Minimax Theory and its Applications
Volume 11 (2026), No. 2s, 1 to 11

A minimax optimization framework for adaptive teaching strategies in multilingual classrooms

Sherzod Yuvashov[1], Ravqat Masharipov[2], Sodiq Majidov[3], Bexzod Isxakov[4], Sanobarkhon Inomkhujaeva[5], Yanglish Kosimova[6], Xurshid Tursunov[7]

[1]Namangan State Technical University, Namangan, Uzbekistan. yuvashov@gmail.com, https://orcid.org/0009-0005-3982-7990

[2]Lecturer, Department of General Sciences, Mamun University, Khiva, Uzbekistan. ravqatmasharipov89@gmail.com, https://orcid.org/0009-0003-0067-6975

[3]Namangan State University, Namangan, Uzbekistan. sodiqboy2013@gmail.com, https://orcid.org/0009-0002-0100-3329

[4]Namangan State Technical University, Namangan, Uzbekistan. behzodisxakov2782886@gmail.com, https://orcid.org/0000-0003-2706-4605

[5]Lecturer, Namangan State University, Namangan, Uzbekistan. s.inomkhujaeva@nsu.uz, https://orcid.org/0009-0009-8802-5571

[6]Lecturer, Jizzakh State Pedagogical University, Jizzakh, Uzbekistan. kyanglish@gmail.com, https://orcid.org/0009-0003-0953-9731

[7]Independent Researcher and Lecturer, Department of Music Education, Termez State University, Termez, Uzbekistan. tursunovv@tersu.uz, https://orcid.org/0000-0002-1002-1943

Abstract:

Multilingual classes have very diverse levels of language abilities of students in the language of instruction. Teacher is required to make his choice between different types of teaching strategies without knowledge of exact numbers of the students of each linguistic profile in a particular unit, and these numbers vary in time during a semester. Almost all researches in this field focus on analyzing particular strategies or optimizing for the average class profile, and it results in suboptimal learning outcomes for the most poorly-performing students when the class is different from the assumed one. In this paper, we propose minimax framework for selection of teaching strategies. Namely, we treat this problem as a zero-sum game where the teacher chooses her mixed strategy, i.e. a proportion of time allocated to each strategy. It is shown that the problem can be reduced to one linear program via duality, and we incorporate an easy multiplicative-weights scheme that enables the teacher to adjust the strategy week to week based on formative feedback. The proposed framework is evaluated in a simulation setting using five different language student profiles and six teaching approaches. For an uncertainty radius of 0.10, the minimax mix (around 48% of scaffolding approach and 52% of translanguaging approach) increased the minimum gain from 0.516 to 0.536 compared with the average optimal approach, and increased the gain of the weaker learner from 0.22 to 0.36 with a cost of 0.024 in the nominal case. For an unanticipated shift in class composition outside the assumed range, the benefit increased further. The minimax initialised adaptive variant performed the best on average weekly gain for a mid-term shift in class composition. The payoffs are modelling assumptions, so the findings only describe the behaviour of the framework.

Keywords: Minimax optimization, multilingual education, adaptive teaching, robust decision making, translanguaging, linear programming, multiplicative weights.

1. Introduction

Multilingual classrooms where learners use various home languages are becoming more and more widespread in many parts of the world (Amirqulov, 2026). To give an example from India, one and the same school can use either regional language or English or both as the medium of instruction whereas the learners themselves are fluent in different first languages and have different levels of experience with the instructional language. According to the Global Education Monitoring Report by UNESCO, many learners are taught in a language that is not spoken

in their homes and that hinders their learning (Mokhtarinejad et al., 2017). The problem facing a teacher in this case is what strategies are to be employed and how much time is to be devoted to each of them.

There are several methods suggested by research for teachers to pick from. Scaffolding helps students understand the material via visual materials and easy language (Gibbons, 2002). Translanguaging enables students to utilize all their languages as a learning tool. Cooperative learning employs peer support, differentiated tasks fit students' readiness levels (Gibbons, 2002), and digital devices provide students with glossaries and exercises on their own pace. Every method is suitable for certain students and not for others. For example, a lecture in the language of instruction can be very helpful for those who are fluent but totally useless for newcomers. A method that helps beginners can have little effect on fluent students.

There are two reasons that complicate the decision-making process. On the one hand, the efficiency of a certain strategy heavily depends on the composition of the class, which is hardly known in advance by the teacher. On the other hand, the standard procedure of selecting the optimal strategy, which is based on maximization of the gain with respect to the expected class, is not robust. It works perfectly well when the expected class arises, but produces bad results when the number of the students from the expected class is underestimated in the plan while being overestimated in practice. The cost of such a misprediction is especially high in the field of education as it hits the group of students, which needs additional help.

In the current paper, the robust approach is considered. Rather than relying on one estimate of the class composition, we assume that it belongs to the specified neighbourhood of the nominal estimate. After that, we look for the teaching plan which guarantees the maximum worst-case learning gain. This is the minimax approach in decision making, (Wald, 1950). The contributions of the paper are as follows.

- We cast the choice of the pedagogical strategy in a multilingual classroom as a zero-sum game between the teacher and the adversarial classroom composition, with the payoff being the sum of learning benefit and teacher effort.

- We prove that the game with the box-type uncertainty set can be solved exactly with one linear program through duality, making the procedure affordable enough for a teacher or school to implement it using a spreadsheet solver.

- We incorporate the adaptive component based on the multiplicative weights (Hedge) algorithm, which allows for updating the strategy based on weekly formative feedback, and show how to initialize it from the minimax solution.

