



Vulnerability-Based Scholarship Allocation with Location-Dependent Service Costs: A Mixed-Integer Optimization Approach

Ignacio Zúñiga¹ Alejandro Cataldo^{*2} Pablo A. Rey³ Antoine Sauré⁴

¹School of Engineering, Pontificia Universidad Católica de Chile, Vicuña Mackenna 4860, Santiago, Chile

²Institute for Mathematical and Computational Engineering, School of Engineering, Pontificia Universidad Católica de Chile, Vicuña Mackenna 4860, Santiago, Chile

³Departamento de Ingeniería Industrial, Universidad Tecnológica Metropolitana, Av. José Pedro Alessandri 1242, Santiago, Chile

⁴Telfer School of Management, University of Ottawa, 55 Laurier Ave. E., Ottawa, ON K1N 6N5, Canada

*Corresponding author; email address: aecatald@uc.cl

Abstract

This paper studies a public resource allocation problem in vulnerability-based scholarship assignment with limited budgets and location-dependent service costs. We propose a mixed-integer linear programming model that integrates direct service costs, fixed location-activation costs, and stepwise service-capacity costs, while enforcing a mandatory-priority rule for highly vulnerable applicants. The approach is evaluated using a large-scale case study from the Chilean School Feeding Program (PAE), with 852,567 applicants and 1,520 institutions. Compared with the ranking-based baseline, the model increases coverage from 382,063 to 390,673 selected applicants while preserving a similar average vulnerability index (87.26 vs. 87.67). Cost-driven policies select more applicants, up to 435,522, but reduce average vulnerability to 55.76. The results show that optimization can improve coverage while maintaining vulnerability prioritization.

Keywords: Public sector, School meal benefits, Scholarship assignment, Service-based benefit allocation, Mixed-integer programming

1. Introduction

The allocation of public resources in education is a global challenge, as governments seek to improve student retention and academic outcomes while operating under limited budgets. Lower-income populations benefit disproportionately from investments in basic healthcare and primary and secondary education [24], which motivates scholarship programs aimed at vulnerable groups. Nevertheless, available resources rarely cover all eligible applicants. For example, conditional scholarships increased enrollment

during Indonesia's 1998 economic crisis [28], while Chile allocated USD 180 million in 2002 to provide meals for 1.3 million children [14].

Despite their relevance, scholarship planning and allocation are still frequently based on intuition, experience, or manual procedures [12]. This motivates quantitative tools that use available budgets more efficiently while increasing coverage and prioritizing those most in need.

In Chile, public social expenditure is allocated across multiple sectors, with education playing a central role. In 2022, the national public budget reached USD 82,135 million, with substantial investments in healthcare, economic recovery, public safety, children, older adults, and education, among others [15]. Institutions such as the Ministry of Education, the National Board of School Aid and Scholarships (JUNAEB), and the National Agency for Research and Development (ANID) provide benefits across different educational levels, from primary education to postgraduate studies. These benefits include food, transportation, housing, and monetary scholarships. Beyond financial support, scholarships also serve as incentives for academic achievement and persistence [18].

Scholarship programs can be broadly classified into two types. The first type provides direct incentives, typically monetary, where a fixed cost is incurred per beneficiary (potentially varying across individuals) and is independent of the total number and composition of recipients [30]. An example is the "Presidente de la República" scholarship in Chile [26]. The second type provides services, where the total cost includes both fixed and variable components that depend on the number and characteristics of beneficiaries [9].

In Chile, JUNAEB administers the School Feeding Program (PAE), which provides breakfast, lunch, and an afternoon snack to selected students. As a service-based program, its cost structure combines a direct cost per selected applicant, a fixed cost for activating each served educational institution, and a stepwise capacity cost when additional food-service staff are required [14].

To maintain consistent terminology, throughout the paper we use the term *applicant* to refer to any individual considered for the program, *selected applicant* to refer to an applicant who receives the benefit, and *school meal benefit* to refer to the service provided in the case study. The term *benefit* is used when referring to the broader class of service-based public programs addressed by the proposed formulation.

Different mechanisms are used to select scholarship recipients. A common approach consists of defining a scoring metric to rank applicants, followed by sequential allocation until the budget is exhausted. However, in practice, some programs may skip candidates in the ranking despite remaining budget. We refer to this phenomenon as *skipping an applicant*. This may occur, for example, when the benefit is non-monetary (e.g., a device with specific technical requirements) and a candidate is not compatible with the resource [11], or when minimum or maximum quotas on certain attributes must be satisfied, and assigning the next candidate would violate these constraints [13].

Given this context, there is a need for data-driven analytical methods and mathematical models to optimize resource allocation [23]. The objective of this work is to develop a modeling framework that incorporates budget constraints, service scale costs, distribution logistics, and applicants' socioeconomic vulnerability to determine an efficient and sustainable allocation of school meal benefits. Rather than using a broad notion of social impact, the proposed approach explicitly maximizes the total vulnerability score of selected applicants, while enforcing a mandatory-priority rule for highly vulnerable applicants and accounting for the cost structure induced by service provision across locations.

This paper makes three main contributions. First, it formulates service-based public benefit allocation as a fixed-charge, logistics-aware problem in which the marginal cost of selecting an applicant depends on the other selected applicants associated with the same location. Second, it proposes a mixed-integer linear programming model that integrates vulnerability scores, mandatory priority for highly vulnerable applicants, location activation, and stepwise capacity costs under a common budget. Third, it evaluates the formulation using a large-scale Chilean PAE case study and quantifies the trade-off among coverage, cost efficiency, baseline alignment, and vulnerability prioritization. By explicitly linking individual priority rules with location-level economies of scale, the framework identifies when deviations from a strict ranking can improve coverage without substantially weakening the protection of highly vulnerable applicants.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature. Section 3 describes the problem setting. Section 4 presents the proposed mathematical model. Section 5 introduces the case study used to evaluate the approach. Section 6 describes the benchmark policies and performance metrics. Section 7 reports the computational results. Section 8 discusses the main limitations of the study. Section 9 presents the main conclusions and directions for future research.

2. Literature Review

The allocation of public resources, particularly in the educational sector, has been extensively studied from multiple perspectives in the literature, encompassing both economic and social criteria. Given resource scarcity, the objective is to optimize their distribution while preserving key considerations such as equity and the prioritization of vulnerable populations.

The present work is positioned at the intersection of three streams of literature. The first stream studies public resource allocation under equity and efficiency considerations, where decision makers must balance aggregate performance with the protection of disadvantaged groups. The second stream addresses logistics and service provision in educational and school feeding programs, where fixed costs, economies of scale, and spatial distribution play a central role. The third stream focuses on scholarship and benefit assignment mechanisms, where applicants are typically ranked according to merit, need, vulnerability, or multi-criteria scores. The methodological gap addressed in this paper arises because these three dimensions are usually treated separately: equity-oriented allocation models often abstract from service-scale costs, logistics models typically focus on suppliers or distribution networks rather than applicant-level selection, and scholarship-selection methods frequently rank applicants without explicitly modeling location-dependent cost structures.

