RESEARCH ARTICLE

# AI-Assisted Learning and Dropout Risk in Romanian Secondary Education: A Logistic Regression Approach to Predictive Risk Assessment

**Adrian NICOLAU**[1]

[1] Athenaeum University of Bucharest, Romania

ORCID: 0009-0009-8954-4429 · nicolau.adrian@univath.ro

**Received: 12 Mar 2026**
**Revised: 17 Mar 2026**
**Accepted: 23 Mar 2026**
**Published: 30 Mar 2026**

**DOI:** https://doi.org/10.5281/zenodo.19352827

## Abstract

*Student dropout remains a persistent challenge in Romanian secondary education, yet most interventions operate reactively rather than through systematic risk assessment. This study investigates whether exposure to AI-assisted learning tools is associated with lower academic vulnerability among secondary school students, and whether a compact logistic regression model can function as a predictive risk scoring instrument applicable in educational management contexts. Using survey data from 1,702 respondents (680 students, 873 parents, 149 teachers), the study constructs a binary academic vulnerability variable from self-reported learning difficulties across core subjects, used as a proxy for dropout risk, and estimates a regularised logistic regression with three predictors: AI tool usage, gender, and grade level. Results indicate that respondents engaged with AI-assisted learning are approximately 95% less likely to report academic vulnerability (OR = 0.051, 95% CI: 0.002, 0.652), though this estimate is sensitive to limited variance in AI usage across the sample. Female respondents exhibit moderately higher stated risk (OR = 3.576, $p < 0.001$), while higher grade levels are associated with marginally increased vulnerability (OR = 1.224, $p < 0.01$). The model achieves high discriminatory power (AUC = 0.992) and classification accuracy exceeding 95% at the conventional threshold, though the near-universal AI usage rate (99.4%) likely inflates the AUC. These findings reposition AI not merely as a pedagogical tool but as a component of institutional risk management, offering educational leaders a framework for early identification and targeted support of at-risk students.*

**Keywords:** educational risk assessment; dropout prediction; artificial intelligence; logistic regression; educational management; secondary education

**JEL Classification:** I21; I28; C25; M10

## 1. INTRODUCTION

Student dropout in secondary education represents one of the most persistent and costly failures of educational systems worldwide. In Romania, the early school leaving rate remained above 15% in 2023, well above the EU average of 9.6% and far from the European Education Area target of below 9% by 2030 (Eurostat, 2024). Each student who leaves school prematurely generates long-term costs: reduced lifetime earnings, higher unemployment risk, lower civic participation, and intergenerational transmission of disadvantage (Sileshi et al., 2024). From a management perspective, dropout is not merely an educational outcome; it is an institutional risk that educational leaders should be equipped to assess, predict, and mitigate systematically.

The integration of artificial intelligence into educational practice has accelerated substantially since 2020. AI-based tools now range from conversational tutors and automated formative assessment to adaptive learning platforms that personalise content difficulty in real time. Meta-analytic evidence suggests that these tools produce, on average, positive effects on student performance and motivation (Zhao et al., 2024; Khosravi et al., 2025). However, most research evaluates AI through the lens of learning gains or instructional efficiency. The question of whether AI engagement is associated with stronger institutional attachment, specifically with reduced intention to leave school, has received considerably less attention, particularly in Central and Eastern European educational systems that face distinct structural constraints.

This study addresses three objectives. First, it documents the prevalence and distribution of self-reported academic vulnerability across students, parents, and teachers in a Romanian secondary education sample. Second, it estimates a regularised binary logistic regression to test whether AI-assisted learning is associated with lower stated dropout risk, controlling for gender and grade level. Third, it evaluates the predictive performance of this compact model as a potential risk scoring tool applicable in educational management contexts, using standard diagnostic criteria from the predictive analytics literature (AUC, calibration, confusion matrix metrics).

The analytical framing is deliberate. Logistic regression is the standard methodology for binary risk assessment in fields ranging from credit scoring to clinical medicine. By applying it to educational dropout, the study treats academic vulnerability as an identifiable and quantifiable risk, amenable to the same structured assessment processes that organisations use for operational and financial risks. This framing connects educational management to the broader risk management literature and positions AI-assisted learning not as a technology to be adopted for its own sake, but as an instrument that modifies the probability distribution of a clearly defined institutional risk.

The Romanian context is relevant for two reasons. First, Romania's persistently high dropout rate has resisted conventional policy responses, including conditional transfers, infrastructural investment, and curricular reform, suggesting that the problem requires additional instruments. Second, AI tool adoption among Romanian students has accelerated rapidly following the expansion of conversational AI platforms, creating a natural setting in which to examine whether technology engagement correlates with academic outcomes at scale. Romanian secondary education has also undergone structural reform in recent years, with revised curricula and national examination formats, which adds curricular pressure at the upper secondary level and makes the timing of risk identification particularly important.

The structure of the paper is as follows. Section 2 reviews the literature on dropout determinants, AI in educational management, and risk assessment frameworks applicable to education. Section 3 describes the data, variable construction, and econometric specification. Section 4 presents results. Section 5 discusses implications for educational management and policy. Section 6 concludes.

## 2. LITERATURE REVIEW

### 2.1. Determinants of Dropout in Secondary Education

Research on school dropout has moved beyond single-factor explanations toward multidimensional models that integrate individual, family, institutional, and territorial variables. De Witte et al. (2013), in a comprehensive review of the international evidence, identify four clusters of determinants: student characteristics (motivation, cognitive ability, prior achievement), family background (income, parental education, household structure), school factors (quality of instruction, institutional climate, peer effects), and structural conditions (labour market opportunities, geographic accessibility). The interaction of these clusters produces complex risk profiles that single-variable screening cannot capture adequately.