- We analyse the proposed framework experimentally in the simulation study with five language configurations and six strategies, considering the worst-case performance, equity, sensitivity to the uncertainty radius, convergence and composition shift during the mid-term period.

The rest of the paper is organized as follows. Section 2 gives the overview of the related work. Section 3 defines the problem and Section 4 introduces our framework. Section 5 describes the experimental setup and Section 6 provides the results. Section 7 discusses the results, Section 8 gives the limitations and Section 9 concludes.

2. Related work

2.1 Teaching strategies for multilingual learners

Cummins' hypothesis of interdependence claims that skills acquired in one language facilitate acquisition of another one (Cummins, 2014), which is a theoretical basis for use of home languages in class (Irmuxamedov et al., 2026; Zaidi, 2025). The idea of translanguaging pedagogy takes this hypothesis further by treating all language abilities possessed by the learner as a source of information (Cenoz & Gorter, 2017). Large scale studies of dual language programs reveal very positive outcomes for language minority students over a long period of time (Collier & Thomas, 2004). Scaffolding, which is based on the theory of zone of proximal development of Vygotsky provides only temporary assistance, which disappears as the learner's independence grows (Gibbons, 2002). Cooperative learning is proven effective in mixed ability groups through many years of research, and differentiated instruction adjusts content and tasks to learners' readiness level (Tomlinson, 2001). Reviews of effectiveness of various teaching strategies indicate that the magnitude of their impact varies greatly from context to context, from learner to learner (Hattie, 2009). Cognitive load theory raises an issue that language complexity increases cognitive load and can block learning process in beginners.

2.2 Optimization and adaptivity in education

Adaptive instructional computing approaches operate mostly at the learner level (Maraimova et al., 2026). Knowledge tracing predicts the current knowledge level of a learner and selects the next task, while multi-armed

bandit algorithms choose between tasks and learn from experience (Clement et al., 2015). Such an approach is highly valuable for digital learning environments, yet it relies on many iterations with one particular learner and much data. The problem of a classroom teacher is quite another; he/she needs to make one decision about one set of strategies to use for an entire group of learners.

2.3 Minimax and robust decision making

Minimax theorem of von Neumann guarantees the existence of the value of the finite zero-sum game and the solution of which can be obtained through linear programming. Wald suggested to use minimax principle in making decisions in the face of uncertainties (Wald, 1950). The minimax regret criterion was studied by Savage. Robust optimization offers a feasible technique for defending a decision from the effects of uncertainties that lie in a known set (Ben-Tal et al., 2009). Bertsimas and Sim proved that this defense could be purchased at low costs in terms of nominal performance (Bertsimas & Sim, 2004). The multiplicative weight update algorithm provides a simple on-line technique for approaching the minimax solution in the repetitive case of the game (Arora et al., 2012). We apply these techniques to the problem of teaching strategy selection, an area to our best knowledge, that has not yet benefited from their combination.

3. Problem formulation

3.1 Setting

Suppose there is a class partitioned into n groups based on their language profiles and a teacher capable of using m different teaching strategies. Group g is characterised by the competence in the language of instruction and literacy in the home language. Let Uag ∈ [0, 1] be the anticipated normalised learning gain of a pupil in group g given the use of teaching strategy a across the entire unit. Let ca ≥ 0 be the additional effort imposed by teaching strategy a on the teacher measured on the same scale as learning gain, such as preparation time which would otherwise have been devoted to other tasks. The payoff from the use of strategy $a$on group $g$is given by:

$$A_{ag} = U_{ag} - c_a \qquad (1)$$

In this case, the teacher employs a mixed strategy $x = (x_1, \dots, x_m)$in the probability simplex Δ^m. It should be noted that $x_a$represents the proportion of instructional time devoted to strategy a.

This text is perfect for teaching purposes since the teacher can combine methods in different classes, while it cannot be done for an isolated case.

3.2 Uncertainty about the class

Suppose that $p = (p_1, \dots, p_n)$ are the true shares of students in every group. The teacher possesses some nominal estimation p^0, which could be obtained from a language survey at the beginning of the term; however, he or she cannot trust it completely. It is assumed that the true shares belong to the following set:

$$P_\delta = \{p \in \Delta^n : |p_g - p_g^0| \le \delta \text{ for all groups } g\} \quad (2)$$

with $\delta \ge 0$ being the uncertainty radius; for example, if $\delta = 0.10$, then the share of every group could deviate from the nominal estimation by up to ten percent. This set is intuitively understandable for teachers and school administrators, and it includes the nominal case ($\delta = 0$).

Define the following notations: let $l_g = \max(p_g^0 - \delta, 0)$ and $u_g = \min(p_g^0 + \delta, 1)$ be the lower and upper bounds.

3.3 The minimax objective

The expected gain for a given plan x and composition p is denoted by $x^T Ap$. The teacher seeks a strategy that works well even under the worst permissible composition. This leads to the following problem:

$$V_\delta = \max_{x \in \Delta^m} \min_{p \in P_\delta} x^T Ap \quad (3)$$

The minimization within represents an adversary (nature) that shifts the composition toward those groups that the plan treats poorly. The maximization outside is the response from the teacher. The value V_δ is the learning gain that can be guaranteed by the teacher against all compositions in P_δ. For contrast, the classical average-case criterion seeks the plan x that maximizes $x^T Ap^0$, referred to as the Bayes-nominal plan.