In the healthcare domain, the fair allocation of scarce resources has been widely analyzed. For instance, Anahideh et al. [1] propose an approach to equitably distribute vaccines during the COVID-19 pandemic by incorporating geographic diversity and socioeconomic vulnerability. Their work introduces an optimization model that minimizes equity gaps across subgroups while ensuring protection for disadvantaged populations. Resource allocation has also been studied in other contexts such as public safety. In this setting, Baycik et al. [4] address the problem of allocating resources between intelligence operations (information gathering on drug trafficking networks) and interdiction operations (arrests and seizures), with the objective of maximizing law enforcement effectiveness at the city level.

The equity–efficiency tension is also closely related to the broader operations research literature on fairness in resource allocation. For example, Bertsimas et al. [5] formalize the concept of the “price of fairness”, showing how imposing fairness-oriented allocation rules may reduce system efficiency relative to purely efficiency-maximizing solutions. This perspective is relevant to the present setting because protecting highly vulnerable applicants may reduce the number of selected applicants, whereas purely cost-driven policies may increase coverage at the expense of selecting less vulnerable populations. The proposed model explicitly captures this tension by combining an efficiency-oriented budget constraint with a mandatory-priority rule for applicants above a vulnerability threshold.

From a broader perspective, Aray et al. [2] analyze the role of political considerations in public resource allocation decisions and highlight cases where operations research solutions are not directly implemented. Similarly, Mayerle et al. [23] develop an optimization framework to improve resource allocation in public education systems, focusing on the interaction between students, schools, classes, teachers, and classrooms. The problem consists of determining an optimal matching under capacity constraints to enhance efficiency in the use of public resources.

Regarding the challenge of managing limited resources while satisfying logistical constraints and addressing the needs of vulnerable populations, Souza et al. [27] study the allocation of individuals to shelters and the distribution of humanitarian aid in disaster and emergency contexts.

In the educational domain, the integration of social and economic criteria is essential. For example, Bouzarth et al. [7] address the assignment of students to schools by minimizing transportation costs and socioeconomic disparities between institutions. The problem is initially formulated as a nonlinear model and subsequently linearized, allowing it to be solved as a mixed-integer linear program. Similarly, Mayerle et al. [23] propose a simultaneous optimization framework that considers both supply (infrastructure, teachers, classrooms) and demand (students), incorporating capacity and spatial distribution constraints within the educational system.

The logistical dimension of school feeding programs has also received increasing attention. For instance, Granillo et al. [16] develop a socially-oriented optimization model for school breakfast delivery, combining clustering techniques with socioeconomic variables and achieving significant reductions in distribution costs. At a larger scale, Mahadevan et al. [22] apply operations research techniques to a feeding program in India, obtaining substantial savings in routing and resource allocation, which enabled program expansion. Similarly, Aydemir et al. [3] propose a multi-objective approach for designing distribution networks in Turkey, emphasizing trade-offs between cost, delivery times, and environmental sustainability. These studies demonstrate that optimization and routing methodologies have a measurable impact on both logistical efficiency and program equity.

The importance of equity and efficiency in educational resource allocation is also reflected in studies focused on scholarship programs. For example, Huang et al. [18] analyze incentive mechanisms to improve academic performance through scholarship allocation, while Puspitasari et al. [25], Latumakulita et al. [20], and Wimatsari et al. [29] employ multi-criteria decision-making approaches (e.g., Analytic Hierarchy Process (AHP) and fuzzy methods) to identify suitable candidates for scholarships. These approaches incorporate both socioeconomic and academic data to maximize social impact. In a related line of work, Durán et al. [13] develop mathematical models for selecting applicants to a graduate program, accounting for equity and socioeconomic vulnerability.

The Chilean case is particularly relevant due to the scale of its student support programs. The National Board of School Aid and Scholarships (JUNAEB) faces the challenge of optimizing the allocation of public resources for school feeding programs, ensuring retention and academic success among vulnerable students. Although advanced analytical models for the direct allocation of these benefits are limited, related optimization tools have been developed. For example, Epstein et al. [14] apply mixed-integer programming techniques to assign contracts to school meal providers, solving a large-scale and complex problem. Subsequently, Catalán et al. [8] improve computational performance through scenario ordering strategies, highlighting the relevance of optimization techniques in public policy applications.

Beyond the educational context, public resource allocation has been addressed using a variety of methodologies. Lozano et al. [21] and Contreras et al. [10] propose optimization models for centralized resource allocation in the Spanish public university system, incorporating lexicographic goal programming and negotiation-based approaches among multiple stakeholders. Similarly, Karakaya et al. [19] present a hierarchical bi-objective model for location and allocation decisions in maternal-neonatal care systems in Turkey, aiming to improve access and reduce transportation costs.

Despite these contributions, scholarship allocation in higher education institutions is often still based on subjective judgment or informal procedures. For instance, Gross et al. [17] show that merit- and need-based allocation criteria are frequently supported by incomplete data and non-rigorous processes, potentially leading to inefficient or inequitable outcomes. Likewise, Boatman et al. [6] argue that even with substantial funding, the absence of a clear analytical framework limits the effectiveness of scholarship programs, highlighting the importance of integrating statistical analysis with expert judgment.

In contrast to these related streams, the allocation problem studied in this paper concerns service-based public benefits assigned at the applicant level but delivered through geographically distributed locations. This creates a distinctive modeling structure. On the one hand, the decision maker must account for individual vulnerability scores and priority rules, as in scholarship and benefit assignment problems. On the other hand, the cost of selecting an applicant is not purely individual: it depends on whether the associated location is activated and on the number of selected applicants already served at that location. Therefore, the problem combines applicant-level selection, location-level fixed charges, stepwise service-capacity costs, and a fixed per-period budget.

Existing work on school feeding programs has demonstrated the value of optimization for contract allocation and distribution-network design [3, 8, 14, 16, 22]. However, these studies primarily address provider selection, routing, logistics, or network design decisions. They do not directly model the applicant-level allocation of benefits under a vulnerability-based priority rule and location-dependent service costs. Similarly, multi-criteria scholarship selection approaches [18, 20, 25, 29] are useful for ranking or scoring candidates, but they generally do not capture the nonseparable cost structure induced by delivering a recurring service through shared locations.

Motivated by this gap, this paper proposes a mathematical optimization framework for allocating service-based public benefits to ranked applicants under a fixed per-period budget. The model allows deviations from a strict ranking-based policy when such deviations improve coverage, but it also enforces a mandatory-priority rule for applicants above a predefined vulnerability threshold. In doing so, the formulation provides a structured way to quantify the trade-off between preserving vulnerability prioritization and exploiting location-level economies of scale.

3. Problem Description

We consider a set $\mathcal{I}$ of applicants to a service-based public benefit program with a known periodicity, where benefits are not automatically renewed (i.e., applicants must reapply in each period). Applicants are ranked according to a nonnegative real-valued vulnerability index, denoted by $V_i \in \mathbb{R}_+$ for each applicant. Higher values of V_i indicate greater socioeconomic vulnerability, whereas lower values indicate lower vulnerability. The mathematical formulation does not require V_i to be integer. In the empirical application, the vulnerability index is normalized on a scale from 1 to 100, allowing applicants to be consistently prioritized according to their socioeconomic vulnerability.