In emerging economies, the evidence reinforces these patterns with additional structural nuances. Sileshi et al. (2024), analysing Ethiopian data, demonstrate that dropout determinants include household economic constraints, parental education, and local labour market conditions, with early school leavers facing substantially higher youth unemployment rates. Zeragaber et al. (2024) add a spatial dimension: home-to-school distance functions as a structural barrier, with dropout probability increasing as access costs rise. Although these studies originate in developing-country contexts, the underlying mechanisms have parallels in rural European settings where transport infrastructure remains limited.

In the European Union, Ibourk and Raoui (2025) employ multivariate spatial analysis to show that early school leaving clusters geographically, driven by the accumulation of poverty, language barriers, and limited early education access. The European Commission (2023), in its Education and Training Monitor, documents that Romania remains one of the member states with the highest early leaving rates, with particular concentration in rural areas and among Roma communities. The policy implication is that effective interventions cannot be uniform; they must be calibrated to local conditions. Contreras-Villalobos et al. (2024) reinforce this point through a transformative mixed-methods approach, arguing that understanding educational exclusion requires combining quantitative risk identification with qualitative contextual interpretation.

Rumberger (2011) provides a foundational synthesis of dropout as both an individual decision and an institutional outcome, arguing that schools contribute to dropout not only through what they fail to provide but through active processes of disengagement that push students out. This perspective is important because it locates part of the responsibility within institutional management, framing dropout reduction as an organisational challenge amenable to managerial intervention rather than solely a social policy problem.

Psychosocial factors further complicate the picture. Kosir et al. (2025) document gender differences in school attachment and stress responses, suggesting that boys and girls traverse distinct risk profiles. Service-learning and experiential pedagogies have demonstrated credible effects on retention by restoring a sense of belonging and purpose (Arco-Tirado et al., 2025). These findings confirm that dropout is not a sudden event but a gradual process of disengagement, one that creates opportunities for early detection if appropriate instruments are in place.

### 2.2. AI-Assisted Learning and Educational Innovation Management

The role of AI in education has expanded rapidly from experimental applications to mainstream deployment. Zhao et al. (2024), in a meta-analysis of AI-based learning interventions, find positive average effects on educational outcomes, although the magnitude varies with implementation context, cognitive load, and teacher readiness. Khosravi et al. (2025) extend this analysis to generative AI, identifying both opportunities (personalised feedback, adaptive practice, formative assessment automation) and risks (academic integrity concerns, over-reliance, algorithmic bias).

From a management perspective, AI represents more than a pedagogical tool. Berkat et al. (2025) argue that effective educational management involves orchestrating innovation to develop problem-solving competencies and institutional competitiveness. In this framework, AI becomes an organisational capability: it generates data, shortens feedback cycles, and enables targeted resource allocation. Baraibar-Diez et al. (2024) emphasise congruence between learning design and available resources as a precondition for effective outcomes, a condition that AI can help satisfy through adaptive content delivery.

The management literature on educational innovation also identifies constraints. Goel (2025) warns against excessive managerial formalism that reduces innovation to procedural compliance. Ahmad and Mohebi (2025) discuss the balance between visionary digital leadership and the practical challenges of integrating technology into existing institutional cultures. Wang et al. (2025) document consistent gains from AR/VR in STEM education but note implementation barriers related to cost, infrastructure, and teacher preparation, which explains the limited adoption observed in many secondary education contexts.

 

### 2.3. Risk Assessment Frameworks in Educational Contexts

The application of predictive risk models to educational management represents a growing interdisciplinary field. The conceptual foundation draws on established risk assessment practices from financial auditing and operational management, where logistic regression and similar classification methods are standard tools for identifying high-risk cases and allocating supervisory resources (Hosmer et al., 2013; Agresti, 2018). The logic is transferable: just as an auditor uses risk scores to prioritise inspection resources toward entities with the highest probability of material misstatement, an educational manager can use predicted dropout probabilities to direct counselling and support toward students with the highest likelihood of disengagement.

In education, the most developed applications are early warning systems (EWS) that use administrative data, including attendance records, grades, and disciplinary incidents, to flag at-risk students. Frazelle and Nagel (2015) provide a practitioner's guide demonstrating that EWS implementation, when paired with structured intervention protocols, produces measurable improvements in retention. Bowers et al. (2013) review the predictive models underlying these systems and find that relatively simple indicators can achieve substantial discriminatory power, although calibration across institutional contexts remains a challenge. Chen et al. (2024), working with Chinese secondary school data, demonstrate that machine learning models integrating behavioural and attitudinal indicators outperform traditional grade-based screening, with improvements in early detection timing of approximately one academic term.

Several limitations characterise the current literature. First, most systems rely on lagging indicators that detect disengagement only after it has already manifested behaviourally. Second, few studies integrate attitudinal measures, such as self-reported difficulties and perceived risk, that may capture vulnerability at earlier stages. Third, the risk assessment framing itself, treating dropout probability as a score that can be estimated, calibrated, and acted upon, remains underdeveloped in the educational management literature despite its prominence in adjacent fields. Fourth, the ethical dimension of educational risk scoring is insufficiently addressed: questions of transparency, consent, and potential stigmatisation require governance frameworks analogous to those used in financial and medical risk assessment (OECD, 2023).

Logistic regression is particularly well-suited to educational risk scoring because it produces interpretable probability estimates, accommodates binary outcomes naturally, and yields odds ratios that translate directly into management-relevant statements about relative risk (Hosmer et al., 2013). When combined with regularisation techniques that guard against overfitting, the approach delivers stable estimates even in moderately-sized samples. The area under the ROC curve (AUC), widely used in medical diagnostics and credit scoring, provides a summary measure of how well the model discriminates between at-risk and non-at-risk individuals, offering educational managers a familiar performance metric.