4. Proposed framework

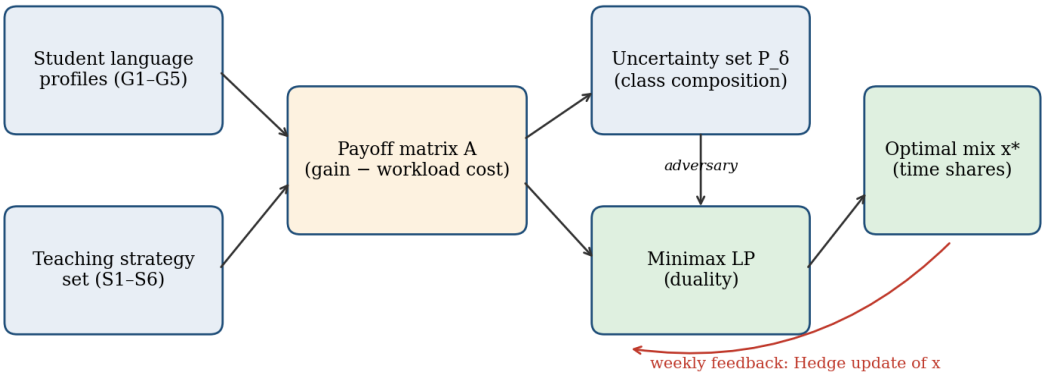


Figure 1: Overview of the proposed framework. Language profiles and the strategy set define the payoff matrix. The minimax model protects against the worst class composition in the uncertainty set, and weekly feedback updates the plan.

4.1 Exact solution by linear programming

For a fixed x, the inner problem in Equation (3) is a linear program in p with the constraints $\Sigma_g\, p_g = 1$ and $l_g \le p_g \le u_g$. The dual has a free variable λ representing the sum constraint and variables αg and βg which are non-negative, representing the lower and upper bounds. Using strong duality, the inner minimisation problem is equal to the maximum of the dual problem. Since the dual is a maximisation problem, the two problems may be combined into one linear program:

$$\max \quad \lambda + \sum_g l_g\, \alpha_g - \sum_g u_g\, \beta_g \quad (4)$$

Such that: $\lambda + \alpha_g - \beta_g \le \sum_a x_a\, A_{ag}$ for all groups $g\lambda + \alpha_g - \beta_g \le \sum_a x_a\, A_{ag}$ for all groups $g\, \sum_a x_a = 1\, \sum_a x_a = 1 x_a, \alpha_g, \beta_g \ge 0 x_a, \alpha_g, \beta_g \ge 0$ The LP formulation requires $m + 2n + 1$variables and $n + 1$constraints, hence the model is very small and can be solved in milliseconds. Our implementation used the HiGHS LP solver through SciPy.

4.2 Properties

Proposition 1 (guarantee). Let x be a solution to the linear programming problem, and let x^B be the Bayes-nominal plan. Then: $\min_{p \in P_\delta} x^T Ap \ge \min_{p \in P_\delta} (x^B)^T Ap$ and the latter inequality is true for any arbitrary fixed plan. This is implied directly from the definition of the maximizing x.

Proposition 2 (monotonicity). Vδ is a nonincreasing function of δ. The sets Pδ are nested, and hence a bigger set provides the adversary with more choices and hence cannot increase the lower bound. The case δ = 0 corresponds to the Bayes-nominal problem, and hence V0 is the optimal expected gain that can be achieved, and V0 – Vδ is the cost of robustness.

Remark (support of the solution). The simple optimal solution of linear programming, having n+1 constraints, has only a few non-zero values. It is therefore obvious that the minimax strategy involves only a few strategies. This makes it easier for teachers to implement because they cannot design six different teaching styles simultaneously. This can be verified by experimentation.

4.3 Adaptive layer

The static plan employs solely the nominal estimation along with the radius δ. Throughout the semester the teacher is informed of progress via quizzes, exit tickets, and observations. We model this as a weekly vector of estimated gains $r^t \in \mathbb{R}^m$, one per strategy, obtained from short formative assessments. The adaptive layer keeps a weight w_a for each strategy, uses the plan $x_a{}^t = w_a{}^t / \Sigma_b\, w_b{}^t$, and updates the weights with the Hedge rule:

$$w_a^{t+1} = w_a^t \exp(\eta r_a^t) \quad (5)$$

The Hedge algorithm has a regret that grows at $\sqrt{(\ln m)/T}$ per step (Arora et al., 2012), and hence the strategy follows the best fixed mixed strategy. In case this algorithm is used with the best-responding adversary in the

repeated game, the mean strategy converges to the minimax solution (Arora et al., 2012).

There are two modifications. The Hedge algorithm starts off with equally weighted strategies. The Minimax-Initialised Hedge (MH) algorithm starts off with weights proportional to the static solution x^*, thus starting off with a safe strategy and then moving on to better strategies for the particular class. A small constant floor of 10^{-4} for each weight ensures availability of all strategies, which will help recover from the change in the class.

4.4 Procedure

1. Define the groups and strategies, and determine the payoff matrix A using existing studies, trial sessions or the teacher's knowledge.

2. Determine the nominal distribution p0 using a linguistic survey, and select the radius δ based on the variation in the enrollment and attendance.

3. Solve the linear programming problem (4) to determine x and the guaranteed value V*δ.

4. Transform x* into a lesson schedule by allocating the lessons according to the proportions.

5. On a weekly basis, solicit formative feedback and use it to update the weights according to equation (5). Adjust the schedule accordingly, and solve the linear programming problem again if there are any changes to the survey.

5. Experimental design

These are simulations. The goal of this research is to evaluate the performance of the framework, its protective abilities towards the most disadvantaged groups and its cost, under an explicitly described model of multilingual classes. All algorithms are implemented in Python using NumPy and SciPy libraries and random seeds are fixed for reproducibility.