A vulnerability threshold $\bar{V}$ defines the mandatory-priority set

$$\mathcal{I}^+ = \{i \in \mathcal{I} : V_i \geq \bar{V}\}.$$

The model requires all applicants in $\mathcal{I}^+$ to be selected whenever this rule is feasible under the available budget. Otherwise, the deterministic formulation becomes infeasible, indicating that the threshold, budget, or priority rule must be revised or relaxed through penalties or a lexicographic formulation.

The benefit provides a standardized service to selected applicants, incurring a constant direct service cost of C monetary units per applicant in the empirical application. This cost is referred to as the *direct service cost*. It represents the unit cost directly associated with serving an applicant, independently of where the service is delivered.

The service is delivered at a set of predefined locations $\mathcal{N}$, such as schools, service centers, local facilities, or other institutional locations. Each applicant is associated ex ante with exactly one location; this association is an input of the problem rather than a decision. The binary parameter U_{in} equals 1 if applicant i is associated with location n , and the corresponding applicant set is $Q_n = \{i \in \mathcal{I} : U_{in} = 1\}$. Activating a location requires incurring a fixed cost of L_n monetary units per period whenever at least one applicant associated with location n is selected. This cost is referred to as the *fixed location-activation cost* or *fixed logistical cost*.

Providing the service may also require discrete capacity resources. At location n , each capacity unit serves up to K_n selected applicants and costs G_n monetary units. Consequently, selecting between 1 and K_n applicants requires one unit, selecting between $K_n + 1$ and $2K_n$ requires two, and so on. We refer to this component as the *stepwise service-capacity cost*.

The decision maker has a non-transferable budget of P monetary units in each allocation period and must select applicants to maximize total vulnerability while covering direct, fixed logistical, and stepwise capacity costs and enforcing the mandatory-priority rule.

4. Solution Approach

The proposed solution approach is based on a mixed-integer linear programming (MILP) model. Its formulation relies on the sets, indices, parameters, and decision variables defined in Table 1.

The formulation can be interpreted as a fixed-charge, stepwise-cost allocation model with applicant-level benefits and location-level cost coupling. In a purely ranking-based policy, each applicant is considered sequentially according to the vulnerability index, and the ranking itself determines the allocation

order. In the proposed model, by contrast, the decision to select an applicant depends jointly on the applicant's vulnerability score, the activation status of the associated location, and the service-capacity units already required at that location. This coupling is precisely what allows the model to exploit economies of scale while preserving a mandatory-priority rule for highly vulnerable applicants.

This structure also clarifies the methodological role of the binary location variables y_n and the integer service-capacity variables H_n . The variables y_n capture the fixed-charge nature of serving a location, while the variables H_n represent the stepwise capacity expansion required as more applicants are selected at the same location. Thus, the model does not merely select applicants independently under a budget constraint; it jointly determines which applicants and locations should be served, and how much location-level capacity is required to support the resulting allocation.

Table 1. Model notation

Set	Definition
$\mathcal{I}$	Set of applicants to the benefit program
$\mathcal{N}$	Set of locations where the service-based benefit is provided
$\mathcal{I}^+$	Subset of mandatory-priority applicants, defined as $\mathcal{I}^+ = \{i \in \mathcal{I} : V_i \geq \bar{V}\}$
Index	Definition
i	Index associated with applicants, $i \in \mathcal{I}$
n	Index associated with locations, $n \in \mathcal{N}$
Parameter	Definition
P	Available budget for the benefit program
C	Constant direct service cost per applicant in the empirical application
K_n	Maximum number of selected applicants that can be served by one service-capacity unit at location n
G_n	Stepwise service-capacity cost per capacity unit at location n
L_n	Fixed location-activation or logistical cost incurred when location n is served
V_i	Nonnegative real-valued vulnerability index of applicant i , $V_i \in \mathbb{R}_+$; higher values indicate greater vulnerability
U_{in}	Binary parameter equal to 1 if applicant i is associated with location n
Q_n	Subset of applicants exogenously associated with location n , $Q_n = \{i \in \mathcal{I} : U_{in} = 1\}$
$\bar{V}$	Vulnerability threshold used to define the mandatory-priority set $\mathcal{I}^+$
Variable	Definition
x_i	Binary variable equal to 1 if applicant i is selected to receive the benefit, and 0 otherwise
y_n	Binary variable equal to 1 if at least one applicant is selected at location n , and 0 otherwise
H_n	Integer variable representing the number of service-capacity units required at location n

The objective function maximizes the total vulnerability of the selected applicants, as defined in (1):

$$\max_{\vec{x}} \sum_{i \in \mathcal{I}} V_i \cdot x_i \quad (1)$$

The constraints ensure that all costs associated with providing the benefit are considered within the available budget, that locations are activated whenever at least one applicant is selected, that the stepwise service-capacity structure is properly represented, and that all applicants in the mandatory-priority set $\mathcal{I}^+$ are selected whenever the resulting allocation is feasible. The model is defined by the following constraints:

$$C \sum_{i \in \mathcal{I}} x_i + \sum_{n \in \mathcal{N}} L_n \cdot y_n + \sum_{n \in \mathcal{N}} G_n \cdot H_n \leq P \quad (2)$$

$$\begin{aligned} \sum_{i \in Q_n} x_i &\geq y_n \quad \forall n \in \mathcal{N} \\ \sum_{i \in Q_n} x_i &\leq |Q_n| \cdot y_n \quad \forall n \in \mathcal{N} \end{aligned} \quad (3)$$

$$\frac{\sum_{i \in Q_n} x_i}{K_n} \leq H_n \quad \forall n \in \mathcal{N} \quad (4)$$

$$x_i = 1 \quad \forall i \in \mathcal{I}^+ \quad (5)$$

$$\begin{aligned} x_i &\in \{0, 1\} \quad \forall i \in \mathcal{I} \\ y_n &\in \{0, 1\} \quad \forall n \in \mathcal{N} \\ H_n &\in \mathbb{Z}_0^+ \quad \forall n \in \mathcal{N} \end{aligned} \quad (6)$$

The role of each constraint is as follows. Constraint (2) enforces the budget limit P by accounting for all cost components: (i) the constant direct service cost per applicant in the empirical application, (ii) the fixed location-activation or logistical cost incurred when a location is served, and (iii) the stepwise service-capacity cost required to provide the benefit at each location.

In the empirical application, the direct service cost is represented by the scalar parameter C because the school meal package is standardized across applicants within the studied region. Nevertheless, the formulation can accommodate heterogeneous direct service costs without modifying its underlying structure. If costs vary across applicants, the term $C \sum_{i \in \mathcal{I}} x_i$ in constraint (2) can be replaced by $\sum_{i \in \mathcal{I}} C_i x_i$, where C_i denotes the direct service cost of applicant i . If direct service costs vary by location or region, they can instead be represented by $\sum_{n \in \mathcal{N}} C_n \sum_{i \in Q_n} x_i$, where C_n denotes the direct service cost applicable to applicants associated with location n .

