Minimax Theory and its Applications
Volume 11 (2026), No. 3s, 956 to 968

Game-Theoretic minimax optimization of educational resource allocation under cultural diversity constraints

Havaskhon Rakhmatullayeva[1], Baxtigul Nurullayeva[2], Nirav Jadav[3], Tursunoy Akhmedova[4], Fazilat Turgunova[5], Mohammed Shamsul Hoque[6], Asilabonu Azizjonova[7]

[1]Lecturer, Department of Pedagogy and Psychology, Faculty of Education, Kokand University, Kokand, Fergana Region, Uzbekistan. havasonbaxromovna@gmail.com, https://orcid.org/0009-0009-4643-0149
[2]Professor, Doctor of Sciences (DSc), Mamun University, Xorazm, Uzbekistan. nurullayeva_baxtigul@mamunedu.uz; lady.bakhtigul@mail.ru, https://orcid.org/0009-0001-8976-3105

[3]Database Architect, Intuitive.AI, Palm Beach Gardens, United States. njadav@asu.edu, https://orcid.org/0009-0000-9316-5506

[4]Tashkent Pharmaceutical Institute, Tashkent, Uzbekistan. tursunoychik81@gmail.com, https://orcid.org/0009-0008-7464-530X

[5]Doctor of Philosophy in Art Studies Kamoliddin Behzod National Institute of Fine Arts and Design, Tashkent, Uzbekistan. azganush1990@gmail.com, https://orcid.org/0009-0002-4982-1216
[6]Professor, Department of English Southeast University, Dhaka, Bangladesh. dr.hoque@seu.edu.bd, https://orcid.org/0000-0001-5695-8742

[7]Department of Primary Education Methodology, Termiz University of Economics and Service, Uzbekistan. asilabonu_azizjonova@tues.uz, https://orcid.org/0009-0009-9933-8290

Abstract:

Resource allocation in culturally and linguistically diverse systems is a problem in conditions of deep uncertainty, where it is uncertain by how much the performance increases from each additional hour of remediation, mother tongue reading materials, digital devices and community liaisons, depending on the learner's cultural and linguistic background, which is the most difficult to measure quantity in the evidence base. Formulate the problem as a two-player zero-sum game in which a planner allocates a budget over schools and learner groups with four categories of resources while nature chooses both the learner's cultural-linguistic background and the resource effectiveness scenario. There are three constraints on cultural diversity in the system, a representation constraint, a mother tongue reading provision constraint and a minimum community liaison requirement, defining the planner's feasible set. The existence of the Nash equilibrium follows from the Sion's theorem, the planner's equilibrium policy is a deterministic plan, and the equilibrium can be found exactly through solving a single exponential cone program, the Lagrange multipliers of which give nature's equilibrium policy. A multiplicative weights learning algorithm converges to the same optimal solution with an optimality certificate in each round. On a 30-school district with 16,116 students split into five groups with eight effectiveness regimes, the GMDC achieved a better guaranteed minimum-group performance by -0.440 versus -0.278 standard deviations above the worst group's best performance under the utilitarian solution, increased the out-of-sample 5th percentile worst-group performance by 0.133 SD, and narrowed the average between-group gap by 39% in return for a loss of only 0.021 SD in district-mean performance. Diversity floors on top of a utilitarian objective did not change anything, while minimax without diversity floors left out the majority group altogether. All 20 additional districts examined gave identical findings (Wilcoxon $p<10 \wedge (-5)$). The nature's equilibrium strategy selected one key assumption, that mother-tongue material is effective in three highest distance groups, turning the robust plan into an agenda for policy makers.

Keywords: Minimax optimisation, zero-sum games, robust resource allocation, educational equity, cultural diversity, mother-tongue education, rawlsian fairness, multiplicative weights.

1. Introduction

Educational administrators distribute budgets among schools and resource types, such as additional teaching time, learning materials, access to technology, and community outreach. Strategic optimization models are useful tools for efficient education resource management and decision making (Banerjee et al., 2007). Under culturally and linguistically diverse circumstances, the productivity of inputs is heterogeneous among different learner communities (Bertsimas et al., 2011a). Remedial classes conducted in the dominant language can be less productive for the learners that speak other languages. Materials written in the mother tongue become more valuable with increasing linguistic distance, while technology-based learning requires connectivity, especially in indigenous and rural settings (Kraft, 2020; Muralidharan et al., 2019). Educational participation and inclusion are impacted by cultural responsiveness and community involvement (Bertsimas et al., 2011a).

Two main issues exist in traditional education resource allocation.

1. Culture-specific uncertainty: Productivity of education interventions is context-dependent, which makes the transportability of estimated effects uncertain. Allocating budgets using point estimates can be inefficient due to uncertainty about culture-specific productivity (Diamond & Boyd, 2016; Ogryczak, 2000).

2. Equity constraints: Allocation guided by utilitarian principle maximizes the aggregate achievement but tends to favor those groups of people that have higher expected gains. Such an approach can exacerbate educational inequities unless it incorporates equity constraints explicitly (Delage & Ye, 2010; Duan, 2025).

The suggested allocation framework creates a two-player, zero-sum game where the player is the planner, and the opponent is nature. The planner allocates the fixed budget among the schools, culture, and resources while respecting diversity requirements. Nature picks the worst possible situation with respect to the selected culture-linguistic group and effectiveness case. The created Game-Theoretic Minimax Optimization under Diversity Constraints (GMDC) ensures the achievement of minimum gains (Arora et al., 2012; Bertsimas et al., 2011b).

The contributions of this paper are as follows.

1. An instance of the game where the payoff function is concave in the planner's allocation and linear in nature's mixed strategy, such that there is a Nash equilibrium which is the same as the max-min solution by virtue of Sion's theorem. Planner's equilibrium is a pure allocation; nature's equilibrium is a mixed strategy indicating which groups and scenarios are binding.

2. Three diversity constraints of cultural significance, all linear, whose meanings are policy relevant: a representation floor based on proportion of population, a mother tongue provision floor as weighted by linguistic distance and a minimum community liaison presence requirement for high linguistic distance groups.

3. An exact epigraph conic program where the dual variables correspond to nature's equilibrium strategy, and a multiplicative weight learning algorithm, both of which yield the same objective value and offer a duality-gap proof at each iteration.