5.1 Groups and strategies

Five groups of students have been created, as outlined in Table 1. The proportionate breakdown shows a class where the majority of students speak the language of instruction while the minority still learn it.

Table 1: Student language-profile groups and nominal shares.

Group	Profile	Nominal share
G1	Fluent in the language of instruction (LoI)	0.35
G2	Balanced bilingual, fluent in LoI and home language	0.25
G3	Intermediate LoI proficiency, literate in home language	0.18
G4	Beginner in LoI, literate in home language	0.14
G5	Beginner in LoI, limited literacy in both languages	0.08

The following six teaching methods were investigated: S1: lecture in the instructional language only; S2: scaffolded instruction; S3: translanguaging pedagogy; S4: collaboration among peers in mixed-language settings; S5: differentiated learning assignments; and S6: technology-enhanced instruction using glossary and translation tools.

5.2 Payoff model

In Table 2, the gain matrix U and the workload cost c are presented. These figures have been chosen by us based on the qualitative analysis of the literature. Lecture is the most effective method for fluent learners and least effective for beginners (Cummins, 2000). Translanguaging offers its biggest benefits to learners literate in their native languages (Collier & Thomas, 2004; Cummins, 2014), while offering fewer benefits to learners who use only the language of instruction. Scaffolding is beneficial to intermediate learners but does not benefit the least literate ones (Gibbons, 2002). Peer collaboration works relatively well for all types of learners with a narrow difference between them. Tiered tasks are most beneficial for all types of learners except those who lack literacy (Tomlinson, 2001). Technology-assisted practice is uniformly beneficial. Workload costs equal 0 for lecture and range between 0.04 and 0.08 for other methods.

Table 2: Gain matrix Uag (normalised learning gain) and workload cost ca. The net payoff is Aag = Uag − ca.

Strategy	G1	G2	G3	G4	G5	Cost

S1 Monolingual lecture	0.84	0.76	0.40	0.18	0.08	0.00
S2 Scaffolded instruction	0.74	0.74	0.60	0.44	0.26	0.04
S3 Translanguaging	0.48	0.70	0.68	0.76	0.56	0.07
S4 Peer collaboration	0.58	0.62	0.60	0.52	0.50	0.05
S5 Tiered tasks	0.68	0.64	0.66	0.58	0.34	0.08
S6 Technology-assisted	0.56	0.60	0.52	0.58	0.60	0.06

5.3 Compared methods

Comparison was made of nine plans: the six individual plans, equal mixture of all six, Bayes-nominal plan (best plan for the nominal composition), and the proposed minimax plan. The primary value of $\delta = 0.10$ was used, and δ varied between 0 and 0.20 in the sensitivity analysis.

5.4 Evaluation

There are four ways of evaluating. (i) Precise worst case: for each policy, the worst-case advantage of the policy over $P\delta$ is obtained by solving the inner linear program, where $\delta = 0.10$ and $\delta = 0.20$. (ii) Monte Carlo: 5,000 compositions were sampled inside P0.10 using rejection sampling from Dirichlet distribution, and another 5,000 compositions were sampled from a wider Dirichlet distribution (concentration $4 \times p0$) to make sure many fall outside the assumed set. For each case, the gain of each group was corrupted with Gaussian noise (standard deviation 0.05) for student-level variance, and we report the mean and the 5th percentile of the class gain. (iii) Equity: the expected gain of the most disadvantaged group (the minimum over groups of $(ATx)g$) and the gap between the best and the worst groups. (iv) Adaptive term simulation: 40 weeks with composition equals p0 and a shock week occurs randomly between week 10 and week 29, with a new arrival composition (0.12, 0.16, 0.20, 0.26, 0.26). This composition lies outside P0.10. In each week, the teacher had a noisy signal of the gain of every strategy (standard deviation 0.05), and Hedge used step size $\eta = 0.8$.

6. Results

6.1 Static plans at the main uncertainty radius

For $\delta = 0.10$, the value obtained from the linear program was $V\delta = 0.536$. In the optimal strategy mix, there were two strategies, with the probability of 48.3% using scaffolded instruction (S2) and 51.7% using translanguaging (S3). The worst-case distribution in this case was (0.35, 0.15, 0.08, 0.24, 0.18), where students in the two groups closest to the midpoint (G2 and G3) had the minimum weight possible while students in the two beginner groups (G4 and G5) had the maximum weights. On the other hand, Bayes-nominal strategy is a pure S2 plan and has the maximum nominal gain (0.594).

The principal comparison is presented in Table 3. Bayes-nominal strategy produced the highest nominal gain (0.594). However, minimax strategy outperformed the former on all the worst-case and tail criteria. With $\delta = 0.10$, the worst-case gain increased from 0.516 to 0.536 (0.020 or 3.9%), and with $\delta = 0.20$, it increased from 0.438 to 0.502 (0.064 or 14.6%). The trade-off was a reduction of 0.024 in the nominal gain (4.0%). The 5th percentile of the transformed composition increased from 0.457 to 0.490. Of the individual strategies, the monolingual lecture was the second-best in terms of nominal gain and the least robust one, reducing from 0.454 at $\delta = 0.10$ and 0.320 at $\delta = 0.20$.

Table 3: Performance of all plans. Worst-case values are exact over the uncertainty set. P5 is the 5th percentile of class gain over 5,000 shifted compositions with noise. The last column is the expected gain of the weakest group.