The two inequalities in constraint (3) jointly impose the correct logical relationship between applicant selection and location activation. The first inequality prevents the activation of an empty location: if $y_n = 1$, at least one applicant in Q_n must be selected. The second inequality ensures that selecting any applicant in Q_n forces $y_n = 1$; equivalently, if $y_n = 0$, it forces $\sum_{i \in Q_n} x_i = 0$. The location-specific coefficient $|Q_n|$ provides a tighter upper bound than the global coefficient $|\mathcal{I}|$. Because each applicant's location is fixed in the input data through U_{in} and Q_n , the single-indexed decision variable x_i is sufficient. A double-indexed variable x_{in} would be required only if assigning applicants to locations were itself a decision. Constraint (4) captures the stepwise structure of service-capacity requirements by determining the number of capacity units of size K_n required at each location, thereby defining the resource usage variable H_n . Constraint (5) enforces the mandatory-priority rule by selecting all applicants whose vulnerability index is at least $\bar{V}$. Finally, constraints (6) define the domain of the decision variables.

5. Case Study

The case study is based on the School Feeding Program (PAE) administered by the National Board of School Aid and Scholarships (JUNAEB) in Chile. The problem addressed concerns the allocation of school meal benefits to students, which are delivered through the educational institutions they attend. The operational data used in this study were obtained directly from JUNAEB's institutional databases, including information on applicants, educational institutions, vulnerability indices, location assignments, budget levels, and the cost components required to represent the allocation problem.

For confidentiality reasons, the data were perturbed before conducting the computational experiments. The perturbation process was designed to preserve the main empirical distributions observed in the original databases, including the distributions of vulnerability scores, the geographic distribution of applicants and educational institutions, and the relative structure of the cost parameters. In addition, the perturbed instance was calibrated to remain consistent with the scale of the actual program, including the order of magnitude of the available budget and the number of selected applicants under the baseline allocation rule. Therefore, the instance used in this paper should be interpreted as a representative and calibrated version of the real operational setting, rather than as an exact reproduction of the confidential JUNAEB databases.

The baseline policy used in the computational experiments, referred to as Policy 1, reflects the operational rules used by JUNAEB to prioritize and assign school meal benefits. In particular, it represents a ranking-based allocation rule in which applicants are ordered according to their vulnerability index and selected sequentially while the available budget allows the service to be provided. This policy was not defined as an ad hoc heuristic for this study; rather, it was constructed from JUNAEB's allocation rules and contrasted against real allocation data to verify that it reproduces the main aggregate patterns of the current assignment procedure.

Although the benefit is conceptually assigned at the individual level, its implementation is carried out at the school level. In other words, the service provider receives a budget to deliver school meal benefits to the selected applicants attending each institution. Consequently, cost structures depend not only on the number of selected applicants but also on their geographic distribution and on how service-capacity and logistical resources are coordinated to provide the benefit.

The case study focuses on the Santiago Metropolitan Region, which includes 852,567 potential beneficiaries (applicants) across 1,520 educational institutions. To deliver the program in this region, JUNAEB operates under a total budget of 300 million monetary units. School meal benefits are allocated based on a nonnegative real-valued vulnerability index V_i , normalized on a scale from 1 to 100, where higher values indicate greater vulnerability. The provision of the benefit involves three main cost components: direct food costs incurred for each selected applicant, fixed delivery or logistical activation costs incurred when an institution is served, and stepwise staffing costs associated with the personnel required to prepare and serve meals. Table 2 summarizes the values, ranges, and empirical specifications of the parameters used in the computational analysis.

The stepwise staffing component introduces an important structural feature of the problem, namely the presence of economies of scale in scholarship allocation. For example, at least one food-service worker is required for any institution with at least one beneficiary, and each worker can serve up to 70 students,


Table 2. Parameters used in the empirical analysis

Parameter	Meaning	Value or empirical specification	Treatment in the case study
$ \mathcal{I} $	Number of applicants	852,567 applicants	Obtained from JUNAEB's operational database
$ \mathcal{N} $	Number of educational institutions	1,520 institutions	Obtained from JUNAEB's operational database
P	Available budget	300,000,000 monetary units	Calibrated to the scale of the operational allocation
V_i	Vulnerability index	Real-valued, $1 \leq V_i \leq 100$	Perturbed while preserving the empirical vulnerability distribution
V	Mandatory-priority threshold	80 in the main experiment	Defines $\mathcal{I}^+ = \{i \in \mathcal{I} : V_i \geq V\}$
C	Constant direct service cost per applicant	Constant perturbed monetary coefficient	Constant because the school meal package is standardized in the studied region
K_n	Applicants served by one service-capacity unit	70 applicants	One food-service staffing unit can serve up to 70 applicants
G_n	Cost per service-capacity unit at location n	Location-specific perturbed monetary coefficients	Preserves the relative empirical structure of staffing costs
L_n	Fixed activation or logistical cost at location n	Location-specific perturbed monetary coefficients	Preserves the relative empirical structure of location-level delivery costs
U_{in}	Applicant–location association	Binary, $U_{in} \in \{0, 1\}$	Each applicant is associated with exactly one observed educational institution
Q_n	Applicants associated with location n	$Q_n = \{i \in \mathcal{I} : U_{in} = 1\}$	Derived from the observed applicant–institution associations

Note. The exact perturbed monetary coefficients used for C , G_n , and L_n are not reproduced because they were derived from confidential operational information. The complete perturbed coefficients were used in all computational experiments and are subject to the conditions stated in the data availability declaration.

as represented by parameter K_n in Table 2. Thus, an institution serving up to 70 beneficiaries requires one staffing unit, whereas serving 71 beneficiaries requires a second staffing unit, increasing the staffing cost discretely. Therefore, both the number of beneficiaries and their distribution across schools have a significant impact on total costs.

Additionally, the region exhibits both urban and rural areas, leading to heterogeneous cost structures. Urban areas typically have higher concentrations of applicants and lower delivery costs, whereas rural areas have lower student density and higher logistical activation costs. Figure 1 presents the spatial distribution of average vulnerability by school (blue indicating higher vulnerability and red lower vulnerability). Figure 2 illustrates the spatial variation in fixed delivery or logistical activation costs for serving each institution, where red indicates lower costs and blue higher costs.

**Figure 1.** Spatial distribution of average vulnerability index by educational institution.

As observed, some rural schools exhibit high transportation costs and relatively few potential beneficiaries, while urban areas concentrate a larger number of applicants with lower logistical costs. The challenge is to allocate scholarships to students—considering both individual characteristics and their as-




Figure 2. Spatial distribution of logistical costs across educational institutions.

sociated institutions—in a way that maximizes the impact of the program without exceeding the available budget.

6. Computational Experiments

To evaluate the performance of the proposed approach, we compare it against several allocation policies, including the one currently implemented by JUNAEB. The comparison is conducted using performance metrics that are relevant to the problem setting. The benchmark policies and evaluation metrics are described below.

6.1. Policies

To assess the performance of the proposed solution, denoted by SP, we compare it with five benchmark allocation policies. Policy 1 represents the ranking-based operational baseline derived from JUNAEB’s assignment rules and validated against real allocation data. Policies 2 and 3 are constructive allocation rules designed to emphasize marginal cost efficiency and location-level economies of scale, respectively. Policies 4 and 5 are optimization-based benchmarks. Detailed pseudocode for Policies 1–3 is provided in Appendix A as supporting material for transparency and reproducibility.

