INTERNATIONAL JOURNAL OF MANAGEMENT AND ORGANIZATIONAL RESEARCH

Advancing Data-Driven, Risk-Based Audit Analytics to Strengthen Fraud Detection, Internal Controls, Financial-Reporting Reliability, and Regulatory Compliance Across U.S. Organizations

Prosper Fuadedzi
Georgia State University Atlanta, GA, USA

* Corresponding Author: **Prosper Fuadedzi**

Article Info

ISSN (online): 2583-6641
Volume: 05
Issue: 06
Received: 24-08-2026
Accepted: 26-09-2026
Published: 28-10-2026
Page No: 18-24

Abstract
Auditors and regulators can review only a small frac-tion of companies in depth, so they need reliable ways to de-cide where to look. This study proposes and evaluates X-RBA, an explainable machine learning framework for risk-based audit analytics that ranks U.S. public companies by misstatement risk and explains each ranking in terms of audit risk areas. X-RBA combines cost-sensitive gradient-boosted trees trained on 28 raw financial-statement items with Shapley-value attributions aggre-gated into seven auditor-facing risk areas. On a public dataset of 146,045 U.S. firm-years labeled with SEC enforcement releases, under a strict out-of-sample protocol over test years 2003–2008, X-RBA achieves a mean AUC of 0.713, significantly higher than a re-implemented RUSBoost benchmark (0.675), and its hit rate among the riskiest 1% of firms is about three times the base rate. Budget-based evaluation shows that model rankings depend on the review budget. The explanations identify equity issuance and financing activity as the dominant risk drivers and show how a flagged firm's profile can direct substantive procedures and internal-control eval-uation, supporting scalable and transparent fraud detection and regulatory compliance.

Keywords: audit analytics, accounting fraud detection, financial misstatement, risk-based auditing, explainable artificial intel-ligence, shapley values, gradient boosting, internal controls, sec enforcement

Introduction

Reliable financial reporting is a public good. Investors, lenders, employees, regulators, and taxpayers all rely on audited finan-cial statements, and when those statements are materially mis-stated the costs extend far beyond the firm that issued them. The accounting scandals of the early 2000s led directly to the Sarbanes–Oxley Act of 2002 (SOX), which created the Pub-lic Company Accounting Oversight Board (PCAOB) and re-quired management of U.S. public companies to report on the effectiveness of internal control over financial reporting (ICFR), with independent auditor attestation for larger filers [1]. Two decades later, misstatement and fraud remain a persistent en-forcement priority of the U.S. Securities and Exchange Com-mission (SEC), which documents its actions against issuers and auditors in Accounting and Auditing Enforcement Releases (AAERs).

Auditing standards are explicitly *risk-based*. PCAOB AS 2110 requires the auditor to identify and assess the risks of material misstatement, and AS 2401 requires specific con-sideration of fraud risk, including the presumption of fraud risk in revenue recognition [2, 3]. In practice, however, the inputs to this risk assessment are often judgmental checklists and a small number of ratio screens. Audit researchers have called for an-alytics that use the growing volume of financial data to direct audit effort where it matters most [4].

Machine learning (ML) offers a natural way to rank firms by misstatement risk. Bao *et al.* [5] showed that an ensem-ble learner (RUSBoost) trained on 28 raw financial-statement items can outperform classic ratio-based models such as the F-score of Dechow *et al.* [6] in detecting AAER-identified fraud out of sample. Two gaps nevertheless limit adoption in risk-based auditing. First, the performance metrics emphasized in much of the literature (e.g., overall AUC) do not directly an-swer the auditor's operational question: *given a fixed review budget, what share of misstatements will be caught?* Second, high-performing ensembles are opaque. Auditing standards re-quire the auditor to understand and document the basis for risk assessments, and

professional skepticism cannot be delegated to an unexplained score [7].

This paper addresses both gaps. We propose and evaluate X-RBA, an explainable machine learning framework for risk-based audit analytics that (i) scores every U.S. public firm-year for misstatement risk using cost-sensitive gradient-boosted deci-sion trees trained only on publicly reported financial-statement data; (ii) evaluates the scores the way an audit team would use them, as a ranked list reviewed under a fixed budget; and (iii) decomposes each score into additive Shapley contributions that are aggregated into auditor-facing risk areas such as revenue and receivables, inventory, and financing. All experiments use a public dataset of 146,045 firm-years with SEC AAER fraud labels released by Bao *et al.* [5], so every result in this paper can be reproduced.