Plan	Nominal	Worst-case δ=0.10	Worst-case δ=0.20	P5 (shifted)	Weakest group
S1 Monolingual lecture	0.588	0.454	0.320	0.354	0.080
S2 Scaffolded instruction	0.594	0.516	0.438	0.458	0.220
S3 Translanguaging	0.547	0.505	0.467	0.447	0.410
S4 Peer collaboration	0.529	0.509	0.489	0.471	0.450
S5 Tiered tasks	0.545	0.503	0.462	0.454	0.260
S6 Technology-assisted	0.509	0.497	0.487	0.451	0.460
Uniform mix	0.552	0.510	0.467	0.470	0.340
Bayes-nominal	0.594	0.516	0.438	0.457	0.220
Minimax (proposed)	**0.570**	**0.536**	**0.502**	**0.490**	**0.360**

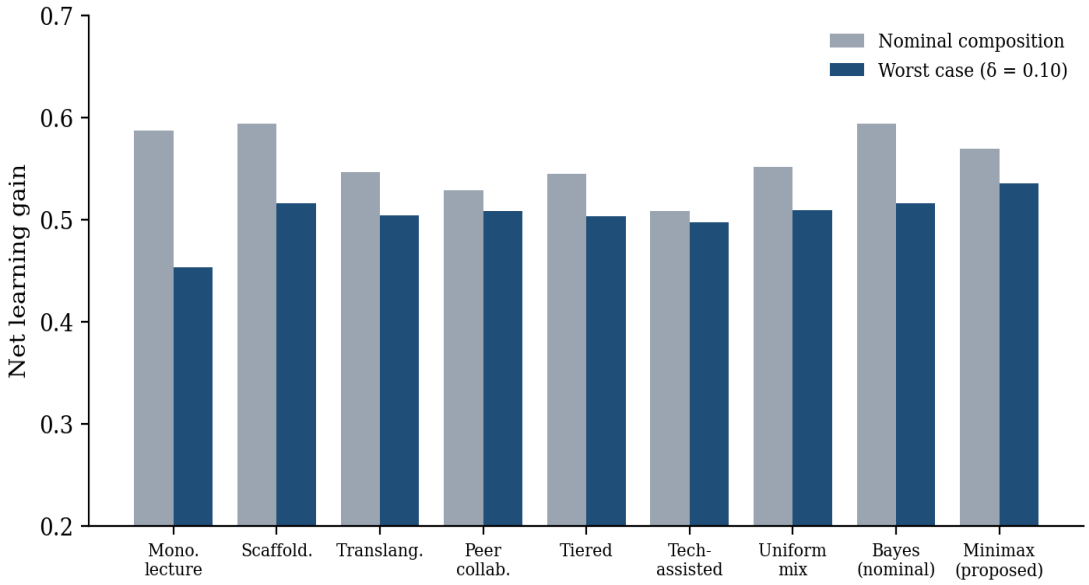


Figure 2: Net learning gain under the nominal composition and in the worst case for $\delta = 0.10$.

6.2 Equity

The expected gain for the least favorable group (G5) in the minimax approach is 0.36 compared to 0.22 in the Bayes-nominal approach, which is an increase of 64%. There is a decrease in the difference between the best and worst groups from 0.48 to 0.30. The expected gains for each group using the minimax approach are 0.55, 0.66, 0.59, 0.55, and 0.36 for groups G1 to G5. Peer learning and technology-based practice had lesser differences (0.12 and 0.08) but had overall less gains. Therefore, they do not offer the best trade-off in the problem. There is no equity clause in the minimax approach; however, improvement has been achieved as it protects any shift towards the beginners' groups. There can be an equity clause included in the linear program in case a school has a minimum gain for each group.

6.3 Sensitivity to the uncertainty radius

The changes in the optimal mix due to the value of δ are given in Figure 3. For the cases where $\delta = 0$ or $\delta \leq 0.05$, the optimal mix consisted purely of S2. From $\delta = 0.075$ onwards, the plan shifted slightly to S3, which means 58.8% S2 and 41.2% S3. From $\delta = 0.10$ to $\delta = 0.15$, the plan remained at 48.3% / 51.7% mix; and from $\delta = 0.175$ to $\delta = 0.20$, technology-assisted practice (S6) started appearing, with the percentage being 40.9%, while S2 and S3 were around 29% to 30%.

The plan is robust to wide ranges of δ values, and this is good since the δ can only be approximated practically.

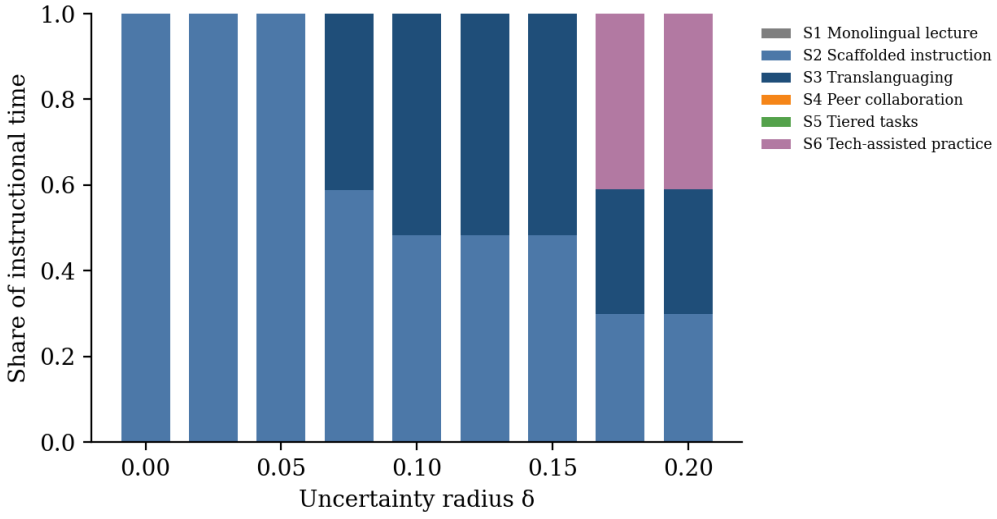


Figure 3: Optimal minimax time shares as the uncertainty radius δ increases.