6.1.1. Allocation Policy 1

Under Policy 1, applicants are sorted in decreasing order of their vulnerability index V_i and selected sequentially while the resulting allocation remains within the available budget. Once the next applicant in the ranking cannot be selected without exceeding the budget, the allocation procedure stops. Consequently, applicants cannot be skipped. Policy 1 is used as the baseline because it represents JUNAEB’s operational ranking-based assignment logic. It was constructed from the institutional allocation rules and contrasted with historical allocation data to verify that it reproduces the principal aggregate characteristics of the current procedure. The implementation is detailed in Algorithm 1 in Appendix A.

6.1.2. Allocation Policy 2

Policy 2 selects, at each iteration, the applicant with the lowest marginal assignment cost. This marginal cost includes the direct service cost and any additional fixed location-activation or stepwise service-capacity cost induced at the applicant's educational institution. The procedure continues while an additional applicant can be selected without exceeding the available budget. Its implementation is detailed in Algorithm 2 in Appendix A.

6.1.3. Allocation Policy 3

Policy 3 seeks to exploit location-level economies of scale while retaining vulnerability prioritization. The procedure first selects the most vulnerable applicant who has not yet been selected and then considers additional high-vulnerability applicants from the same educational institution, up to the corresponding service-capacity block. Applicants are added while the remaining budget permits, preserving the vulnerability order within the institution. The procedure is repeated until no further feasible selection can be made. Its implementation is detailed in Algorithm 3 in Appendix A.

6.1.4. Allocation Policy 4

This allocation policy seeks to maximize the total vulnerability of the selected applicants. To do so, it uses the optimization model proposed in Section 4, removing the mandatory-priority vulnerability constraint (constraint (5)).

6.1.5. Allocation Policy 5

Finally, this policy seeks to maximize the total number of selected applicants without considering the vulnerability index V_i . To do so, it uses the optimization model proposed in Section 4, replacing the objective function with the following:

$$\max \sum_{i \in \mathcal{I}} x_i$$

6.2. Performance Metrics

To compare the performance of the proposed approach against the five policies described above, we define the following evaluation metrics:

1. **Direct cost percentage:** This metric measures the proportion of total expenditure corresponding to direct costs.
2. **Number of selected applicants:** This metric counts the total number of applicants selected to receive the benefit.
3. **Baseline-displaced applicants:** This metric measures the deviation of a policy from the baseline ranking-based allocation. Let $\mathcal{S}^1$ denote the set of applicants selected under Policy 1, and let $\mathcal{S}^p$ denote the set of applicants selected under policy p . The number of baseline-displaced applicants is defined as $|\mathcal{S}^1 \setminus \mathcal{S}^p|$. Therefore, this metric does not measure “skipping” in an absolute ranking

sense; rather, it counts how many applicants selected by the baseline policy are not selected by the policy being evaluated.

4. **Average vulnerability index of selected applicants:** This metric computes the average vulnerability index of the applicants selected to receive the benefit.
5. **Position of the last selected applicant:** This metric reports the ranking position of the last selected applicant in the list sorted by vulnerability index.

It is important to distinguish between baseline displacement and ranking-based skipping. The former is measured relative to Policy 1, which is used as the baseline allocation rule. The latter refers to selecting lower-ranked applicants while leaving higher-ranked applicants unselected. In this study, the baseline-displacement metric is used to quantify how much each policy departs from the current ranking-based allocation, whereas the position of the last selected applicant provides complementary information on how far into the ranking each policy reaches.

7. Results

To evaluate the performance of the proposed policy relative to the policies described in Section 6.1, we consider the instance associated with the JUNAEB case study. As described in Section 5, this instance was constructed from JUNAEB operational databases, perturbed to protect confidentiality while preserving the main empirical distributions and calibrated against the current allocation procedure. All models and algorithms were implemented in Python 3.10 and executed on the Cedar cluster of the Digital Research Alliance of Canada. The optimization models were solved using GUROBI 9.5.

The resulting optimization model is large-scale, but its variable structure follows directly from the formulation presented in Section 4. Specifically, the model includes one binary variable x_i for each applicant $i \in \mathcal{I}$, one binary activation variable y_n for each location $n \in \mathcal{N}$, and one nonnegative integer variable H_n for each location $n \in \mathcal{N}$. Thus, the formulation does not require a binary decision variable for each applicant–location pair; applicant–location associations are captured through the input sets Q_n and the assignment parameters U_{in} .

The computational effort was mainly associated with preparing the large-scale instance, including data preprocessing, construction of the applicant–location sets, parameter generation, and model instantiation. Once the instance was constructed, the optimization model was solved within approximately two hours. This distinction is important because the reported preprocessing and model-generation effort should not be interpreted as solver time.

For the case study, with 852,567 applicants and 1,520 locations, this structure leads to 852,567 applicant-selection binary variables, 1,520 location-activation binary variables, and 1,520 integer service-capacity variables, in addition to the corresponding budget, activation, capacity, and priority constraints.

7.1. Metrics

Table 3 summarizes the performance of each policy across the defined metrics. The proposed solution is denoted as SP. In the main experiment, the vulnerability threshold was set to $\bar{V} = 80$. This value should be interpreted as a policy parameter rather than as an estimated statistical cutoff. It defines the

set of applicants who receive mandatory priority in the allocation process and, consequently, controls the trade-off between protecting highly vulnerable applicants and allowing flexibility to improve overall coverage.

The choice of $\bar{V} = 80$ was used as the main scenario because it represents a high-vulnerability threshold on the normalized vulnerability scale and provides a clear policy interpretation: applicants in the upper segment of the vulnerability distribution are treated as mandatory-priority applicants. Before solving the proposed model, we verified that the mandatory-priority set $\mathcal{I}^+ = \{i \in \mathcal{I} : V_i \geq 80\}$ was feasible under the available budget. Therefore, constraint (5) represents an enforceable policy requirement in the evaluated instance.

Because the threshold directly affects the feasible set and the allocation pattern, the results reported below should be interpreted as conditional on the policy choice $\bar{V} = 80$. Different values of $\bar{V}$ would induce different mandatory-priority sets and, therefore, potentially different coverage-equity trade-offs. Thus, in practical implementations, $\bar{V}$ should be selected by the decision maker according to the desired level of protection for highly vulnerable applicants, and sensitivity analyses over alternative threshold values should be used to support that choice.

Table 3 reports the main results using Policy 1 as the reference point. Rather than highlighting only the best value in each column, the table includes the change in coverage and average vulnerability relative to the ranking-based baseline. This format is intended to make the policy trade-offs explicit: for example, a policy may increase the number of selected applicants while substantially reducing the average vulnerability of the selected population. The proposed solution (SP) is emphasized because it corresponds to the model proposed in this paper, and its performance should be interpreted as a compromise between improving coverage, preserving vulnerability prioritization, and remaining close to the baseline allocation.