The contributions are as follows:

- We propose X-RBA, a four-layer framework (data, learn-ing, explanation, and audit decision) that links ML-based misstatement detection to the risk-assessment workflow re-quired by U.S. auditing standards.
- Under a strict out-of-sample protocol with a two-year label-disclosure gap and serial-fraud correction, we bench-mark X-RBA against a logistic ratio model, a re-implementation of the RUSBoost model of [5], and an un-supervised isolation forest over six test years (2003–2008).
- We report budget-based metrics (fraud capture and preci-sion at the top 1% and 5% of firms) alongside AUC and NDCG@1%, and quantify uncertainty with a firm-level bootstrap.
- We provide global and firm-level explanations, showing which financial-statement areas drive flagged risk and how these explanations can direct substantive audit procedures, strengthen ICFR evaluation, and support regulatory com-pliance.

2. Related Work

2.1. Financial Misstatement and Fraud Prediction

Early detection models relied on a small number of theoreti-cally motivated ratios. Beneish [8] proposed the M-score, a probit model of eight ratios capturing receivables growth, mar-gin deterioration, asset quality, and accruals. Dechow *et al.* [6] built a comprehensive AAER-based sample and developed the F-score from accrual quality, financial performance, nonfinan-cial, off-balance-sheet, and market-based measures; the 14 ratio features used in this paper follow their definitions. Cecchini *et al.* [9] applied support vector machines with a financial kernel, and Perols [10] found that simpler classifiers such as logistic regression and support vector machines were competitive under realistic class imbalance.

2.2. Ensemble Learning on Raw Financial Data

Bao *et al.* [5] argued that ratios constructed by experts may dis-card information and showed that RUSBoost [11], a boosting algorithm that undersamples the majority class at each round, trained on 28 raw financial-statement items, outperformed both logistic regression on ratios and SVM-based models. They also introduced an evaluation protocol that has become a reference point: a two-year gap between training and test periods to reflect the delay in AAER disclosure, and the recoding of serial frauds that span the training and test periods. Bertomeu *et al.* [12] used gradient-boosted regression trees with a broad set of account-ing, capital-market, and governance variables to detect misstate-ments and examined which variables carry the most predictive information. Our work adopts the protocol of [5] and focuses on public financial-statement data only, so the framework can be applied to any SEC registrant.

Table 1: Sample periods and fraud prevalence

Period	Firm-years	Fraud firm-years	Fraud rate (%)
Full sample (1990–2014)	146,045	964	0.66
Primary test (2003–2008)	35,166	261	0.74

For each test year t in 2003–2008, training uses fiscal years 1991 to t − 2.

2.3. Explainability in Audit Analytics

Explainability is particularly important in auditing because the auditor must document the reasoning behind risk assessments. Shapley values from cooperative game theory provide additive feature attributions with desirable consistency properties [13, 14]. Rudin [7] cautions that post-hoc explanations of black-box models can be misleading in high-stakes settings. We therefore use unbiased Monte Carlo Shapley estimates computed on the deployed model, aggregate them to the level at which auditors plan procedures, and present them as decision support rather than as a substitute for professional judgment.

3. Data

3.1. Sample and Fraud Labels

We use the public dataset released by Bao *et al.* [5]. It cov-ers U.S. public firms from fiscal years 1990 to 2014, drawn from Compustat fundamentals, and contains 146,045 firm-years for 18,444 distinct firms. A firm-year is labeled as misstated ($y_{it} = 1$) if it is identified as fraudulent in an SEC AAER, as compiled by the Center for Financial Reporting and Manage-ment at the University of California, Berkeley [5,6]. The sample contains 964 misstated firm-years belonging to 402 firms and 412 distinct AAER cases, for an overall fraud rate of 0.66%. The extreme class imbalance is representative of the real audit setting, in which the overwhelming majority of filings are not fraudulent. Table I summarizes the sample periods used in this study.

3.2. Features

We consider two feature sets, both computable from public fi-nancial statements filed with the SEC.

Raw financial-statement items (28): Following [5], these are line items taken directly from the balance sheet, income state-ment, and cash flow statement, such as total assets, receivables, inventory, sales, cost of goods sold, current liabilities, long-term debt, net income, retained earnings, income taxes, interest ex-pense, issuance of stock and debt, and year-end share price.