4. Computational experiments in a 30-school district with five cultural-linguistic groups and eight effectiveness scenarios, comparing GMDC to seven other allocation rules on criteria of achievement of worst group, average achievement, achievement gap and robustness out-of-sample, along with sensitivity analysis of budget and constraints in twenty replications.

Section 2 presents previous research on the topic. The model is provided in Section 3 and its game-theoretic analysis in Section 4. Section 5 explains the computational experiment while Section 6 describes its results. The policy implications and limitations are discussed in Section 7, and Section 8 concludes.

2. Related work

2.1. Minimax decisions and games against nature

The minimax theorem for finite zero-sum games provides the basis for decision making under adversarial situations (Arora et al., 2012). The statistical decision theory extends this concept to the game against nature, wherein decisions minimize maximum risks (Bertsimas et al., 2011b). Robust optimization deals with the uncertainty using bounded uncertainty sets, giving decision makers control over the amount of conservatism and worst-case performances (Bertsimas & Sim, 2004; Diamond & Boyd, 2016). Distributionally robust optimization extends robust optimization by dealing with uncertainties in the distribution of parameters instead of the actual values of the parameters (Wald, 1945). The online learning method offers a constructive way of equilibrium determination through multiplicative weights updating and best response actions leading to the value of zero-sum games (Hanushek & Woessmann, 2008).

2.2. Fairness and equity in allocation

The maximin criterion takes precedence over the well-being of the most disadvantaged people. The equitable optimization approach takes into account inequality measurements and max-min fairness criteria to ensure an effective distribution of resources among conflicting populations (von Neumann, 1928). The price of fairness is defined as the cost to aggregate utility that is incurred when fairness constraints are imposed on an optimization problem (Duan, 2025). Inequality in education is also another factor in formulating policies, thus demonstrating the need to incorporate distributional equity in institutional resource allocations (Delage & Ye, 2010).

2.3. Culture, language and the effectiveness of educational resources

Culturally responsive education acknowledges the dependence of learning efficiency on the language and culture of learners. The inclusive education policies from various cultural settings underscore equal participation and differentiated education (Bertsimas et al., 2011b). Interventions in education, such as remediation, reveal significant improvements in learning efficiency, despite varying effectiveness across learner groups (Kraft, 2020). Instruction using technology may enhance educational attainment; however, its effectiveness may vary based on specific conditions and availability of learning resources (Muralidharan et al., 2019). Optimization-based educational management offers supplementary help for efficient planning of learning resources (Banerjee et al., 2007). The interpretation of standardized intervention effect sizes is crucial for assessing educational performance across different settings (Ogryczak, 2000). Economic development is another aspect influenced by educational attainment due to the impact of resource misallocation in the future (Isroilov et al., 2025). Computational realization of the suggested constrained optimization framework may be performed using convex optimization models like CVXPY (Uchaipichat, 2025).

2.4. Gap addressed

In educational resource allocation models, the approach usually taken is to employ point estimates for effectiveness and handle equity via ex post reporting. These are not aware of any existing model that treats effectiveness uncertainty in terms of culture as an adversary, incorporates diversity constraints that have legal significance, and reports nature's equilibrium strategy as a diagnostic measure. GMDC tackles this combination.

3. Model

3.1. Schools, groups and resources

A district has S schools indexed by s and C cultural-linguistic groups indexed by c. School s enrols $n_{sc} \geq 0$ learners of group c; $N_c = \sum_s n_{sc}$ is the district total of group c and $N = \sum_c N_c$ the overall total. Each group has a cultural-linguistic distance $d_c \in [0,1]$ from the language and cultural norms of the curriculum, with $d_c = 0$ for the majority-language group. Each school has a baseline mean achievement b_{sc} for each group, in district standard-deviation (SD) units, and a connectivity index $\omega_s \in (0,1]$.

The planner allocates a budget B over $R = 4$ resource types (Table 1). The decision is $x_{rsc} \geq 0$, the budget spent on resource r for group c at school s, with $x_{rsc} = 0$ wherever $n_{sc} = 0$. It is common practice when creating remedial programs and language services to target assistance to different groups within the school.

Table 1: Resource types and cost scales.

Index	Resource	Cost scale κ_r (budget units per learner per unit dose)

$r = 1$	Targeted teaching hours (remedial and small-group instruction)	1.0
$r = 2$	Mother-tongue and bilingual learning materials	0.4
$r = 3$	Digital learning access (devices, adaptive software)	0.8
$r = 4$	Community liaison and culturally responsive pedagogy (CRP) training	0.5

3.2. Learning-response function

Per-learner resource r dose equals the expenditure per learner in relation to its scale of cost, while learning depends on dose with diminishing returns. Achievement mean for group c under effectiveness scenario k is

$$A_{kc}(x) = \sum_{s=1}^{S} \frac{n_{sc}}{N_c} \left[b_{sc} + \sum_{r=1}^{4} \theta_{k,cr}\ \rho_{sr} \ln\left(1 + \frac{x_{rsc}}{\kappa_r\, n_{sc}}\right) \right] \qquad (1)$$

Where $\theta_{k,cr}$ is the achievement gain per unit of log-dose, $\rho_{s3} = \omega_s$ scales digital effectiveness by connectivity, and $\rho_{sr} = 1$ for $r \neq 3$. The logarithm captures diminishing returns and makes A_{kc} concave in x.

Nominal effectiveness is embedded with four patterns found in the literature analyzed in Section 2.3: general teaching has less effect when there is a language mismatch; mother tongue has more effect as distance increases; digital effectiveness decreases as distance increases; and community relations matter when the distance is high:

$$\theta_{c1} = 0.18\,(1 - \gamma d_c), \quad \theta_{c2} = 0.06\,(0.3 + 2.2\,d_c), \quad \theta_{c3} = 0.10\,(1 - 0.4\,d_c), \quad \theta_{c4} = 0.05\,(0.5 + 1.5\,d_c) \qquad (2)$$

With γ for nominal mismatch being equal to 0.6. The scenarios k adjusts the coefficients in the way specified in Table 2. The scenarios 2-8 reduce the value of resources that a culturally-informed strategy depends upon for a particular group; collectively, they are an illustration of the uncertainties that the planner encounters regarding transfer across cultures.