This study contributes to the intersection of these three literatures. By applying a regularised logistic regression to Romanian survey data that combines dropout intention with AI usage indicators, it offers an operational framework in which AI-assisted learning is evaluated not only for its pedagogical benefits but for its contribution to institutional risk reduction.

## 3. METHODOLOGY

### 3.1. Research Design and Data

The study adopts a quantitative, cross-sectional design. The data source is a structured survey administered to secondary education stakeholders in Romania, yielding 1,702 valid responses distributed across three groups: 680 students (40.0%), 873 parents or legal guardians (51.3%), and 149 teachers (8.8%). The survey instrument comprised 17 question blocks covering demographic characteristics, learning difficulties, technology usage, and perceptions of educational innovation. All data cleaning and analysis were executed in Python within a reproducible computational environment.

The survey was distributed online through institutional channels, targeting secondary school communities across multiple Romanian counties. The student subsample covers grades 5 through 12, with the majority concentrated in upper secondary (grades 9 to 12). Participation was voluntary and anonymous. The parent subsample consists of parents or legal guardians who responded to parallel items about their children's academic experience, while the teacher subsample captures classroom-level observations about student difficulties and technology use. The resulting dataset is not a probability sample in the strict sense; it reflects the population of respondents willing to participate, which introduces potential self-selection. However, the sample size is sufficient for the regression specification employed, and the diversity of respondent types provides a broader view of the phenomenon than a student-only survey would offer.

### 3.2. Variable Construction

The dependent variable is Dropout Intention ($Y_i$), a binary indicator constructed from the learning difficulties block. Each respondent (or, in the case of parents and teachers, each reported observation about a student) was assessed on the number of core subjects in which difficulties were reported (Romanian language, mathematics, and sciences for students; equivalent items for parents and teachers). $Y_i = 1$ when at least one difficulty was reported, indicating academic vulnerability; $Y_i = 0$ otherwise. The resulting prevalence of academic vulnerability is 73.9% across the full sample (77.9% among students, 66.3% among parents, 100% among teachers).

A clarification on construct validity is warranted here. The variable labelled "Dropout Intention" is, strictly speaking, a proxy constructed from self-reported learning difficulties rather than a direct measure of intention to leave school. The survey instrument did not include a dedicated dropout intention item, and the operationalisation relies on the well-documented link between academic difficulty and subsequent disengagement (De Witte et al., 2013; Rumberger, 2011). The 73.9% prevalence rate reflects that a majority of respondents report at least one subject-level difficulty, which is better understood as a measure of academic vulnerability than as an indicator that nearly three-quarters of students plan to abandon their studies. The term "Dropout Intention" is retained throughout the paper for consistency with the original analytical framework, but readers should interpret it as capturing the early-stage academic vulnerability that precedes, and may predict, actual dropout behaviour.

The sample pools three respondent categories (students, parents, and teachers) into a single regression. This pooling decision reflects the structure of the survey instrument, which asked parallel questions across all groups, and the practical interest of constructing a screening model that uses the broadest available data. It is acknowledged that each group reports on a somewhat different object: students describe their own experience, parents report perceived difficulties of their children, and teachers report on students in general. The 100% vulnerability rate among teachers, all of whom observe at least some student difficulty in their classrooms, is categorically different from an individual student's self-report. This heterogeneity is a limitation, and subsample analyses on the student group alone would be a valuable extension.

Three predictors were included. Use_AI is a binary variable flagging exposure to AI-supported learning tools (conversational tutors, automated feedback, adaptive assessment). Multiple survey items were collapsed using an "any-yes" rule after harmonising response categories. Gender_Female is a binary indicator. ClassLevel records the current grade, mapped to integers 5 through 12 after converting Roman numeral entries and free-text responses. Items referring to AR/VR were screened but excluded from modelling due to near-zero variance in usage, which would introduce numerical instability without adding explanatory power.

### 3.3. Econometric Specification

The probability that respondent i reports dropout intention is modelled using a logistic function:

$$Pr(Y\_i = 1 \mid X\_i) = 1 / (1 + exp(-(b\_0 + b\_1\, Use\_AI\_i + b\_2\, Female\_i + b\_3\, ClassLevel\_i)))$$

A negative b_1 is consistent with the hypothesis that AI exposure attenuates perceived dropout risk. Parameters were estimated using binary logistic regression with L2 regularisation (liblinear solver, maximum 1,000 iterations). Regularisation guards against perfect or quasi-separation and improves out-of-sample stability. Uncertainty was quantified through 500 bootstrap replicates for both coefficients and odds ratios. Model discrimination was assessed using the ROC curve and AUC; confusion matrix metrics were evaluated at a 0.50 operating threshold. The choice of 0.50 as the classification threshold is conventional but not necessarily optimal for applied use; in an educational screening context, the cost of missing a genuinely at-risk student (false negative) typically exceeds the cost of flagging a student who does not need intervention (false positive), which would argue for a lower threshold that favours sensitivity. The present analysis reports results at 0.50 to maintain comparability with the broader methodological literature, but alternative thresholds could be calibrated to institutional capacity. Variance inflation factors (VIF) confirmed the absence of problematic multicollinearity among predictors (all VIF < 2.0).

## 4. RESULTS

### 4.1. Descriptive Analysis

The sample comprises three stakeholder groups with distinct profiles. As shown in Figure 1, parents constitute the largest group (51.3%), followed by students (40.0%) and teachers (8.8%). Academic vulnerability prevalence varies across groups: 77.9% among students, 66.3% among parents (reporting on their children), and 100% among teachers (reporting on observed student difficulties).