Expert risk ratios (14): These follow the misstatement pre-dictors of Dechow *et al.* [6]: working-capital and RSST ac-cruals, changes in receivables and inventory, the share of soft assets, changes in cash sales, cash margin, return on assets and free cash flow, an indicator of securities issuance, book-

to-market, a depreciation index, retained earnings over assets, and EBIT over assets.

Missing values are handled natively by the gradient-boosted trees. For models that require complete inputs, we impute the training-period median, computed separately for each test year to avoid information leakage.

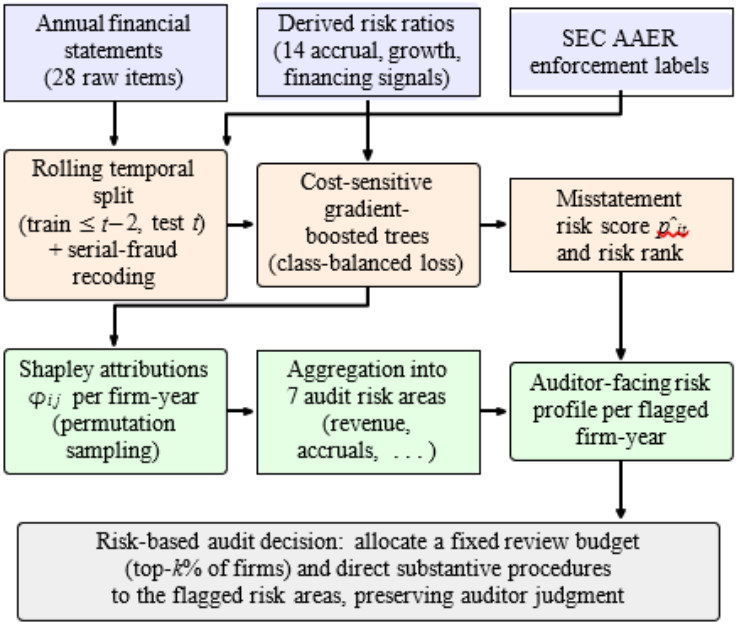


Fig 1: The proposed X-RBA framework. Blue: data layer built from public SEC filings and enforcement labels. Orange: learning layer producing a mis-statement risk score and rank for every firm-year. Green: explanation layer that decomposes each score into audit risk areas. Gray: the audit decision layer, where the ranked, explained output supports, rather than replaces, the auditor's risk assessment.

4. Methodology

4.1. The X-RBA Framework

Fig. 1 showsthe proposed explainable risk-based audit analytics (X-RBA) framework. It has four layers.

Data layer: Public financial statements and SEC AAER labels are assembled at the firm-year level as described in Section III.

Learning layer: Let $x_{it} \in \mathbb{R}^d$ be the feature vector of firm i in fiscal year t. X-RBA estimates the misstatement risk $\hat{p}_{it} = \Pr(y_{it} = 1 \mid x_{it})$ with a gradient-boosted decision tree (GBDT) ensemble [15, 16]. The model $F(x) = \sum_{m=1}^{M} \eta h_m(x)$ is an additive sum of M regression trees h_m with learning rate η, fitted stagewise to minimize the class-weighted logistic loss

$$\mathcal{L} = -\sum_{i,t} w_{y_{it}} [y_{it} \log \sigma(F(x_{it})) + (1 - y_{it}) \log(1 - \sigma(F(x_{it})))], \quad (1)$$

where $\sigma(\cdot)$ is the logistic function and the class weights $w_1 = N/(2N_1)$ and $w_0 = N/(2N_0)$ give the rare fraud class the same total weight as the non-fraud class. We use histogram-based GBDT with up to M = 400 trees and early stopping (training stops when the loss on a random 10% validation split of the training data fails to improve for 10 rounds), η = 0.03, at most 15 leaves per tree, a minimum of 40 samples per leaf, and an ℓ_2 penalty of 1.0. These settings were fixed before the test period was examined and were not tuned on test data. All models use a fixed random seed (0).