Figure 4 shows the worst-case values. The guaranteed value decreased steadily from 0.594 when $\delta = 0$ to 0.509

when δ = 0.20, as expected under Proposition 2. The Bayes-nominal strategy coincided with the minimax strategy until δ = 0.05, but thereafter trailed behind and finally had a worst-case value of 0.438 at δ = 0.20, 0.070 less than that of the minimax. Monolingual lecturing was the fastest decline, ending at 0.320. Translanguaging alone was slower but remained worse than the minimax for all δ > 0.

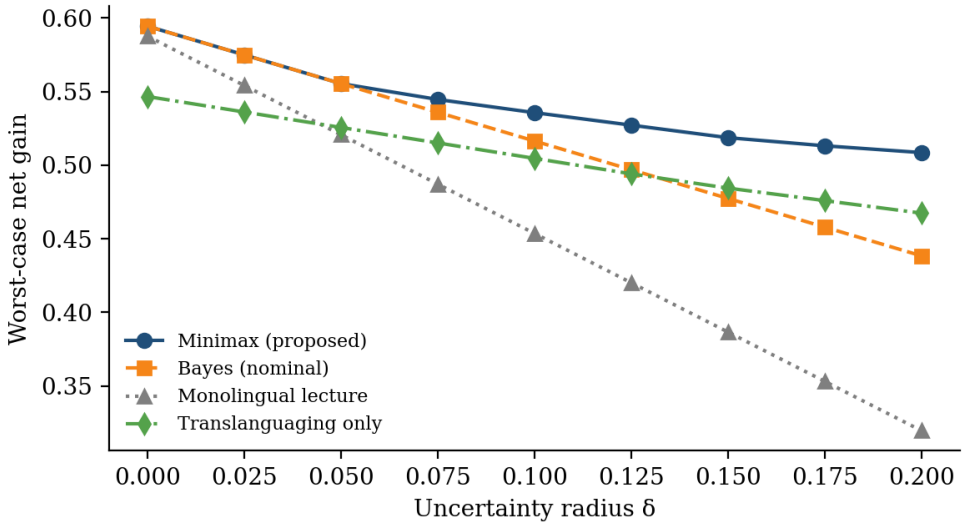


Figure 4: Worst-case net gain as a function of δ for four plans.

6.4 Convergence of the repeated game

We have tested whether the Hedge strategy, playing against a best-responding adversary, finds the linear programming solution. After 2,000 iterations, the expected plan achieved the worst-case value of 0.534, while the solution has the value of 0.536, which makes the difference 0.3%. It was 1.1% after 500 iterations and 2.6% after 100 iterations (see Figure 5). This convergence justifies the application of Hedge updates as an online approximation, when the teacher cannot, or does not want to, recompute the linear program.

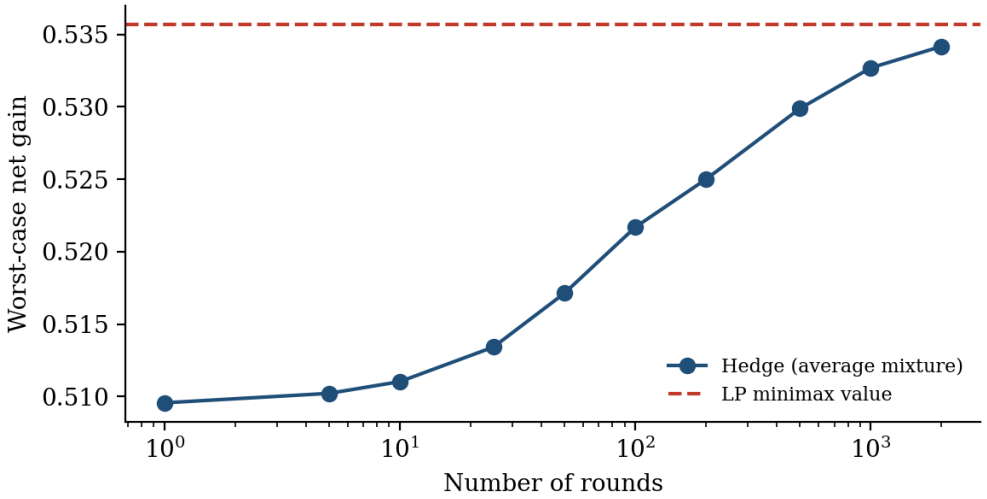


Figure 5: Worst-case value of the average hedge plan against the exact linear programming value (δ = 0.10).

6.5 Adaptive teaching under a mid-term shift

Table 4 presents the results of the 40-week simulation run with composition changes not falling into the set of those which were supposed by the design of the experiment. The fixed Bayes-nominal rule performed best at the start (average gain of 0.595 in the first ten weeks), but its gain declined to 0.470 in the last ten weeks after the change. The fixed minimax strategy had a slightly lower start (average gain of 0.570) but kept 0.526 gain after the change. Hedge with the starting weights of one-half had a slow convergence and a lower average gain (0.530). The variant of the minimax rule using hedge (MH) showed the highest average weekly gain (0.552) and the highest gain in the last ten weeks (0.538), and the first ten weeks was close to the fixed minimax strategy (0.571). But the fixed minimax strategy demonstrated the best gain in the worst week (0.492 vs. 0.456 of MH) and the 5th percentile week (0.513 vs. 0.503 of MH).