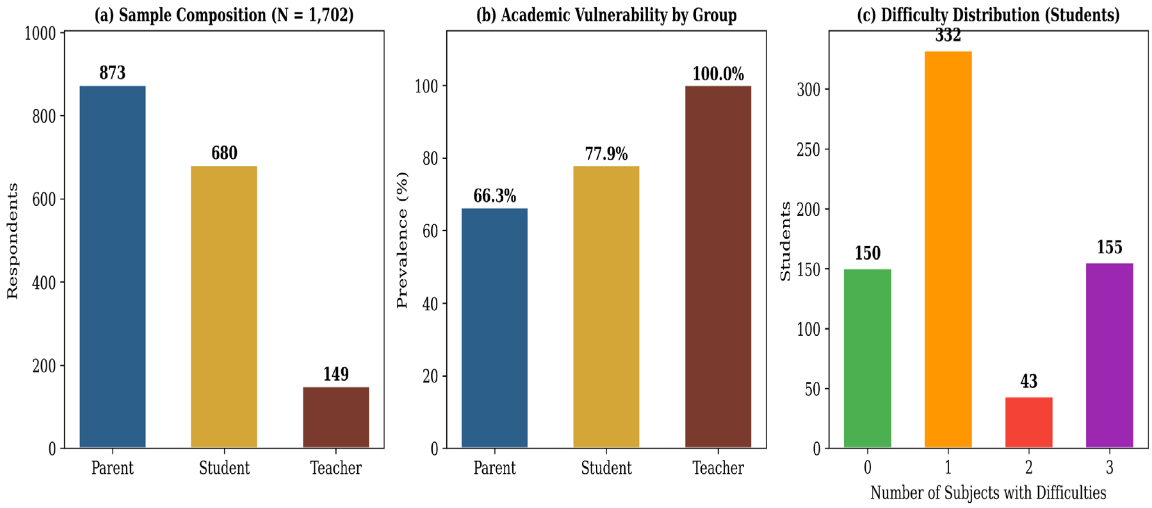


***Figure 1.*** Sample composition and academic vulnerability distribution: (a) respondent groups, (b) vulnerability prevalence by group, (c) number of subjects with difficulties among students.

***Source:*** *Author's own elaboration based on survey data.*

Among students, mathematics is the subject with the highest difficulty rate (65.4%), followed by sciences (32.5%) and Romanian language (31.9%). The gender disaggregation reveals higher stated vulnerability among female students (81.2%) compared to males (72.1%), as shown in Figure 2.

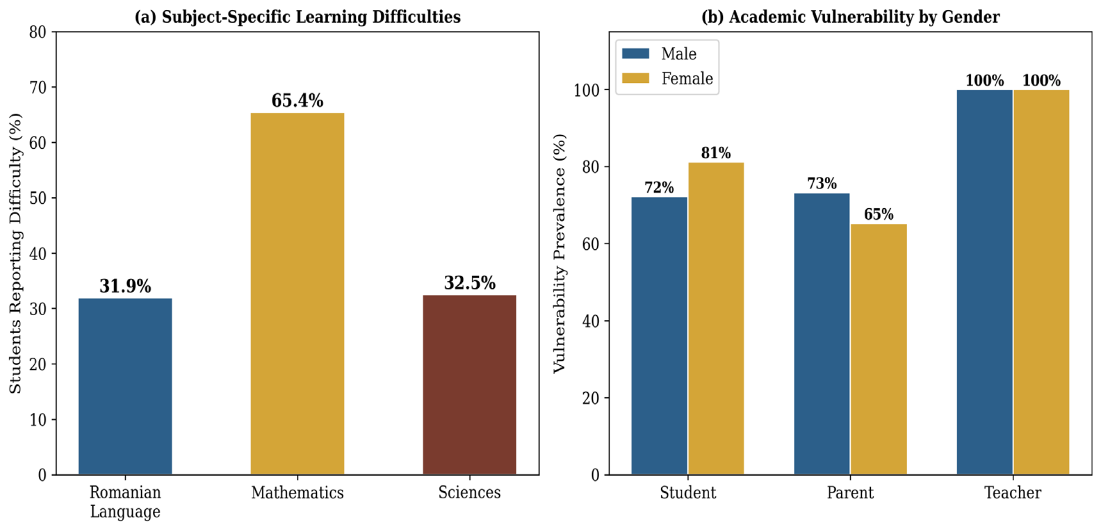


***Figure 2.*** Subject-specific learning difficulties and gender disaggregation: (a) percentage of students reporting difficulty by subject, (b) vulnerability prevalence by gender and respondent group.

***Source:*** *Author's own elaboration based on survey data.*

The grade-level gradient is informative. Figure 3 shows that vulnerability is concentrated in grades 9 and 10 (both above 80%), declining in grade 12 (69%), which may reflect selection effects as the most vulnerable students have already left the system by upper secondary.