Table 3. Comparison of allocation policies relative to the ranking-based baseline

Policy	Direct cost (%)	Selected applicants	Δ selected vs. P1	Avg. vuln. index	Δ vuln. vs. P1	Baseline-displaced applicants	Last rank selected
P1	68.6	382,063	–	87.67	–	N/A	382,063
P2	77.0	428,335	+46,272	55.76	-31.91	266,176	852,566
P3	70.0	393,736	+11,673	86.36	-1.31	19,952	852,552
P4	74.5	414,666	+32,603	85.23	-2.44	40,782	578,954
P5	78.2	435,522	+53,459	57.90	-29.77	248,936	852,566
SP	70.2	390,673	+8,610	87.26	-0.41	12,334	482,052

Note. Changes are computed relative to Policy 1, the ranking-based baseline. Positive values in Δ selected indicate additional selected applicants relative to the baseline, whereas negative values in Δ vulnerability indicate a reduction in the average vulnerability index. The SP row is emphasized because it corresponds to the proposed solution. Its interpretation is not that it maximizes a single metric, but that it increases coverage while preserving an average vulnerability index close to the baseline and producing the lowest number of baseline-displaced applicants among the non-baseline policies.

The results reveal clear trade-offs between efficiency (coverage and cost structure) and equity (vulnerability prioritization).

Policies 2 and 5 achieve the highest number of selected applicants, with 428,335 and 435,522 selected applicants, respectively. This is consistent with their objective of prioritizing cost efficiency, which is

also reflected in the highest proportion of direct costs (up to 78.2%). However, this comes at the expense of equity, as these policies select significantly less vulnerable applicants, resulting in substantially lower average vulnerability indices (55.76 and 57.9, respectively).

In contrast, Policy 1 prioritizes vulnerability, achieving the highest average vulnerability (87.67) and selecting applicants strictly according to the ranking. However, this results in the lowest coverage, with 382,063 selected applicants, indicating an efficiency loss in terms of the number of applicants served.

Policies 3 and 4 provide intermediate behaviors, partially improving coverage while still prioritizing vulnerable applicants. However, they also displace a non-negligible number of applicants selected by the baseline policy, as measured by the baseline-displacement metric. This indicates that their coverage gains are achieved by modifying the composition of the baseline allocation.

The proposed solution (SP), under the threshold $\bar{V} = 80$, achieves a balanced performance. It maintains a high average vulnerability (87.26), very close to Policy 1, while increasing the number of selected applicants compared to Policy 1. Additionally, it results in the lowest number of baseline-displaced applicants (12,334), indicating strong alignment with the baseline allocation while improving overall efficiency.

The redesigned comparison in Table 3 also shows why selecting a single “best” policy based only on column-wise maxima can be misleading. Policy 5, for instance, provides the largest increase in coverage relative to Policy 1 (+53,459 applicants), but this gain is accompanied by a large reduction in average vulnerability (-29.77 points) and a substantial displacement of applicants selected by the baseline policy. In contrast, SP produces a smaller but meaningful coverage increase (+8,610 applicants), while reducing the average vulnerability index by only 0.41 points and displacing 12,334 baseline applicants. Thus, SP should be interpreted as a policy-oriented compromise rather than as a solution that dominates all alternatives in every metric.

Overall, these results suggest that, for the selected threshold value, the proposed approach provides a favorable compromise between maximizing coverage and preserving vulnerability prioritization. Therefore, the performance of SP should not be interpreted as unconditional dominance over all alternatives, but rather as the outcome of a policy design that explicitly protects applicants above the mandatory-priority threshold while using the remaining budget more efficiently.

7.2. Spatial analysis

This section presents the spatial distribution of scholarship allocations across schools in the Santiago Metropolitan Region. Each point represents an educational institution. Blue indicates that more than 70% of applicants at the institution were selected, yellow represents between 51% and 70%, red between 31% and 50%, and black indicates less than 30%.

Figure 3 presents the allocation maps for all policies. From left to right and top to bottom, the figures correspond to Policies 1 through 5, followed by the proposed solution.

The spatial patterns highlight important differences between policies. Policy 1 concentrates allocation in rural areas, which tend to have higher vulnerability levels. In contrast, Policies 2 and 5 prioritize urban areas, where logistical costs are lower, leading to a higher number of beneficiaries but at the expense of excluding vulnerable rural populations.

Policies 3 and 4 partially improve upon the baseline allocation, increasing coverage across institutions




Figure 3. Comparison of Spatial Allocations across Policies

while maintaining some focus on vulnerable areas.

The proposed solution exhibits a hybrid behavior. It preserves strong coverage in rural areas—similar to Policy 1—while also increasing allocation in urban zones, particularly in the southern part of Santiago. This demonstrates its ability to balance geographical equity and cost efficiency.

Compared to Policy 1, the proposed approach expands coverage without significantly compromising the allocation to highly vulnerable regions. In contrast, Policies 2 and 5 achieve higher coverage by concentrating almost exclusively on low-cost urban areas, effectively neglecting the initial target population prioritized by Policy 1.

7.3. Comparison of allocation dynamics

Figure 4 compares the cumulative allocation across policies relative to Policy 1. The horizontal axis represents applicants ordered by vulnerability, while the vertical axis shows the cumulative number of selected applicants. The blue line corresponds to Policy 1.

The allocation dynamics provide additional insight into the behavior of each policy.

Policies 2 and 5 exhibit strongly convex (parabolic-like) curves, indicating that they allocate scholarships primarily to applicants located further down in the vulnerability ranking. As a result, they reach a much larger portion of the ranking but also displace many applicants selected by the baseline ranking-based policy.

Policies 3 and 4 deviate moderately from Policy 1. They select more applicants than the baseline policy, but this improvement is accompanied by changes in the composition of selected applicants. This is reflected both in their baseline-displacement values and in the fact that their last selected applicants are located further down in the vulnerability ranking.




Figure 4. Comparison of allocation dynamics across policies

The proposed solution closely follows the allocation pattern of Policy 1 for applicants in the mandatory-priority segment defined by $\bar{V} = 80$. Once this requirement is satisfied, the model uses the remaining budget to improve coverage, which may involve selecting some lower-ranked applicants instead of some applicants selected by the baseline policy. This explains why SP reaches position 482,052 in the ranking while still maintaining a low number of baseline-displaced applicants relative to the other non-baseline policies.

This behavior explains its strong performance in the main scenario: it preserves the allocation of highly vulnerable applicants protected by the threshold, remains close to the baseline policy, and still improves overall coverage. At the same time, these conclusions are conditional on the selected value of $\bar{V}$, since a higher or lower threshold would modify the mandatory-priority set and the degree of flexibility available to the model.

8. Limitations

Several limitations should be considered when interpreting the results. First, although the case study uses operational data from JUNAEB, costs, prices, and vulnerability indices were perturbed to protect confidentiality. The perturbation preserved the main empirical distributions and relative structures, and the instance was calibrated against the actual allocation procedure. Therefore, the results are representative of the operational setting but do not reproduce the confidential database exactly.

Second, the empirical analysis is limited to the Santiago Metropolitan Region. Although this region contains heterogeneous urban and rural areas, it does not capture all geographic, logistical, and administrative conditions present elsewhere in Chile. Consequently, the numerical results should not be

generalized to the national system without additional regional analyses.