Table 4: Weekly net gain over 300 simulated 40-week terms with a mid-term composition shift.

Policy	Mean	First 10 weeks	Last 10 weeks	P5 week	Worst week
Uniform mix (fixed)	0.518	0.552	0.485	0.472	0.446
Bayes-nominal (fixed)	0.531	0.594	0.470	0.456	0.434
Minimax (fixed)	0.547	0.570	0.526	0.513	0.492
Hedge (equal start)	0.530	0.555	0.511	0.467	0.432
Minimax-initialised Hedge	**0.552**	**0.571**	**0.538**	**0.503**	**0.456**

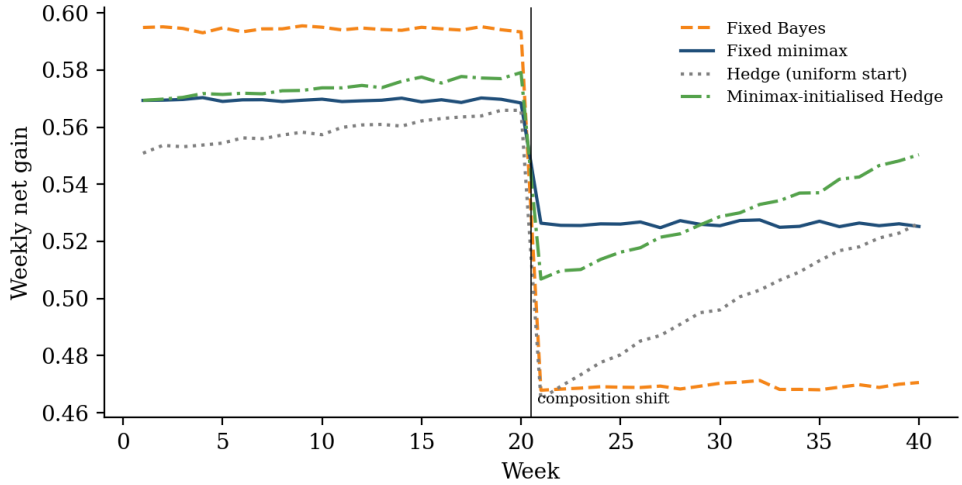


Figure 6: Mean weekly net gain with a composition shift at week 20 (200 runs). The fixed Bayes-nominal plan falls after the shift, while the minimax-based plans hold up.

7. Discussion

The results demonstrate a consistent pattern. If a plan is selected based on the anticipated class, it produces the best outcome only if the composition is as anticipated. Otherwise, any deviation brings loss mostly to the learners who are supported the least, which is precisely the subgroup of interest to the school. In contrast, the minimax plan involves an acceptable sacrifice in the nominal case (0.024 in our problem) and pays back in improvement of the worst case and gain of the least favored subgroup. This is indeed a classical tradeoff of robust optimization called the price of robustness (Bertsimas & Sim, 2004); and as can be seen from the sensitivity analysis results, the price is relatively low compared to the protection gained.

The finding that the minimax plan used mixed scaffolded instruction and translanguaging is also confirmed by the literature; the complementary nature of these two pedagogies has been stated many times in previous work (Gibbons, 2002). Of course, the framework does not necessarily produce this optimal combination. The mixing depends on the particular structure of the payoff matrix and on δ, as is shown by the presence of the technology-assisted practice at large δ.

The narrow support of the solution (only two strategies at the main radius) provides for practical implementation. For example, a teacher can choose to implement "roughly half of the lessons in scaffolded mode and half in translanguaging mode" without significant additional preparation. The linear program can be implemented using a simple spreadsheet, and therefore, the framework does not require any special software.

The adaptive experiment offers another practical implication. Starting from a robust plan and then using the feedback (MH) provided a safer start and a better long-run average than either the fixed plan or learning from scratch. With noisy feedback and small classes, the robust plan can be a solid foundation. Thus, a school may start the term with a static plan and switch to an adaptive one after several weeks.

In planning, the difference between the guaranteed value Vδ and the nominal value V0 shows how costly lack of information about the class is. If this difference is substantial, it makes sense to invest effort in conducting a better survey since δ will be decreased. Otherwise, no further measurement is required, and the robust plan works just fine.

8. Limitations and future work

However, the main drawback is that the payoff matrix is a modelling assumption. Values are in line with the results

in the literature, but are not calculated based on classroom observations and actual results depend on the subject matter, grade level, competence of the teacher and languages used. The numbers should be seen as an illustration of the method operation and not as proof of the real scale of gains. The conclusions, such as that the average case strategy is fragile and the minimax strategy is a tradeoff between minor loss and security, are the result of the modelling assumptions. However, the exact numbers do not translate into the real life school situation.

Further limitations are the following. First, the model considers the unit as the time share of independent strategies and neglects order effects and the carry-over effect of learning. Second, the number of groups is small, and the division into them may not be valid for individual students. Third, the adaptive experiment assumes the teacher knows the estimate of gains from every strategy every week, which means she knows more than the usual amount of information. Thus, a version of the problem with partial feedback (bandit) seems more realistic. Fourth, the box-type uncertainty set separates every group. Other uncertainty sets can be applied, e.g., the distance from the nominal distribution. Finally, minimax is by definition conservative, and another criterion, minimax regret or mixed one, might be considered by a teacher with less aversion to risks.