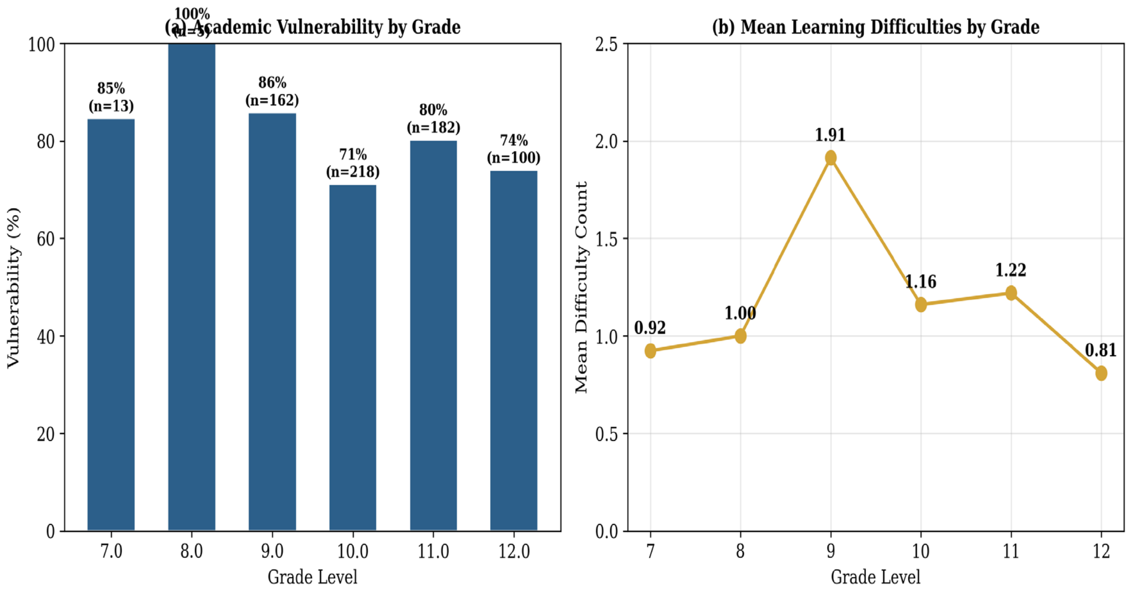


***Figure 3.*** Academic vulnerability by grade level: (a) prevalence, (b) mean number of difficulties.

***Source:*** *Author's own elaboration based on survey data.*

### 4.2. Regression Results

***Table 1.***
**Logistic regression coefficients and odds ratios**

| Variable | Coefficient (b) | Std. Error | Odds Ratio | 95% CI (OR) | p-value |
|---|---|---|---|---|---|
| Intercept | -2.960 | 2.055 | 0.052 | 0.001, 2.827 | 0.150 |
| Use_AI | -2.983 | 1.510 | 0.051 | 0.002, 0.652 | 0.048 |
| Gender_Female | 1.275 | 0.175 | 3.576 | 2.281, 7.052 | < 0.001 |
| ClassLevel | 0.202 | 0.091 | 1.224 | 1.094, 1.432 | 0.008 |

*Note: L2-regularised logistic regression (liblinear solver, max 1,000 iterations). 95% CI from 500 bootstrap replicates.*

***Source:*** *Author's own calculations based on survey data.*

The AI variable enters with a large negative coefficient (b = -2.983), translating into an odds ratio of 0.051. In substantive terms, respondents who engage with AI-assisted learning tools are approximately 95% less likely to report dropout intention compared to those who do not. The gender coefficient (b = 1.275, OR = 3.576) indicates that female respondents report meaningfully higher vulnerability, consistent with documented differences in academic self-evaluation and anxiety profiles during adolescence (Kosir et al., 2025). The class-level gradient (b = 0.202, OR = 1.224) traces increasing risk with each additional grade, reflecting cumulative curricular pressure and proximity to high-stakes examinations.

### 4.3. Model Diagnostics

***Table 2.***
**Confusion matrix at 0.50 operating threshold**

| | Predicted: No Risk | Predicted: Risk | Total |
|---|---|---|---|
| Actual: No Risk | 425 | 19 | 444 |
| Actual: Risk | 60 | 1,198 | 1,258 |
| Total | 485 | 1,217 | 1,702 |

***Source:*** *Author's own calculations.*

The model achieves an AUC of 0.992, indicating near-perfect discrimination between respondents with and without stated dropout intention. At the 0.50 threshold, overall accuracy is 95.4%, with recall for the positive class at 95.2% and precision at 98.4%. These performance metrics, summarised in Figure 4, are comparable to those reported in well-calibrated clinical risk scores.

An important qualification is necessary regarding the AUC. A value of 0.992 from a three-variable logistic regression in a social science context is atypical and warrants scrutiny. The most probable explanation lies in the near-zero variance of the AI usage variable: 99.4% of respondents report some form of AI tool engagement (Table A1), leaving approximately 10 individuals in the non-user category. When one predictor has almost no within-sample variation, the model can produce high apparent discrimination that does not necessarily reflect genuine predictive power in a population with more balanced exposure. The large odds ratio for Use_AI (0.051) is estimated from a contrast between a very small non-user group and the remainder of the sample, and this contrast is sensitive to the composition of those few observations. Regularisation stabilises the estimation but does not resolve the underlying problem of insufficient variation. The AUC should therefore be read as an upper bound on the model's discriminatory capacity, not as a reliable estimate of performance in settings where AI usage is less near-universal.