Table 2: Mapping of raw financial-statement items to audit risk areas

Audit risk area	Financial-statement items
Revenue & receivables	Sales; receivables
Inventory & cost of sales	Inventory; cost of goods sold
Profitability & earnings	Net income; income before extraordinary items; retained earnings
Financing & leverage	Liabilities (total, current); current and long-term debt; debt issuance; stock issuance; interest expense; preferred stock; accounts payable
Asset base & investments	Assets (total, current); cash; gross PP&E; investments; depreciation and amortization
Equity & market valuation	Common equity; common shares outstanding; year-end share price
Taxes	Income taxes payable; total income taxes

Explanation layer: For every firm-year in the top 1% of each test year (352 firm-years in total), we compute Shapley values φ_j of the model's log-odds output $F(x_i)$ for every feature j [13, 14]:

$$\phi_{ij} = \sum_{S \subseteq D \setminus \{j\}} \frac{|S|!\,(d - |S| - 1)!}{d!} \Big[v_i(S \cup \{j\}) - v_i(S) \Big], \quad (2)$$

where D is the feature set and $v_i(S)$ is the expected model output when the features in S take the firm's values and the others are drawn from a background sample of 500 training firm-years. The values are additive, so $F(x_i) = E[F] + \sum_j \varphi_{j,i}$. We estimate them with 150 random feature permutations per firm-year, which yields unbiased estimates [13]. Because auditors plan procedures by account area rather than by

individual line item, we sum the attributions within G = 7 predefined audit risk areas (Table II), $\Phi_{g,i} = \sum_{j \in g} \varphi_{j,i}$. Additivity guarantees that area contributions also sum to the firm's total deviation from the average risk.

Audit decision layer: Firms are ranked by ρ_i. Given a review budget of 1% of firm-years, the top-ranked firms are selected for heightened scrutiny, and their area-level risk profiles indicate where substantive testing and control evaluation should focus.

4.2. Baselines

We compare X-RBA against three models that represent the main streams of the literature:

- **Logit (ratios):** class-weighted logistic regression on the 14 standardized expert ratios, representing the F-score approach [6].
- **RUSBoost (raw):** our re-implementation of the model of Bao *et al.* [5]: AdaBoost.M1 with 500 decision trees (minimum leaf size 5), learning rate 0.1, and random undersampling of non-fraud firm-years to a 1:1 ratio at every boosting round [11].
- **Isolation forest (unsupervised):** an anomaly detector that uses no fraud labels [17], representing approaches that flag statistically unusual filings.

To isolate the effect of the feature set, we also train the X-RBA learner on the combined 42 raw and ratio features (X-RBA raw+ratios). All models were implemented in Python with scikit-learn [18].

4.3. Out-of-Sample Evaluation Protocol

We follow the rolling protocol of Bao *et al.* [5]. For each test year $t \in \{2003..., 2008\}$, models are trained on fiscal years 1991 to $t - 2$ and tested on year t. The two-year gap reflects the fact that frauds are typically revealed in AAERs years after they occur, so labels for year $t - 1$ would not be available to an auditor in year t. In addition, a fraud case can span several years. If a fraud case that appears in the test year also appears in the training period, its training-period fraud observations are re-coded as non-fraud so that the model cannot learn firm-specific patterns of the very case it is asked to detect [5]. The test period 2003–2008 matches that of [5]. We restrict the study to this window because AAER labels for later years are increasingly incomplete: frauds are often disclosed several years after they occur, so the most recent years in the dataset contain only a fraction of the frauds that will eventually be revealed.

4.4. Metrics

We report five metrics, each computed within a test year and then averaged across years.

- **AUC:** the probability that a randomly chosen fraud firm-year is ranked above a randomly chosen non-fraud firm-year.
- **Average precision (AP):** the area under the precision–recall curve, which is more informative than AUC under extreme imbalance.
- **NDCG@1%:** normalized discounted cumulative gain of the top 1% of firms [19], the ranking metric used in [5].
- **Capture@*k*% (sensitivity):** the share of all fraud firm-years in the test year that fall within the top *k*% of ranked firms, for $k \in \{1, 5\}$.
- **Precision@1%:** the share of firms in the top 1% that are actually fraudulent.

Capture@*k*% directly answers the audit budgeting question. Random selection would capture *k*% of frauds in expectation. To measure uncertainty, we use a firm-level bootstrap with 500 replications: within each test year, firm-years are resam-pled with replacement, the metric difference between two mod-els is recomputed, and the average difference across years is recorded.

5. Results

5.1. Detection Performance

Table III reports the out-of-sample performance of all mod-els averaged over the six test years. The proposed X-RBA model achieves the highest mean AUC (0.713), followed by the ratio-augmented X-RBA variant (0.684), the logistic ratio model (0.682), RUSBoost (0.675), and the unsupervised isola-tion forest (0.561). X-RBA also has the highest average preci-sion and the highest precision in the top 1% of ranked firms. Its top-1% precision of 2.26% is roughly three times the base fraud rate of 0.74% in the test period.