Third, the model is deterministic and assumes that applicants, vulnerability scores, costs, capacities, and institutional assignments are known when decisions are made. In practice, attendance, demand, operational costs, and supply conditions may be uncertain. Moreover, the cost structure abstracts from routing, procurement constraints, supplier heterogeneity, menu-specific costs, and detailed labor scheduling. Finally, although applicant-level decisions are linked to location-level costs, the formulation does not capture within-school factors such as daily attendance variability, grade-specific service requirements, or additional administrative rules. These elements constitute relevant extensions for future decision-support models.

9. Conclusions

This paper addresses a public resource allocation problem, specifically the assignment of school meal benefits under budget, logistical, and service-capacity constraints. We develop and evaluate several allocation policies, including a mixed-integer linear programming model that incorporates budget constraints, location-dependent costs, stepwise service-capacity requirements, and applicants' socioeconomic vulnerability.

The results show that alternative policies may perform well with respect to specific metrics, but often at the expense of other relevant dimensions. In contrast, the proposed approach does not necessarily achieve the best value in any single metric; however, it consistently delivers strong performance across all evaluated criteria. This makes it a more robust and balanced solution compared to the benchmark policies.

Furthermore, the results highlight the importance of an appropriate mathematical formulation. Although Policy 4 also relies on an optimization model, it does not achieve the same performance as the proposed approach. The inclusion of the mandatory-priority vulnerability constraint plays a key role, as it ensures that applicants above the predefined vulnerability threshold are selected, while still allowing flexibility in the allocation of the remaining budget. In particular, the model behaves similarly to a ranking-based policy within a critical vulnerability range, while retaining the ability to improve overall efficiency beyond that range.

From a policy perspective, the proposed methodology provides a flexible and practical decision-support tool. It enables decision makers to balance equity and efficiency, ensuring that highly vulnerable applicants are not excluded while increasing overall coverage. In addition, the parameter $\bar{V}$ can be adjusted to control the degree of mandatory protection assigned to highly vulnerable applicants, allowing the model to adapt to different policy objectives and budget scenarios. Since this parameter directly shapes the feasible set and the resulting allocation, its value should be selected as part of the policy-design process and supported by sensitivity analyses.

Overall, the proposed approach demonstrates that incorporating cost structures, logistical considerations, and vulnerability-priority constraints into the allocation process can improve coverage while preserving the prioritization of highly vulnerable applicants in large-scale public benefit programs.

Future research may explore stochastic demand, dynamic budget allocation across periods, or the integration of additional fairness criteria.

A particularly relevant extension is to combine the proposed optimization framework with multi-criteria decision-making methods such as TOPSIS, COPRAS, SAW, or related approaches. The current formulation already incorporates several policy dimensions directly: vulnerability is represented in the objective function, budget and service costs are included in the constraints, location-dependent logistics are explicitly modeled, and equity is protected through the mandatory-priority rule. However, the model does not elicit decision-maker preferences over the relative importance of coverage, cost efficiency, vulnerability prioritization, and baseline alignment, nor does it produce an aggregate ex post ranking of the evaluated policies. A multi-criteria layer could therefore complement the optimization model by ranking feasible allocation policies according to explicitly stated preferences.

Declarations

Ethics approval and consent to participate. Not applicable.

Consent for publication. Not applicable.

Availability of data and materials. The data and materials used in this study are available from the corresponding author, Alejandro Cataldo, upon reasonable request.

Acknowledgement

Not applicable.

Funding

This research was partially supported by the Chilean National Agency for Research and Development (ANID-Fondecyt) [Regular Grant 2023-1231320] and the Natural Sciences and Engineering Research Council of Canada (NSERC) [Grant RGPIN-2018-05225].

References

- [1] ANAHIDEH, H., KANG, L., AND NEZAMI, N. Fair and diverse allocation of scarce resources. *Socio-Economic Planning Sciences* 80 (2022), 101193.
- [2] ARAY, H., AND PACHECO-DELGADO, J. Public investment allocation across ecuadorian provinces. *Socio-Economic Planning Sciences* 71 (2020), 100830.
- [3] AYDEMIR-KARADAG, A., AND AKDERE, E. Design of a distribution network for the school lunch program. *Journal of Engineering Research* 11, 2B (2023).
- [4] BAYCIK, N. O., SHARKEY, T. C., AND RAINWATER, C. E. A markov decision process approach for balancing intelligence and interdiction operations in city-level drug trafficking enforcement. *Socio-Economic Planning Sciences* 69 (2020), 100700.
- [5] BERTSIMAS, D., FARIAS, V. F., AND TRICHAKIS, N. The price of fairness. *Operations Research* 59, 1 (2011), 17–31.
- [6] BOATMAN, A., AND LONG, B. T. Does financial aid impact college student engagement? evidence from the gates millennium scholars program. *Research in Higher Education* 57, 6 (2016), 653–681.
- [7] BOUZARTH, E., FORRESTER, R., HUTSON, K., AND REDDOCH, L. Assigning students to schools to minimize both transportation costs and socioeconomic variation between schools. *Socio-Economic Planning Sciences* 64 (2018), 1–8.
- [8] CATALÁN, J., EPSTEIN, R., GUAJARDO, M., YUNG, D., AND MARTINEZ, C. Solving multiple scenarios in a combinatorial auction. *Computers & Operations Research* 36, 10 (2009), 2752–2758.
- [9] CATMA, S. The management of scholarship programs at universities: a cost-benefit analysis. *International Journal of Management* 29, 4 (2012), 519.
- [10] CONTRERAS, I., AND LOZANO, S. Allocating additional resources to public universities. a dea bargaining approach. *Socio-economic planning sciences* 71 (2020), 100752.
- [11] CRAFT, E. When a traditional scholarship is simply not enough: Addressing the digital divide to recruit and motivate engineering technology students through graduation. In *2016 ASEE Annual Conference & Exposition* (2016).
- [12] DELIMA, A. Predicting scholarship grants using data mining techniques. *Int. J. Mach. Learn. Comput* 9 (2019), 513–519.