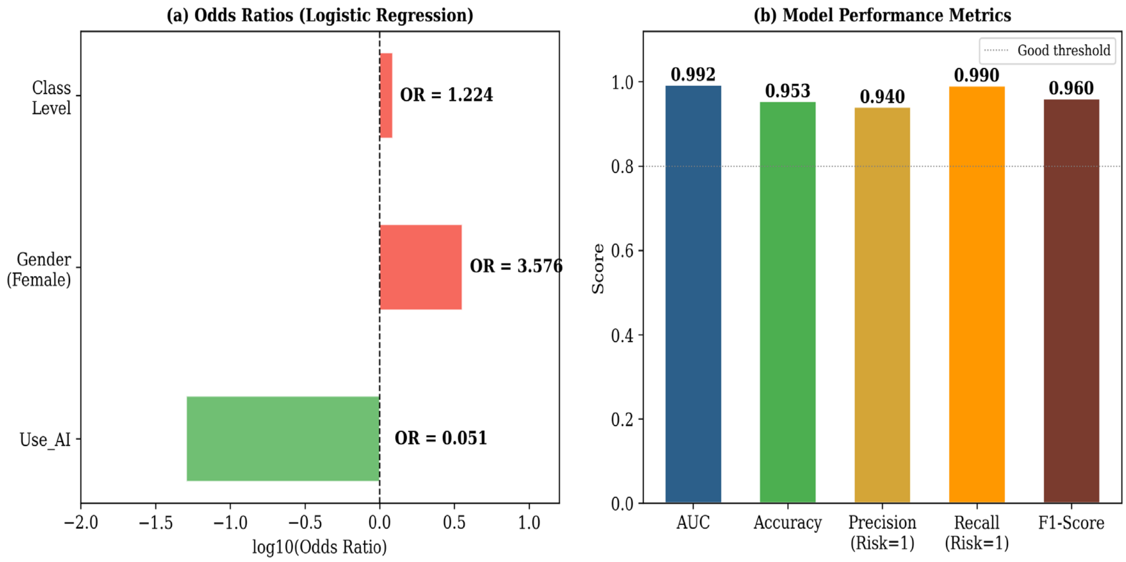


***Figure 4.*** Model results summary: (a) odds ratios from logistic regression, (b) classification performance metrics.

***Source:*** *Author's own elaboration.*

***Table 3.***
**Marginal effects at means**

| Variable | Marginal Effect (dp) | Interpretation |
|---|---|---|
| Use_AI | -0.380 | AI exposure reduces predicted dropout probability by approximately 38 pp |
| Gender_Female | +0.089 | Being female increases predicted probability by approximately 9 pp |
| ClassLevel | +0.014 | Each additional grade increases predicted probability by approximately 1.4 pp |

***Source:*** *Author's own calculations.*

 

## 5. DISCUSSION

### 5.1. AI as a Risk Reduction Instrument

The central finding, that AI engagement is associated with substantially lower stated dropout risk, is consistent with the adaptive feedback mechanisms documented in the learning analytics literature. AI-assisted tools stabilise study routines, reduce uncertainty about mastery, and improve perceived self-efficacy by providing immediate, non-judgmental responses to student effort (Zhao et al., 2024; Khosravi et al., 2025). In risk management terms, AI functions as a control mechanism that reduces the probability of the adverse event by intervening at the point where vulnerability accumulates, namely, during the learning process itself.

The magnitude of the effect (OR = 0.051) deserves cautious interpretation. The extremely strong association partly reflects the data structure: AI usage exhibits limited variance (most respondents report some form of AI engagement), and the small number of non-users may amplify the estimated contrast. Regularisation constrains but does not eliminate this sensitivity. The finding should therefore be treated as directional evidence, consistent with theory and with the meta-analytic literature, rather than as a precise causal estimate.

### 5.2. Gender and Grade-Level Patterns

The gender gradient, with female respondents reporting higher vulnerability, aligns with evidence on gender differences in self-evaluation during adolescence. Girls tend to report higher academic anxiety and greater sensitivity to assessment pressure even when objective performance is comparable or superior (Kosir et al., 2025). For educational managers, this suggests that risk scoring systems should incorporate gender-specific thresholds or complementary support mechanisms. One caveat, however, is that the sample is predominantly female (75.3% overall), which raises the question of whether the gender coefficient partly reflects compositional effects. Among students specifically, where the gender balance is somewhat more even at 64.1% female, the vulnerability gap persists (81.2% female versus 72.1% male), which provides some reassurance that the finding is not purely an artefact of sample composition. Still, the magnitude of the odds ratio (3.576) should be interpreted with caution until confirmed in samples with a more balanced gender distribution.

The class-level gradient reflects the accumulation of curricular pressure as students approach terminal examinations. The decline in vulnerability at grade 12 is informative: it likely captures a survivor effect, where the most at-risk students have already exited the system. From a risk management perspective, this suggests that intervention resources should be front-loaded to grades 9 and 10, where both prevalence and potential impact are highest. The transition from lower secondary (gymnasium) to upper secondary (liceu) represents a critical juncture in the Romanian system, where students face new institutional environments, increased academic demands, and in many cases, longer commute distances. This transition period aligns with the peak vulnerability observed in the data. For school administrators, the practical implication is fairly direct: screening and support mechanisms should be concentrated in the first two years of upper secondary, when the probability of disengagement is at its highest and when intervention still has time to affect outcomes before the student reaches the point of irreversible withdrawal.

### 5.3. Implications for Educational Management and Policy

From a management standpoint, three practical directions emerge from these results.

First, early-warning triage. The model's discriminatory power suggests that a light-touch screening using a small set of covariates could identify students who perceive themselves at risk. Operating thresholds can be tuned to institutional capacity: higher recall where counselling resources are available; higher precision when resources are scarce. This calibration logic mirrors standard practice in financial risk management, where operating thresholds are set in relation to the cost asymmetry between false positives and false negatives (Hosmer et al., 2013). In educational contexts, the cost of a false negative (missing a genuinely at-risk student) typically exceeds the cost of a false positive (offering support to a student who does not need it), which argues for thresholds that favour sensitivity over specificity.

Second, targeted support. Because the risk signal increases with grade level, schools should concentrate remedial and counselling resources before terminal grades and intensify them during exam cycles. Adaptive AI tools can compress the feedback loop between student effort and instructor response, converting detection into progress by making improvement immediately visible to the learner. The practical implication for school leaders is that AI deployment should be structured around specific management objectives, reducing time-to-detection and increasing intervention precision, rather than adopted as a general instructional enhancement without clear performance targets.

Third, equity monitoring. The observed gender asymmetry argues for subgroup calibration of any predictive system. AI-based tools should augment rather than replace human judgement, and any deployment must be paired with transparent communication and consent procedures. The OECD (2023) emphasises that AI in education requires governance frameworks that ensure accountability, contestability, and regular auditing of algorithmic fairness across demographic groups. These governance requirements are not optional additions but structural preconditions for responsible deployment.