Table 2: Effectiveness scenarios available to nature.

k	**Scenario**	**Change from nominal**
1	Nominal	none
2	Severe language mismatch	$\gamma = 0.95$
3	Weak mother-tongue uptake	$\theta_{c2} \times 0.50$
4	Low digital uptake	$\theta_{c3} \times 0.40$
5	Teacher attrition	$\theta_{c1} \times 0.60$
6	Weak community linkage	$\theta_{c4} \times 0.45$
7	Mismatch and low digital uptake	$\gamma = 0.85$, $\theta_{c3} \times 0.60$
8	Mother-tongue and liaison shortfall	$\theta_{c2} \times 0.75$, $\theta_{c4} \times 0.75$

3.3. Feasibility and cultural diversity constraints

The feasible set $\mathcal{X}(B)$ contains every allocation $x \geq 0$ that satisfies the following constraints.

- Budget: $\sum_{r,s,c} x_{rsc} \leq B$.
- Staffing ceiling: $\sum_c x_{1sc} \leq 0.36\,B\,n_s/N$, where $n_s = \sum_c n_{sc}$, reflecting the number of teachers a school can absorb.
- (D1) Representation floor: $\sum_{r,s} x_{rsc} \geq \beta\,(N_c/N)\,B$ for every group c. No group receives less than a fraction β of its population share; by default, $\beta = 0.5$.
- (D2) Mother-tongue floor: $x_{2sc} \geq \alpha \cdot 0.25\,B \cdot n_{sc} d_c / \sum_{s',c'} n_{s'c'}\, d_{c'}$. Mother-tongue provision is guaranteed in proportion to linguistic need; the default $\alpha = 0.5$ reserves 12.5% of the budget.
- (D3) Liaison presence: $\sum_s x_{4sc} \geq 0.03\,B$ for every group with $d_c > 0.3$.
- All constraints are linear in (x, B), so $\mathcal{X}(B)$ is a polyhedron.

3.4. The game

The planner chooses $x \in \mathcal{X}(B)$. Nature chooses a scenario–group pair (k, c), or a mixed strategy $\mu \in \Delta_{K\times C}$ over the $K \times C$ pairs. The payoff to the planner is

$$U(x, \mu) = \sum_{k=1}^{K} \sum_{c=1}^{C} \mu_{kc}\ A_{kc}(x) \qquad (3)$$

And nature receives $-U(x, \mu)$. Choice of group by nature reflects the Rawlsian principle that an allocation should be judged based on its least advantaged group (Hanushek & Woessmann, 2008); choice of scenario reflects culturally specific effectiveness ambiguity. GMDC allocation is the planner's equilibrium strategy.

4. Game-Theoretic analysis and solution method

4.1. Existence and characterization of equilibrium

Theorem 1. The game of Section 3.4 has a Nash equilibrium (x^*, μ^*), the planner's strategy x^* can be taken pure, and the value of the game is

$$V(B) = \max_{x\in\mathcal{X}(B)} \min_{\mu\in\Delta_{K\times C}} U(x, \mu) = \min_{\mu\in\Delta_{K\times C}} \max_{x\in\mathcal{X}(B)} U(x, \mu) = \max_{x\in\mathcal{X}(B)} \min_{k,c} A_{kc}(x) \qquad (4)$$

Proof. Each A_{kc} is a constant plus a nonnegative combination of terms $\ln(1 + \text{affine})$, so it is concave and continuous in x; hence U is concave in x and linear in μ. $\mathcal{X}(B)$ is a nonempty compact polyhedron, being bounded by the budget, and the simplex is compact and convex. Sion's minimax theorem implies equality of max-min and min-max value and existence of a saddle point. Linear function on the simplex is minimised at a vertex, hence the final formula. Planner's equilibrium strategy is the maximiser of a concave function, thus pure.

The planner need never randomize, which is important for public policy because a published allocation constitutes one plan only. On the other hand, the equilibrium strategy of nature is a mixture and diagnostic.

4.2. Exact computation and nature's strategy as dual variables

The equilibrium allocation solves the epigraph program

$$\max_{x,\,t} t \quad \text{s.t.} \quad t \le A_{kc}(x)\ \ \forall (k, c), \qquad x \in \mathcal{X}(B) \qquad (5)$$

This being an example of a convex exponential cone program with KC = 40 epigraph constraints and RSC = 600 allocation variables.

Proposition 1. Let $\lambda_{kc} \ge 0$ be optimal Lagrange multipliers of the epigraph constraints in (5). Then $\sum_{k,c} \lambda_{kc} = 1$, and $\mu^* = \lambda$ is an equilibrium strategy of nature. By complementary slackness, $\mu^*_{kc} > 0$ only for pairs with $A_{kc}(x^*) = V(B)$.

Proof. Strong duality holds because (5) is convex and strictly feasible. Stationarity of the Lagrangian in t gives $1 - \sum_{k,c} \lambda_{kc} = 0$. Stationarity in x states that x^* maximises $\sum_{k,c} \lambda_{kc}\, A_{kc}(x)$ over $\mathcal{X}(B)$, so x^* is a best response to λ. Complementary slackness confines the support of λ to binding pairs, on which x^* attains the minimum, so λ is a best response to x^*.

The support of μ^* names the groups and assumptions that determine the plan. A decision-maker can read it directly: these are the populations and scenarios where better evidence would change the allocation.

4.3. Learning dynamics and a certificate of optimality

In addition, nature plays the multiplicative weights game while the planner responds optimally to nature's actions:

1. Initialise $\mu^{(1)}$ uniform over the KC scenario–group pairs.
2. In round t, the planner computes $x^{(t)} = \text{argmax}_{x \in \mathcal{X}(B)} \sum_{k,c} \mu_{kc}^{(t)} A_{kc}(x)$, a concave program.
3. Nature updates $\mu_{kc}^{(t+1)} \propto \mu_{kc}^{(t)} \exp\left(-\eta\, A_{kc}\left(x^{(t)}\right)\right)$.
4. After T rounds, report the averages $\bar{x}_T = \frac{1}{T}\sum_t x^{(t)}$ and $\bar{\mu}_T = \frac{1}{T}\sum_t \mu^{(t)}$.