Future research should (i) use real data on payoff matrices obtained from classroom experiments or school tests, (ii) conduct an experimental verification of the model on real classroom data, (iii) introduce the equity constraints and workload restrictions to the linear programming problem, (iv) consider partial-feedback models for the adaptive layer, and (v) expand groups for specific subject language requirements.

9. Conclusion

In this paper, we developed a minimax optimization model for selecting teaching strategies in multilingual classrooms in case the composition of the class is not known. The model was formulated as a two-person zero-sum game between the teacher and the worst possible composition of the class in a bounded space, it could be solved analytically by one linear programming problem and was supplemented by an adaptive multiplicative-weights layer for weekly tuning. In our numerical experiment that considered five language profiles and six teaching strategies, the minimax solution, which consisted of scaffolded instruction and translanguaging, outperformed the average-case solution in the worst-case gain, gain of the weakest class and its tail performance at a small cost. Furthermore, the minimax solution was robust in a large interval of uncertainty radius and with the initial solution set as minimax solution, the adaptive layer performed best in case of a mid-term change of the class composition. The proposed framework is easy to use and interpret and can be computed cheaply. The next goal is to calibrate the gain function based on real-world data.

References

1. Amirqulov, I. (2026). Facilitating Translingual Practices in Multilingual STEM Classrooms to Bridge Content and Language Gaps. *International Journal of English and Education, 15*(2), 1-9. https://doi.org/10.67050/IJEE/V15I2/IJEE262001

2. Arora, S., Hazan, E., & Kale, S. (2012). The multiplicative weights update method: a meta-algorithm and applications. *Theory of computing*, *8*(1), 121-164. http://dx.doi.org/10.4086/toc.2012.v008a006

3. Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). Robust Optimization. Princeton University Press, Princeton, NJ.

4. Bertsimas, D., & Sim, M. (2004). The price of robustness. Operations Research, 52(1), 35–53. https://doi.org/10.1287/opre.1030.0065

5. Cenoz, J., & Gorter, D. (2017). Minority languages and sustainable translanguaging: Threat or opportunity?. *Journal of Multilingual and Multicultural development*, *38*(10), 901-912. https://doi.org/10.1080/01434632.2017.1284855

6. Clement, B., Roy, D., Oudeyer, P.-Y., & Lopes, M. (2015). Multi-armed bandits for intelligent tutoring systems. *Journal of Educational Data Mining, 7*(2), 20–48.

7. Collier, V. P., & Thomas, W. P. (2004). The astounding effectiveness of dual language education for all. *NABE Journal of Research and practice*, *2*(1), 1-20.

8. Cummins, J. (2000). Language, Power and Pedagogy: Bilingual Children in the Crossfire. Multilingual Matters, Clevedon (Vol. 23).

9. Cummins, J. (2014). Linguistic interdependence and the educational development of bilingual children. In *The new immigrant and language* (pp. 72-101). Routledge.

10. Gibbons, P. (2002). Scaffolding Language, Scaffolding Learning: Teaching Second Language Learners in the Mainstream Classroom. Heinemann, Portsmouth.

11. Hattie, J. (2009). Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement. Routledge, London.

12. Irmuxamedov, T., Maxmudov, J., Ibragimova, G., Muminova, A., Muradova, A., Safarova, F., & Sapaev, B. (2026). Developing Cross-Cultural Communicative Competence in Advanced Language Learners Through Generative AI-Based Socio-Pragmatic Support. *Archives for Technical Sciences, 1*(35), 751–762. https://doi.org/10.70102/afts.2026.1835.751

13. Maraimova, M., Pardayeva, K., Satimova, D., Turakulova, Z., Kurambayeva, G., Navruzova, M., & umayev, F. (2026). An Intelligent Adaptive Learning Framework for Uzbek Language Teaching Using Dynamic Parameter Adaptation. *International Academic Journal of Science and Engineering, 13*(2), 447-456. https://doi.org/10.71086/IAJSE/V13I2/IAJSE1386

14. Mokhtarinejad, A., Mokhtarinejad, O., Kafaki, H. B., & Ebrahimi, S. M. H. S. (2017). Investigating German Language Education through Game (Computer and non-Computer) and its Correspondence with Educational Conditions in Iran. *International Academic Journal of Innovative Research, 4*(2), 1–9.

15. Tomlinson, C. A. (2001). How to Differentiate Instruction in Mixed-Ability Classrooms (2nd ed.). ASCD, Alexandria, VA.

16. Wald, A. (1950). Statistical decision functions. In *Breakthroughs in Statistics: Foundations and Basic Theory* (pp. 342-357). New York, NY: Springer New York. https://doi.org/10.1007/978-1-4612-0919-5_22

17. Zaidi, Z. F. (2025). Developing a Predictive Model for Language Acquisition in Multilingual Environments Based on Cognitive and Social Factors. *International Academic Journal of Humanities, 12*(3), 1-8. https://doi.org/10.71086/IAJH/V12I3/IAJH1208

18. Dönük, D., & Demir, Ö. K. (2026). Developing Productive Language Skills Using Artificial Intelligence Tools. *International Journal of Education Technology and Scientific Researches*, *11*(33), 58-64.

19. Poorani.D. (2026). A Systematic Review of Applied Nonlinearity in Modeling Information Service Dynamics for Academic Library and Research Data System Networks. *Applied Nonlinearity in Science and Technology*, *2*(1), 42-50.