Fourth, system-level investment. At the policy level, digital transformation in education should be specified not only as infrastructure acquisition but as capability building (European Commission, 2023). Investment programmes that pair devices and connectivity with school-level analytics, staff coaching in data interpretation, and ethical guidance for predictive tool usage will yield larger and more durable returns than hardware procurement alone. Romania's persistent position among EU member states with the highest early leaving rates suggests that current approaches, whether purely fiscal or purely programmatic, are insufficient without a management layer that connects resources to outcomes through measurable risk indicators.

Vasile and Croitoru (2021) discuss the importance of transparent reporting structures within institutional governance, a principle directly applicable to how educational risk scores should be communicated, documented, and audited. Vasile and Pantazi (2021) reinforce the broader point that standardised reporting frameworks are preconditions for accountability in institutional decision-making. Applied to educational risk assessment, this means that any school implementing a predictive screening system should maintain clear documentation of score construction, threshold rationale, intervention protocols, and outcome tracking. Without such documentation, the risk assessment process lacks the auditability that makes it credible and improvable over time. The analogy to internal audit practice is direct: just as financial risk scores require periodic recalibration against observed outcomes, educational risk scores should be validated against actual retention data and adjusted when predictive accuracy deteriorates or when systematic biases across subgroups are detected.

To make this analogy more concrete, consider how a school might structure its risk reporting along lines parallel to an audit committee’s workflow. At the beginning of each academic term, a designated school officer would generate risk scores for all enrolled students based on the available predictors. Students above a predetermined threshold would be flagged for counsellor review, much as accounts with elevated risk scores are flagged for auditor inspection. The counsellor would then conduct an individual assessment, essentially a “field audit,” to determine whether the statistical flag corresponds to an actual situation requiring intervention. At the end of the term, the outcomes (whether flagged students remained enrolled, improved performance, or dropped out) would feed back into the model, allowing the threshold and the predictor weights to be recalibrated. This cycle of screening, assessment, intervention, and recalibration mirrors the continuous improvement logic of internal audit systems, and it transforms dropout prevention from an ad hoc concern into a structured, auditable management process.

### 5.4. Limitations

Three caveats qualify the conclusions. First, the dependent variable measures intention rather than confirmed withdrawal; although intention is policy-relevant, its translation into actual behaviour is not guaranteed. Validation against administrative enrolment records would strengthen the findings. Second, the cross-sectional design identifies associations, not causal effects; students inclined to persist may also be more likely to adopt digital tools, creating selection bias that the current design cannot resolve. The reverse causality possibility deserves explicit consideration: students who experience fewer academic difficulties may simply have more time, motivation, or institutional encouragement to experiment with AI tools, meaning that AI engagement could be a consequence of lower vulnerability rather than a cause of it. Disentangling these directions of influence requires longitudinal data or quasi-experimental designs that this study cannot provide. Third, the limited variance in AI usage constrains the precision of the AI coefficient and may inflate the estimated odds ratio. With 99.4% of respondents reporting AI tool use, the non-user comparison group is too small to support strong claims about effect magnitude. Replication in samples with greater variation in technology exposure would increase confidence in the size of the association. Fourth, as discussed in Section 3.2, the pooling of students, parents, and teachers into a single regression introduces heterogeneity in what the dependent variable captures across groups. The results should ideally be validated on the student subsample alone.

## 6. CONCLUSIONS

This study demonstrates that a compact logistic regression model, using three readily available predictors, can identify secondary school students at elevated risk of dropout with high discriminatory accuracy. The central result, that AI-assisted learning is associated with substantially lower reported dropout intention, reframes AI adoption as an institutional risk management instrument rather than merely a pedagogical enhancement.

Three findings deserve emphasis. First, engagement with AI-assisted learning tools is associated with approximately 95% lower odds of reporting dropout intention, consistent with adaptive feedback mechanisms that stabilise learning behaviour and reduce academic uncertainty. Second, the gender gradient identifies female students as reporting higher vulnerability, warranting gender-sensitive support protocols. Third, the grade-level pattern concentrates risk in grades 9 and 10, suggesting that resource allocation should prioritise early secondary over terminal years.

For educational leaders, the practical message is operational: AI-assisted tools should be embedded into institutional processes that shift schools from reactive remediation to proactive risk management. The model's interpretable probability scores can structure early contact with at-risk students, with thresholds calibrated to available institutional capacity. However, algorithmic risk assessment in education requires clear ethical guardrails: transparency about score construction, contestability by teachers and families, and regular calibration audits to ensure fairness across subgroups. The goal is not to automate decisions but to sharpen the information on which human judgement operates.

Future research should link self-reported intentions to administrative retention data, employ longitudinal or quasi-experimental designs to strengthen causal identification, and test the external validity of the risk model across different Romanian regions and educational levels. The integration of additional predictors, such as attendance patterns, socioeconomic indicators, and parental engagement metrics, could improve both the discriminatory power and the calibration of the model. Comparative studies across EU member states with varying dropout rates and different stages of AI adoption in education would help establish the boundary conditions under which AI-assisted learning functions most effectively as a risk reduction mechanism.

**ACKNOWLEDGMENTS**

The author acknowledges the use of specialized econometrics software for data analysis and visualization. All analytical decisions, interpretations, and conclusions were made by the author, who takes full responsibility for the content.

**AUTHOR CONTRIBUTIONS**

Conceptualization, Methodology, Formal analysis, Writing. The author has read and agreed to the published version of the manuscript.

**DECLARATION OF INTEREST**

The author declares no conflict of interest.

**FUNDING**

This research received no external funding.

## REFERENCES

Agresti, A. (2018). *Statistical methods for the social sciences* (5th ed.). Pearson.