Proposition 2. With step size $\eta = O\left(\sqrt{\ln(KC)/T}\right)$, the averaged strategies form an ϵ_T-equilibrium with $\epsilon_T = O\left(\sqrt{\ln(KC)/T}\right)$ [11, 12]. Moreover, at every round

$$L_T = \min_{k,c} A_{kc}(\bar{x}_T) \;\le\; V(B) \;\le\; \max_{x \in \mathcal{X}(B)} U(x, \bar{\mu}_T) = U_T \qquad (6)$$

so, the gap $U_T - L_T$ certifies how far the averaged plan is from optimal. The lower bound holds because $\bar{x}_T$ is feasible, and the upper bound because nature can always play $\bar{\mu}_T$.

4.4. Comparative statics

Proposition 3. $V(B)$ is nondecreasing and concave in B. Relaxing any diversity constraint weakly increases $V(B)$.

Proof. $\mathcal{X}(B)$ grows with B, so V is nondecreasing. For budgets B_1, B_2 with optimal plans x_1, x_2 and $\vartheta \in [0,1]$, the combination $\vartheta x_1 + (1-\vartheta)x_2$ is feasible for $\vartheta B_1 + (1-\vartheta)B_2$ because every constraint is linear in (x, B). Concavity of $\min_{k,c} A_{kc}$ then gives $V(\vartheta B_1 + (1-\vartheta)B_2) \ge \vartheta V(B_1) + (1-\vartheta)V(B_2)$. Removing constraints enlarges the feasible set.

The two-efficiency metrics. The first one is called the price of robustness (PoR), which is the gap between the achievement of the district under the robust solution and the achievement of the district under the expected value solution in SD units. The second one is the price of diversity constraints (PoD).

4.5. Implementation

Models were implemented using CVXPY and solved using the Clarabel interior-point conic solver. The time taken to compute one equilibrium of the problem is around 3 seconds for the two-core processor.

5. Computational study design

5.1. District instance

The study instance is a district of $S = 30$ schools serving $N = 16{,}116$ learners in $C = 5$ cultural-linguistic groups. This is because the district has been modeled on a parametric model with random seeds of 7. This implies that all the values can be replicated and changed. The schools are rural with 0.55 probability, making 16 rural schools and 14 urban schools that together enroll 58.0% of the learners in rural areas in Table 3 and Table 4. The group composition is based on a Dirichlet distribution that assigns higher weights to indigenous students in rural schools and migrant students in urban schools. Baseline achievement is $b_{sc} = -0.55\, d_c - 0.15 \cdot \mathbb{1}[s \text{ rural}] + e_{sc}$ with $e_{sc} \sim \mathcal{N}(0, 0.08^2)$. Connectivity is drawn from $\text{Beta}(2,4)$ in rural schools and $\text{Beta}(5,2)$ in urban schools, giving mean connectivity of 0.26 and 0.72 respectively.

Table 3: Learner groups in the study district.

Group	Distance d_c	Learners	Share	Schools present	Baseline mean (SD)
Majority-language	0.00	9,307	57.7%	30	−0.086
Regional minority A	0.35	2,459	15.3%	29	−0.302
Regional minority B	0.45	1,762	10.9%	28	−0.311
Indigenous/tribal	0.80	876	5.4%	22	−0.580
Migrant	0.65	1,712	10.6%	25	−0.373

The initial gap between the majority language and the indigenous communities is 0.49 standard deviation (SD), similar to the academic achievement gaps that have been found for majority and minority language learners in assessments. The default budget is B=0.6 budget units per student, totalling 9,670 budget units.

5.2. Allocation rules compared

Table 4: Allocation rules compared.

Rule	Objective	D1–D3
Proportional	Equal spend per learner; fixed resource split 45/15/25/15%	no
Need-based	Spend proportional to $n_{sc}\max(0.1 - b_{sc}, 0)$; same split	no
Util-Nominal	Maximise district mean achievement under the nominal scenario	no
Util-SAA	Maximise district mean achievement averaged over the eight scenarios	no
Util-SAA+div	Util-SAA subject to D1–D3	yes
Rawls-Nominal+div	Maximise the worst group under the nominal scenario only	yes
Minimax (no div)	Game value (4) without D1–D3	no
GMDC (proposed)	**Game value (4) subject to D1–D3**	**yes**

5.3. Evaluation protocol

Each plan is scored in four ways.

1. In-scenario guarantee: $\min_{k,c} A_{kc}(x)$ over the eight scenarios of Table 2.

2. Out-of-sample robustness: $M = 4{,}000$ effectiveness draws not seen by any optimiser. Each draw is a mixture $\sum_k \pi_k \Theta_k$ of the eight scenario coefficient matrices with weights $\pi \sim \text{Dirichlet}(0.5 \cdot 1)$, after which every coefficient is multiplied by an independent factor e^z with $z \sim \mathcal{N}(0, 0.15^2)$. The average value of the district means, the average value of the worst group mean, and the difference between the best-off and worst-off groups will be obtained. It will provide average values over draws and 5th percentiles of the worst group mean.

3. Composition: share of the budget by groups, by types of resources and by rural areas, and expenditure per pupil by groups.

4. Replication: The full comparison is repeated for 20 additional samples drawn from the same parametric distribution, with seeds 100-119, with M=1,000 samples from out-of-sample distributions. The paired comparisons are statistically assessed using two-sided Wilcoxon signed rank tests.

The sensitivity analysis is performed for the range 0.2 to 1.2 for the budget of each individual, the value of β ranging from 0 to 1 and the value of α ranging from 0 to 1.5. The equilibrium is also calculated using 40 rounds of the learning process of Section 4.3.

6. Results and analysis

6.1. Main comparison

The eight rules are listed in table 5 in their entirety. GMDC offers an absolute worst-group mean of -0.278 SD in all eight cases as opposed to -0.440 SD for the best utilitarian rule and -0.527 SD for proportional allocations. Exogenously, the 5th percentile performance of the worst group is -0.288 SD, which is 0.133 SD better than Util-SAA. The penalty of robustness is small: the district mean of achievement goes down from -0.126 to -0.146 SD, a loss of PoR=0.021 SD.

Table 5: Performance of allocation rules on the study district (achievement in sd units; higher is better except for the gap). out-of-sample columns average 4,000 effectiveness draws.