- [13] DURÁN, G., AND WOLF-YADLIN, R. A mathematical programming approach to applicant selection for a degree program based on affirmative action. *Interfaces* 41, 3 (2011), 278–288.
- [14] EPSTEIN, R., HENRÍQUEZ, L., CATALÁN, J., WEINTRAUB, G. Y., AND MARTÍNEZ, C. A combinational auction improves school meals in chile. *Interfaces* 32, 6 (2002), 1–14.
- [15] GOBIERNO DE CHILE. Presupuesto para el desarrollo inclusivo y sostenible, 2022.
- [16] GRANILLO-MACÍAS, R. Logistics optimization through a social approach for food distribution. *Socio-Economic Planning Sciences* 76 (2021), 100972.
- [17] GROSS, J. P. K., HOSSLER, D., AND ZISKIN, M. Institutional merit-based aid and student persistence: An analysis of effects. *The Journal of Student Financial Aid* 37, 1 (2007), 28–39.
- [18] HUANG, D., GU, Y., WANG, H., LIU, Z., AND CHEN, J. An incentive dynamic programming method for the optimization of scholarship assignment. *Discrete Dynamics in Nature and Society* 2018 (2018).
- [19] KARAKAYA, Ş., AND MERAL, S. A biobjective hierarchical location-allocation approach for the regionalization of maternal-neonatal care. *Socio-Economic Planning Sciences* 79 (2022), 101093.
- [20] LATUMAKULITA, L. A., AND USAGAWA, T. Indonesia scholarship selection model using a combination of back-propagation neural network and fuzzy inference system approaches. *International Journal of Intelligent Engineering and Systems* 11, 3 (2018), 79–90.
- [21] LOZANO, S., AND CONTRERAS, I. Centralised resource allocation using lexicographic goal programming, application to the spanish public university system. *Socio-Economic Planning Sciences* 84 (2022), 101419.
- [22] MAHADEVAN, B., SIVAKUMAR, S., DINESH KUMAR, D., AND GANESHRAM, K. Redesigning midday meal logistics for the akshaya patra foundation: Or at work in feeding hungry school children. *Interfaces* 43, 6 (2013), 530–546.
- [23] MAYERLE, S. F., RODRIGUES, H. F., DE FIGUEIREDO, J. N., AND CHIROLI, D. M. D. G. Optimal student/school/class/teacher/classroom matching to support efficient public school system resource allocation. *Socio-Economic Planning Sciences* 83 (2022), 101341.
- [24] MOLINA, C. G., ET AL. *Gasto social en América latina*. Departamento de Integración y Programas Regionales, Instituto Interamericano . . . , 2003.
- [25] PUSPITASARI, T., SARI, E., DESTARIANTO, P., AND RISKIawan, H. Decision support system for determining scholarship selection using an analytical hierarchy process. In *Journal of Physics: Conference Series* (2018), vol. 953, IOP Publishing, p. 012119.
- [26] REDONDO, S., CATALDO, A., MARQUINEZ, J. T., REY, P. A., AND SAURÉ, A. Improving scholarship assignment using approximate dynamic programming: A chilean case study. *Socio-Economic Planning Sciences* (2025), 102296.
- [27] SOUZA, J. S., LIM-APO, F. A., VARELLA, L., COELHO, A. S., AND SOUZA, J. C. Multi-period optimization model for planning people allocation in shelters and distributing aid with special constraints. *Socio-Economic Planning Sciences* 79 (2022), 101087.
- [28] SPARROW, R. Protecting education for the poor in times of crisis: An evaluation of a scholarship programme in indonesia. *Oxford Bulletin of Economics and Statistics* 69 (2007), 99–122.
- [29] WIMATSARI, G. A. M. S., PUTRA, I. K. G. D., AND BUANA, P. W. Multi-attribute decision making scholarship selection using a modified fuzzy topsis. *International Journal of Computer Science Issues (IJCSI)* 10, 1 (2013), 309.
- [30] YE, H.-J., ZHAN, D.-C., LI, X., HUANG, Z.-C., AND JIANG, Y. College student scholarships and subsidies granting: A multi-modal multi-label approach. In *2016 IEEE 16th International Conference on Data Mining (ICDM)* (2016), 559–568.

A. Benchmark Policy Pseudocode

The following pseudocode is included as supporting methodological documentation.

Algorithm 1. Policy 1

-
- 1: Start with the set of selected applicants empty ($\mathcal{S} = \emptyset$) and let $\mathcal{I}$ be the set of all applicants.
 - 2: Let $\mathcal{P}$ denote the available budget, and let $\mathcal{C} = 0$ be the current allocation cost.
 - 3: Construct a priority queue $\mathcal{D}$ from $\mathcal{I}$ using the vulnerability index.
 - 4: **while** $\mathcal{D} \neq \emptyset$ **do**
 - 5: Extract the first applicant from queue $\mathcal{D}$.
 - 6: Create a copy $\mathcal{T}$ of set $\mathcal{S}$ and add the extracted applicant.
 - 7: Compute the allocation cost $\mathcal{C}$: direct service cost + fixed location-activation cost + stepwise service-capacity cost.
 - 8: **if** $\mathcal{C} \leq \mathcal{P}$ **then**
 - 9: Add the applicant to set $\mathcal{S}$.
 - 10: **else**
 - 11: Terminate the allocation loop because there is not enough budget to select the applicant.
 - 12: **end if**
 - 13: **end while**
 - 14: Return the set $\mathcal{S}$ of selected applicants.
-

Algorithm 2. Policy 2

-
- 1: Start with the set of selected applicants empty ($\mathcal{S} = \emptyset$) and let $\mathcal{I}$ be the set of all applicants.
 - 2: Let $\mathcal{P}$ denote the available budget, and let $\mathcal{C} = 0$ be the current allocation cost.
 - 3: Create one applicant set for each educational institution $n \in \mathcal{N}$.
 - 4: **while** $\mathcal{I} \neq \emptyset$ **do**
 - 5: Extract one applicant from each nonempty institution set and store the extracted applicants in list $\mathcal{E}$.
 - 6: For each applicant in $\mathcal{E}$, compute the marginal assignment cost.
 - 7: Identify the applicant with the lowest marginal assignment cost.
 - 8: **if** the resulting total cost does not exceed $\mathcal{P}$ **then**
 - 9: Add the corresponding applicant to $\mathcal{S}$ and remove the applicant from the remaining applicant and institution sets.
 - 10: **else**
 - 11: Terminate the allocation loop.
 - 12: **end if**
 - 13: Return all remaining applicants in $\mathcal{E}$ to their corresponding institution sets.
 - 14: **end while**
 - 15: Return the set $\mathcal{S}$ of selected applicants.
-

Algorithm 3. Policy 3

- 1: Start with the set of selected applicants empty ($\mathcal{S} = \emptyset$) and let $\mathcal{I}$ be the set of all applicants.
 - 2: Let $\mathcal{P}$ denote the available budget, and let $\mathcal{C} = 0$ be the current allocation cost.
 - 3: Create one applicant set for each educational institution $n \in \mathcal{N}$.
 - 4: Construct a priority queue $\mathcal{D}$ from $\mathcal{I}$ using the vulnerability index.
 - 5: **while** $\mathcal{I} \neq \emptyset$ **do**
 - 6: Extract the most vulnerable remaining applicant from queue $\mathcal{D}$.
 - 7: Compute the allocation cost after tentatively adding the applicant to $\mathcal{S}$.
 - 8: **if** the resulting total cost does not exceed $\mathcal{P}$ **then**
 - 9: Add the applicant to $\mathcal{S}$ and remove the applicant from $\mathcal{I}$.
 - 10: **for** $K_n - 1$ iterations **do**
 - 11: Extract the remaining applicant with the highest V_i from the same educational institution.
 - 12: Compute the allocation cost after tentatively adding the applicant.
 - 13: **if** the resulting total cost does not exceed $\mathcal{P}$ **then**
 - 14: Add the applicant to $\mathcal{S}$ and remove the applicant from queue $\mathcal{D}$.
 - 15: **else**
 - 16: Terminate the allocation loop.
 - 17: **end if**
 - 18: **end for**
 - 19: **else**
 - 20: Terminate the allocation loop.
 - 21: **end if**
 - 22: **end while**
 - 23: Return the set $\mathcal{S}$ of selected applicants.
-