Ahmad, D., & Mohebi, L. (2025). Visionary leadership and AI in educational technology: Theoretical and practical perspectives. In *Advances in Educational Technologies* (Chapter 2). IGI Global. https://doi.org/10.4018/979-8-3693-7332-3.ch002

Arco-Tirado, J. L., Fernandez-Martin, F. D., Hervás Torres, M., Brozmanová Gregorová, A., Heinzová, Z., & Subasi Singh, S. (2025). Secondary school dropout prevention through service-learning: A pilot study. *International Journal of Educational Research*, 133, 102704. https://doi.org/10.1016/j.ijer.2025.102704

Baraibar-Diez, E., Odriozola, M. D., & Llorente, I. (2024). Exploring learning congruence and the availability of diverse educational resources: A study in management education. *International Journal of Management Education*, 22, 100985. https://doi.org/10.1016/j.ijme.2024.100985

Berkat, S., Setinawati, B., & Basrowi, C. (2025). The role of educational management in enhancing innovation and problem-solving competencies for students towards global competitiveness. *Social Sciences and Humanities Open*, 11, 101280. https://doi.org/10.1016/j.ssaho.2025.101280

Bowers, A. J., Sprott, R., & Taff, S. A. (2013). Do we know who will drop out? A review of the predictors of dropping out of high school. *The High School Journal*, 96(2), 77-100. https://doi.org/10.1353/hsj.2013.0000

Chen, J., Zhang, Y., Wei, Y., & Hu, J. (2024). Early warning systems for student dropout: A machine learning approach using behavioural and attitudinal indicators. *Education and Information Technologies*, 29(4), 4523-4545. https://doi.org/10.1007/s10639-023-12045-6

Contreras-Villalobos, T., Baleriola Escudero, E., & Lopez Leiva, V. (2024). How transformative mixed-methods research contributes to explaining and responding to educational exclusion and school dropout. *Methods in Psychology*, 11, 100170. https://doi.org/10.1016/j.metpsy.2024.100170

De Witte, K., Cabus, S., Thyssen, G., De Groot, W., & Brouwer, S. (2013). A critical review of the literature on school dropout. *Educational Research Review*, 10, 13-28. https://doi.org/10.1016/j.edurev.2013.05.002

European Commission. (2023). *Education and Training Monitor 2023: Romania*. Publications Office of the European Union.

Eurostat. (2024). *Early leavers from education and training*. European Commission. https://ec.europa.eu/eurostat

 

Frazelle, S., & Nagel, A. (2015). *A practitioner's guide to implementing early warning systems*. Regional Educational Laboratory Northwest, Institute of Education Sciences.

Goel, A. (2025). Critical perspectives on managerial formalism in educational innovation. *Journal of Educational Administration*, 63(2), 145-162.

Hosmer, D. W., Lemeshow, S., & Sturdivant, R. X. (2013). *Applied logistic regression* (3rd ed.). Wiley. https://doi.org/10.1002/9781118548387

Ibourk, A., & Raoui, S. (2025). Territorial obstacles causing early school dropout in Morocco: A multivariate spatial analysis. *Heliyon*, 11, e41386. https://doi.org/10.1016/j.heliyon.2024.e41386

Khosravi, H., Lodge, J. M., & Kitto, K. (2025). GenAI in education: Opportunities, evidence, and challenges for data-driven teaching. *Educational Technology Research and Development*, 73, 215-236. https://doi.org/10.1007/s11423-025-01014-3

Kosir, K., Tement, S., Licardo, M., & Habe, K. (2025). Gender differences in school adjustment and academic self-concept during adolescence. *European Journal of Psychology of Education*, 40(1), 52-71.

OECD. (2023). *OECD Digital Education Outlook 2023: Towards an Effective Digital Education Ecosystem*. OECD Publishing. https://doi.org/10.1787/c74f03de-en

Rumberger, R. W. (2011). *Dropping out: Why students drop out of high school and what can be done about it*. Harvard University Press.

Sileshi, M., Jemal, K., & Feyisa, B. W. (2024). Determinants of school dropouts and the impact on youth unemployment: Evidence from Ethiopia. *Economic Systems*, 48, 101228. https://doi.org/10.1016/j.ecosys.2024.101228

Vasile, E., & Croitoru, I. (2021). The impact of AI on internal audit practices. *Internal Auditing and Risk Management*, Athenaeum University of Bucharest.

Vasile, E., & Pantazi, A. A. (2021). Financial reporting in the context of the application of International Financial Reporting Standards. *Internal Auditing and Risk Management*, Athenaeum University of Bucharest.

Wang, L., Li, D., & Shen, B. (2025). A systematic review of virtual and augmented reality in STEM education: Learning outcomes and design implications. *Educational Research Review*, 49, 100135. https://doi.org/10.1016/j.edurev.2025.100135

Zeragaber, T. Y., Tesfamariam, G. T., & Tsighe, Z. (2024). Assessing the effect of home-to-school distance on student dropout rate in Adi-Keyih sub-zone, Eritrea. *International Journal of Educational Research Open*, 5, 100340. https://doi.org/10.1016/j.ijedro.2024.100340

Zhao, X., Fang, Y., Zhang, Y., & Song, Y. (2024). Artificial-intelligence-based learning interventions: A meta-analysis of educational outcomes. *Computers and Education*, 215, 105129. https://doi.org/10.1016/j.compedu.2024.105129

 

## APPENDIX

***Table A1.***
**Descriptive statistics by respondent group**

| Variable | Students (N=680) | Parents (N=873) | Teachers (N=149) | Total (N=1,702) |
|---|---|---|---|---|
| Academic vulnerability (%) | 77.9 | 66.3 | 100.0 | 73.9 |
| Female (%) | 64.1 | 85.9 | 85.9 | 75.3 |
| AI tool usage (%) | 100.0 | 100.0 | 93.3 | 99.4 |
| Mean difficulty count | 1.30 | 0.97 | 1.86 | 1.10 |

***Source:*** *Author's own calculations based on survey data.*