Rule	Guarantee $\min_{k,c} A_{kc}$	Nominal worst group	OOS district mean	OOS worst-group mean	OOS worst-group P5	OOS group gap	Majority budget share (%)	Indigenous budget share (%)
Proportional	−0.527	−0.515	−0.145	−0.522	−0.530	0.492	57.8	5.4
Need-based	−0.475	−0.449	−0.146	−0.464	−0.480	0.413	36.0	12.2

Util-Nominal	−0.441	−0.361	**−0.126**	−0.380	−0.422	0.349	37.2	14.1
Util-SAA	−0.440	−0.360	**−0.126**	−0.379	−0.421	0.348	37.5	14.2
Util-SAA+div	−0.441	−0.361	**−0.126**	−0.379	−0.421	0.349	37.5	14.2
Rawls-Nominal+div	−0.340	−0.237	−0.139	−0.277	−0.324	0.217	28.9	30.0
Minimax (no div)	**−0.246**	**−0.234**	−0.147	**−0.245**	**−0.256**	**0.160**	0.0	49.7
GMDC	−0.278	−0.266	−0.146	−0.275	−0.288	0.214	28.9	40.9

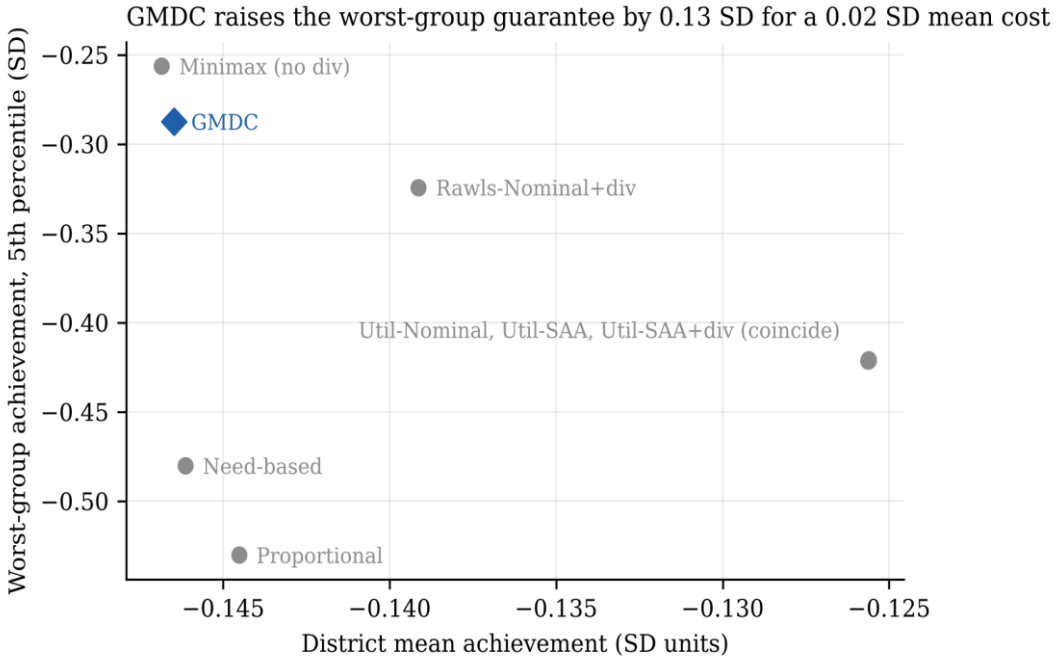


Figure 1: Trade-off between district-mean achievement and the 5th-percentile worst-group achievement across 4,000 out-of-sample effectiveness draws. the three utilitarian rules produce nearly identical plans.

Four findings stand out

However, diversity constraints alone do not sufficiently take care of the worst-off group. D1-D3 constraints do not affect Util-SAA because it satisfies all of them and still retains a guarantee of -0.440. Floors constraint how much a group is getting, but not how good the allocation is in terms of its cultural relevance. The benefit is from Minimax, which locates resources where their effectiveness is preserved in scenarios harmful for each group.

It is better to be robust to scenarios than to optimize a nominal case. Rawls-Nominal+div maximizes the worst group in terms of nominal effectiveness and gives the best value of worst group among constrained rules: -0.237. However, its scenario guarantee is -0.340 and out-of-sample 5th percentile is 0.036 SD below GMDC. Optimal design based on point estimates of culturally specific effectiveness has a weakness exactly where there is the least evidence.

Unconstrained minimax does not include the majority group. If we remove D1–D3, the minimax program will allocate zero to learners from the majority language group (Table 5, Figure 1). This is the famous problem of pure maximin criteria – when marginal returns for some groups are very low, they receive nothing. Representation threshold avoids this problem, but at the cost of having diversity constraint of PoD=0.032 SD.

Utilitarian programs underserve the group with the largest gap. In Util-SAA, the indigenous group ends up with -0.379 SD, 0.35 SD lower than the majority-language group. GMDC improves it to -0.229 SD, surpassing both regional minority groups and migrants (Figure 2).

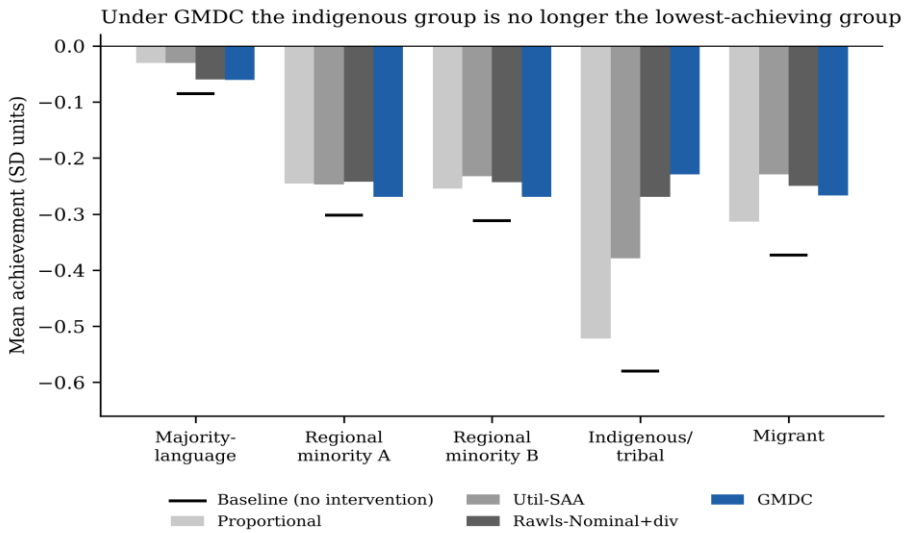


Figure 2: Mean post-intervention achievement by group (out-of-sample average). Black marks show baseline achievement without intervention.

