student transitions exhibit a power-law distribution where certain schools attract massive inflows while others only attract a few. This leads to overdispersion, where the variance significantly exceeds the mean. As shown in Table 3, moving from a Poisson to an NBR specification results in a massive improvement in log-likelihood (from −64,568 to −53,653). Furthermore, the estimated dispersion parameter ($\alpha = 0.3948$) is highly significant ($p < 0.001$), formally rejecting the Poisson model and confirming the NBR as the only robust choice for this dataset.

B Model Specification

To determine the optimal configuration for the predictive framework, we conducted a hierarchical stepwise evaluation. This process involves starting with a parsimonious gravity model and incrementally adding school-quality attributes, regional fixed effects, and socioeconomic controls. We evaluate each specification using three criteria:

- (1) Information Criteria (AIC/BIC): Lower values indicate a better balance between model fit and complexity.
- (2) Likelihood Ratio Test (LRT): A χ^2 test to determine if the improvement in Log-Likelihood from adding new variables is statistically significant.
- (3) Dispersion Parameter (α): Monitoring the overdispersion to ensure the negative binomial specification remains appropriate as controls are added.

Table 4: Hierarchical Model Comparison and Specification Tests

Model	Log-Lik.	AIC	BIC	LRT Stat	p-value	α
0. Baseline (Dist. + Cost)	-53,838.5	107,684.9	107,718.0	—	—	0.4022
1. + School Rating	-53,837.2	107,684.4	107,725.8	2.49	0.115	0.4021
2. + Origin Region	-53,800.1	107,614.3	107,672.3	74.13	< 0.001	0.4014
3. + Dest. Region	-53,790.0	107,598.0	107,672.6	20.28	< 0.001	0.4007
4. + Origin Income Only	-53,653.5	107,326.9	107,409.7	273.11	< 0.001	0.3948
5. + Dest. Income Only	-53,653.5	107,326.9	107,409.7	0.00	1.000	0.3948
6. + Dual Income	-53,592.4	107,206.8	107,297.9	122.08	< 0.001	0.3925

Note: LRT statistics compare each model to the specification in the immediately preceding row. Model 6 is the final specification used for policy simulations.

The results are presented in Table 4. Our selection of Model 6 as the final specification is informed by the following observations:

- The inclusion of regional fixed effects (Models 2 and 3) significantly improved model fit, confirming that student flows are largely shaped by local school networks and regional boundaries. While the school rating (Model 1) was statistically significant on its own, it was retained to control for perceived school quality.
- Adding the income of the local government unit (LGU) of the origin school provided the largest jump in log-likelihood, while adding that of the destination school yielded no statistical gain.
- The simultaneous inclusion of the LGU income of both origin and destination resulted in a significant improvement. This means that while origin income influences a student's starting capacity, the negative coefficient for destination income suggests that students often avoid schools that are economically out of reach, even when tuition is subsidized.

C Model Validation and Robustness

To evaluate the reliability of our expected flow predictions, we performed a non-parametric cluster bootstrap ($n = 1,000$) on the evaluation metrics. By resampling at the origin-school level, we account for the correlated error structures inherent in our administrative data. The narrow confidence intervals for our primary friction variables confirm that our policy simulations are grounded in consistent student behaviors, rather than being sensitive to localized data fluctuations.

Metric	Bootstrap Median	95% CI
MAE	1.5980	[1.5602, 1.6359]
RMSE	3.8782	[3.6566, 4.0998]

We also assess the stability of our flow determinants by re-estimating the chosen model (Model 6 in B) across 1,000 bootstrap iterations. This ensures that our behavioral elasticities—specifically distance and net cost—are representative of the entire school system and not driven by influential outliers.

Variable	Lower 2.5%	Median	Upper 97.5%
Intercept	3.1406	3.3946	3.6407
ln(distance)	-0.4693	-0.4506	-0.4292
ln(net cost)	-0.1260	-0.1007	-0.0762
ESC School Rating	-0.0372	-0.0207	-0.0050
ln(origin LGU income)	-0.0283	-0.0205	-0.0128
ln(dest. LGU income)	-0.0575	-0.0488	-0.0403
<i>Regional Effects (Ref: NCR)</i>			
Origin: Region III	-0.0743	-0.0208	0.0291
Origin: Region IV-A	-0.0897	-0.0504	-0.0099
Dest: Region III	-0.0724	-0.0221	0.0252
Dest: Region IV-A	-0.0176	0.0178	0.0573

The cluster bootstrap results confirm the structural integrity of the model, with key determinants of student flow exhibiting high stability across iterations. The 95% confidence intervals for spatial and economic friction are notably narrow and do not cross zero, providing robust empirical evidence for a non-compensatory hierarchy of priorities where geographic proximity serves as the primary filter. Furthermore, the confidence interval for destination LGU income remains entirely negative, statistically confirming the existence of hidden cost barriers, such as elevated prices for meals and supplies, that deter subsidy-eligible students from wealthier districts. Finally, while students from Region IV-A show a significantly lower propensity to participate, the destination fixed effects for both Region III and Region IV-A cross zero, suggesting that the attractiveness of a destination is effectively captured by institutional quality and economic variables rather than a pure regional preference.