6.2. Nature's equilibrium strategy

These duals define the optimal policy of nature under the equilibrium. Under the benchmark budget, nature assigns weight 0.408 to the indigenous group, weight 0.353 to the migrant group and weight 0.238 to the regional minority B, all under scenario 3, weak mother-tongue materials. The weights are zero for all other combinations of scenario and group. This means that the GMDC policy is constrained by the efficiency of the mother-tongue materials for the three groups with largest distances to the native language, and a revelation about this one parameter would alter the weights more than any other. In the absence of diversity requirements, the constraint set is wider: it includes all four minority groups under scenario 3 and both regional minorities under scenario 5, teacher attrition.

The learning process of Section 4.3 is in line with the precise answer. After 40 periods, the confidence interval (6) is [–0.2803, –0.2729] with a margin of 0.0074 SD and the true value V(B)=–0.2784 is contained within that interval (Figure 3). The average nature's strategy after 40 periods puts 62% of the total probability weight on scenario 3 and 14% on scenario 8, while 58% of the total weight goes to scenario 3 for the three farthest away groups.

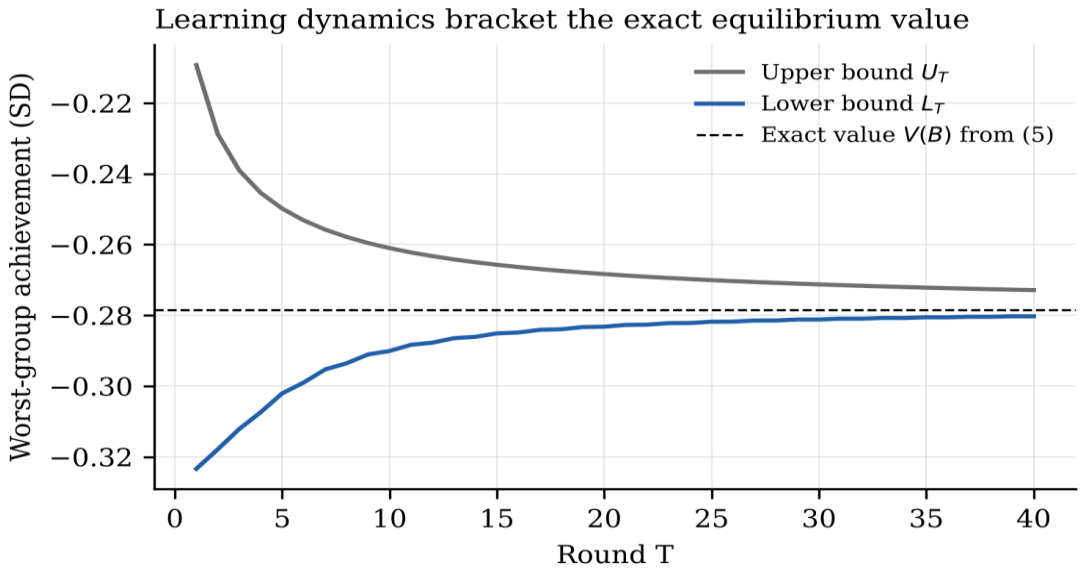


Figure 3: Lower and upper bounds of the multiplicative-weights dynamic and the exact equilibrium value from the conic program (5).

6.3. Composition of the equilibrium plan

GMDC alters the distribution of both the recipient population and resources used (Figure 4). In Util-SAA, 43% of the budget is used for mother tongue material, 36% for teaching hours, 17% for liaison, and 3.5% for digital access. GMDC re-aligns to use more liaison (33%) and teaching hours (29%), maintains mother tongue materials at 29%, and increases digital access to 8.5%. Mother tongue material is the most efficient resource with nominal efficiency, while nature's equilibrium states that it is also the principal source of danger, and hence GMDC diversifies by adding an alternative channel for liaison and teaching for such populations.

The expenditure per learner is 0.30 for majority language learners, 0.30 and 0.35 for the two regional minorities, 0.92 for migrant learners, and 4.51 for indigenous learners. The floor constraint is binding for the majority language learners and regional minority A. GMDC assigns 63.6% of its budget to rural schools that have 58.0% of the learners due to the concentration of indigenous learners and poor connectivity.

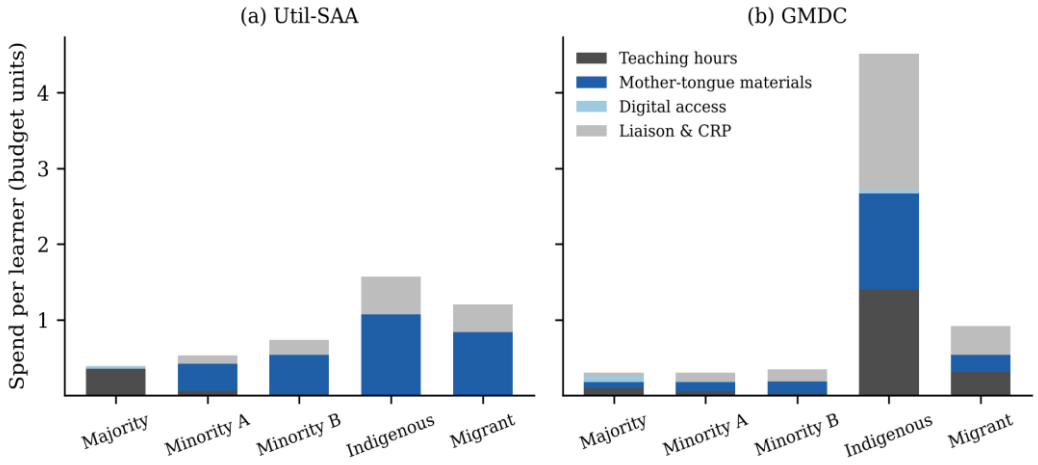


Figure 4: Spend per learner by group and resource type under (a) the stochastic utilitarian plan and (b) the GMDC equilibrium plan.

6.4. Budget sensitivity

Figure 5 depicts the V(B) function for the worst-case value and mean achievement level for different values of the budget ranging from 0.2 to 1.2 units per student. As per Proposition 3, V(B) is an increasing and a concave function and it increases by 0.083 SD between 0.2 to 0.4 units and by 0.021 SD between 1.0 and 1.2 units. GMDC at 0.4 units per learner promises -0.324 SD, which is a better promise than Util-SAA thrice as costly (-0.391 at 1.2 units). The cost of mean achievement for GMDC versus Util-SAA varies between 0.011 and 0.028 SD.

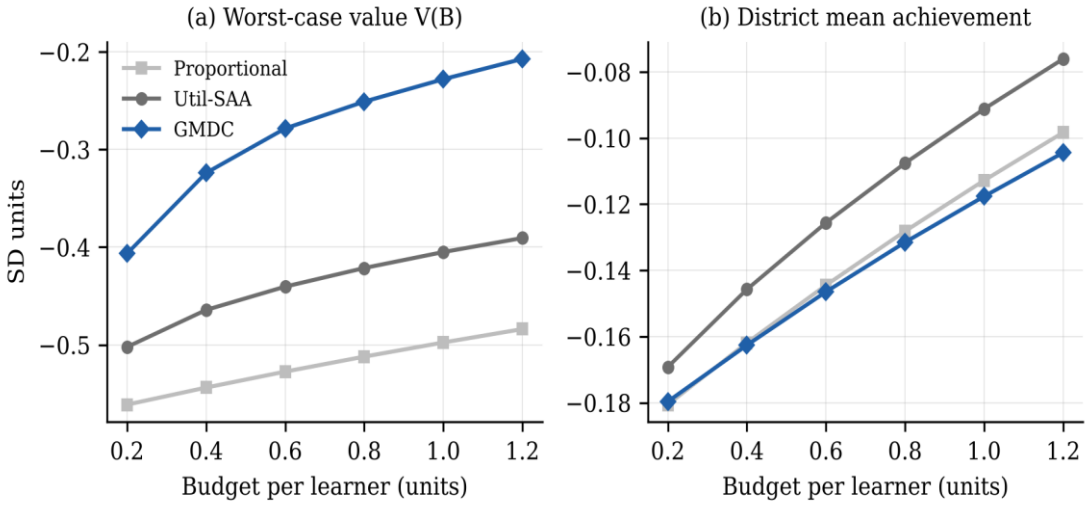


Figure 5: Budget sensitivity of (a) the worst-case value over scenarios and (b) district-mean achievement.

6.5. Sensitivity to the diversity constraints

Table 6 varies the representation floor β and the mother-tongue floor α. The representation floor bears the bulk of the cost of diversity. The increase of β from 0 to 0.5 reduces the worst-case value from -0.246 to -0.278; increasing it to 0.9 reduces it to -0.421, near the utilitarian level, as virtually the entire budget would be allocated according to population proportions. For β=1, the sum of both floors plus the mother tongue reservation exceeds the budget, making it impossible to design any feasible plan. The average achievement by district goes in the opposite direction but far more modestly, from -0.147 to -0.141 SD. The mother tongue floor is nearly costless until α=0.5, as the optimal allocation already exceeds it; for α=1.5, reserving 37.5% of the budget, the worst-case value decreases by 0.039 SD.

Table 6: Sensitivity of the GMDC equilibrium to the diversity constraints (default $\alpha = 0.5$, $\beta = 0.5$; out-of-sample columns use 1,500 draws).

Setting	$V(B)$	OOS district mean	OOS worst-group P5	OOS group gap	Majority share (%)	Indigenous share (%)
$\beta = 0$	−0.246	−0.147	−0.257	0.160	0.0	49.7
$\beta = 0.25$	−0.260	−0.148	−0.270	0.184	14.4	45.6
$\beta = 0.5$ (default)	−0.278	−0.147	−0.288	0.214	28.9	40.9
$\beta = 0.75$	−0.333	−0.143	−0.338	0.265	43.3	29.1
$\beta = 0.9$	−0.421	−0.141	−0.422	0.352	52.0	14.9
$\beta = 1.0$	infeasible	–	–	–	–	–
$\alpha = 0$	−0.278	−0.147	−0.287	0.214	–	–
$\alpha = 1.0$	−0.290	−0.143	−0.297	0.213	–	–
$\alpha = 1.5$	−0.317	−0.141	−0.323	0.225	–	–

6.6. Replication across districts

The ranking is robust with respect to another 20 districts created using the same model (Table 7). GMDC dominated Util-SAA with respect to the 5th-percentile worst-group achievement in all 20 districts, gaining 0.113 SD on average (Wilcoxon p=1.9×10^(-6)) and dominating Rawls-Nominal+div in all 20 districts by 0.026 SD (p=1.9×10^(-6)). Its cost was 0.020 SD lower on average than that of Util-SAA and it closed the gap by 0.117 SD. GMDC was better than need-based allocation with respect to the worst-group achievement by 0.171 SD while there is no difference in mean achievement (p=0.87). The additional costs of the diversity constraints compared with unconstrained minimax were 0.039 SD on the worst-group measure in every district, the replicated price of diversity constraints, while maintaining the representation floor for each group. GMDC's gap was 0.003 SD higher than that of Rawls-Nominal+div on average (p=0.03). So the robustness gain is not achieved at the expense of reducing the gap; it is achieved by increasing the representation floor under uncertainty.

Table 7: Replication over 20 independently generated districts (mean ± standard deviation across districts; 1,000 out-of-sample draws per district).

Rule	Guarantee $\min_{k,c} A_{kc}$	OOS district mean	OOS worst-group P5	OOS group gap
Proportional	−0.517 ± 0.025	−0.143 ± 0.016	−0.520 ± 0.025	0.484 ± 0.035
Need-based	−0.464 ± 0.023	−0.145 ± 0.016	−0.470 ± 0.024	0.404 ± 0.028
Util-Nominal	−0.433 ± 0.028	−0.125 ± 0.016	−0.413 ± 0.029	0.343 ± 0.041
Util-SAA	−0.432 ± 0.029	−0.125 ± 0.016	−0.412 ± 0.029	0.342 ± 0.041
Util-SAA+div	−0.432 ± 0.029	−0.125 ± 0.016	−0.413 ± 0.029	0.343 ± 0.041
Rawls-Nominal+div	−0.342 ± 0.018	−0.139 ± 0.016	−0.325 ± 0.018	0.223 ± 0.023
Minimax (no div)	−0.248 ± 0.019	−0.147 ± 0.016	−0.259 ± 0.019	0.164 ± 0.024
GMDC	−0.291 ± 0.024	−0.145 ± 0.016	−0.299 ± 0.023	0.226 ± 0.026

7. Discussion

7.1. Policy implications

Read nature's strategy as an evidence agenda. The equilibrium condition determines the set of scenario-group pairs that constrain the plan. For the district of the analysis, this is a single variable, the effectiveness of mother tongue material for the three highest distance groups. A ministry deciding where to evaluate would benefit the most by an evaluation of mother tongue material for the indigenous, migrant, and regional minority B learner groups. This converts the strong plan to a claim on the value of information that point estimate planning cannot generate.

Floors protect inclusion; the objective protects outcomes. Representation and mother tongue floors are the tools that policy frameworks usually prescribe. These findings indicate that they are essential, as they exclude the majority group in the unconstrained minimax approach, but insufficient, since the utilitarian approach could meet all the floors and yet leave the indigenous group 0.35 SD behind. Equity implies both a constraint of no excluded groups and a criterion based on the welfare of the worst-off group.

Hedge with a second channel. GMDC doesn't ignore the most fruitful resources. The company retains mother tongue learning materials at around 29 percent of the budget while including community liaisons and focused instruction for those who rely heavily on them. The approach works like diversification of portfolios; if the usefulness of a culturally specific input is uncertain, the strong strategy creates alternative paths to the same audience.

The price of equity is small at the margin. With budgets between 0.2 and 1.2 units per pupil, GMDC reduces district-mean achievement in comparison with the utilitarian solution by 0.011 to 0.028 standard deviations, but increases the worst-case outcome by 0.10 to 0.18 standard deviations. With respect to Kraft's [21] summary of the standardized-effect sizes scale, the former difference is small while the latter is large.

Choose the representation floor deliberately. Value judgements are involved when it comes to the determination of β. Table 6 helps the decision makers to establish the threshold at which protecting the majority share conflicts with the worse-off group's welfare.

7.2. Limitations

It is a computational experiment. The district itself is created from the parametric model rather than being pulled from the administrative record, and the response coefficients are embedded with qualitative regularities from the literature rather than estimated values from any particular program. This means that the results pertain to the properties of the allocation rule under specified uncertainty, not to the size of any effects within any particular system. The application of GMDC to an actual district would need effective estimates, ideally with ranges of uncertainty for each group, derived from local studies.

The payoff function is additive in resources. Complementarities, like those between liaison and teaching, render the payoff non-additive; complementarities are possible while retaining the concavity, for example by having concave functions of the weighted sum of resources used, at the cost of increasing the number of elicited parameters. The model is static and one-period; dynamic allocation and learning about the effectiveness of programs, like bandit style experimentation, would be an obvious extension, with nature's strategy dictating adaptive evaluation. Last, the groups are considered homogenous; intersectional groups like indigenous girls at remote schools could be included as additional agents at the expense of expanding nature's strategy space.

8. Conclusion

In this paper, allocation of educational resources in systems characterized by cultural diversity was modelled as a zero-sum game between the planner and nature, which selects both the cultural-linguistic population to be considered as well as the scenario under which the effectiveness of the plan is assessed. Because the payoff function is a concave function of the allocation and a linear function of the mixed strategy chosen by nature, there is equilibrium, the strategy of the planner consists of a unique published plan, and the plan can be computed via one convex optimization problem with the dual variables of which providing the strategy of nature.

For a 30-school, five-group district, GMDC increased the minimum-worst group performance from -0.440 to -0.278 SD in comparison with the optimal utilitarian plan, reducing the average group difference by 39% at a sacrifice of 0.021 SD in the district mean performance. Neither diversity floors alone, nor an effectively optimized

plan without diversity floors – that excluded the majority group – was able to achieve these results. In addition, the minimax approach pinpointed the only necessary assumption upon which the entire plan is based. Taken together, these features make the game-theoretic minimax allocation feasible for equity-oriented planning with unknown cross-cultural effectiveness of education.

References

1. Arora, S., Hazan, E., & Kale, S. (2012). The multiplicative weights update method: A meta-algorithm and applications. Theory of Computing, 8(1), 121–164.

2. Banerjee, A. V., Cole, S., Duflo, E., & Linden, L. (2007). Remedying education: Evidence from two randomized experiments in India. The Quarterly Journal of Economics, 122(3), 1235–1264.

3. Bertsimas, D., Brown, D. B., & Caramanis, C. (2011). Theory and applications of robust optimization. SIAM Review, 53(3), 464–501.

4. Bertsimas, D., Farias, V. F., & Trichakis, N. (2011). The price of fairness. Operations Research, 59(1), 17–31.

5. Bertsimas, D., & Sim, M. (2004). The price of robustness. Operations Research, 52(1), 35–53.

6. Delage, E., & Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. Operations Research, 58(3), 595–612.

7. Diamond, S., & Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. Journal of Machine Learning Research, 17(83), 1–5.

8. Duan, Q. (2025). Special education and inclusion policies in different cultural contexts. International Academic Journal of Social Sciences, 12(4), 10–17.

9. Hanushek, E. A., & Woessmann, L. (2008). The role of cognitive skills in economic development. Journal of Economic Literature, 46(3), 607–668.

10. Isroilov, S., Zokirkhanovich, K. Z., Adkhamjanovich, A. C., Ugli, S. K. B., & Ugli, K. I. J. (2025). Developing the edu-opt model for strategic management in educational institution. Archives for Technical Sciences, 2(33), 1042–1051.

11. Kraft, M. A. (2020). Interpreting effect sizes of education interventions. Educational Researcher, 49(4), 241–253.

12. Muralidharan, K., Singh, A., & Ganimian, A. J. (2019). Disrupting education? Experimental evidence on technology-aided instruction in India. American Economic Review, 109(4), 1426–1460.

13. Ogryczak, W. (2000). Inequality measures and equitable approaches to location problems. European Journal of Operational Research, 122(2), 374–391.

14. Uchaipichat, A. (2025). Examining social inequality and its effects on policy formation using a multi-level analysis of education systems. International Academic Journal of Humanities, 12(3), 9–15.

15. von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele. Mathematische Annalen, 100(1), 295–320.

16. Wald, A. (1945). Statistical decision functions which minimize the maximum risk. Annals of Mathematics, 46(2), 265–280.