Table 3: Out-of-sample detection performance, test years 2003–2008

Model	AUC (s.d.)	AP (%)	NDCG @1%	Prec. @1% (%)	Capture @1%	Capture @5%
Isolation forest	0.561 (0.051)	1.00	0.002	0.28	0.2	5.9
Logit (ratios)	0.682 (0.027)	1.68	0.026	1.69	3.2	9.9
RUSBoost (raw) [5]	0.675 (0.049)	1.69	0.021	1.97	2.4	16.1
X-RBA (raw+ratios)	0.684 (0.035)	1.58	0.010	1.12	1.2	10.4
X-RBA (proposed, raw)	0.713 (0.036)	1.84	0.026	2.26	2.6	11.9
Random review	0.500	0.74	—	0.74	1.0	5.0

Note: Metrics are computed within each test year and averaged over 2003–2008 (s.d. = standard deviation of AUC across years). Capture@k% is the share of all fraud firm-years found when the top k% of firm-years are reviewed. Best value per column in bold.

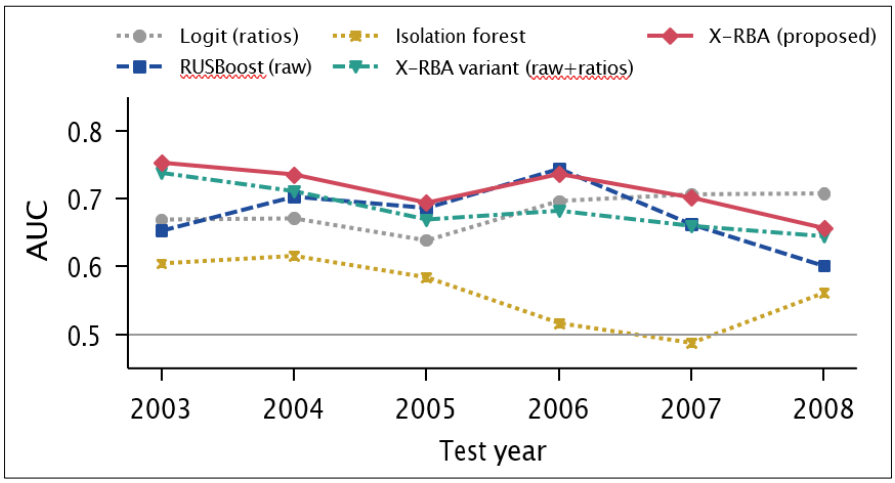


Fig 2: Out-of-sample AUC by test year. Each model is trained on fiscal years 1991 to t − 2 and tested on year t.

Fig. 2 shows that X-RBA ranks first in AUC in three of the six test years (2003–2005) and second in the other three. The tree-based models decline in 2008, the year of the financial crisis, while the logistic ratio model remains stable, suggesting that the ratio model is less sensitive to the sharp shifts in reported financials during that year. The isolation forest, which uses no fraud labels, is only modestly better than random: statistically unusual statements are not, by themselves, a reliable fraud sig-nal.

Table IV quantifies the uncertainty of these differences. The AUC gain of X-RBA over the RUSBoost re-implementation (+0.037) and over its own ratio-augmented variant (+0.029) is statistically significant, with bootstrap 95% confidence inter-vals that exclude zero. The gain over the logistic ratio model (+0.031) is borderline, and the gain over the isolation forest (+0.151) is large and significant.

5.2. Performance Under a Fixed Audit Budget

AUC summarizes ranking quality over the entire population, but an audit team can only review a small fraction of firms. Fig. 3 plots the share of frauds captured as the review budget grows from 0.25% to 20% of firm-years. All supervised models cap ture frauds at roughly two to three times the rate of random re-view across the range. Reviewing the 10% of firm-years ranked riskiest by X-RBA captures about one quarter of all frauds, and reviewing 20% captures about 44%.

Table 4: Difference between x-rba and each baseline (firm-level bootstrap, 500 replications, 2003–2008)

X-RBA minus	ΔAUC [95% CI]	ΔCapture@1%	ΔCapture@5%
RUSBoost (raw)	+0.037 [+0.009, +0.065]	±0.2 [−2.0, +3.1]	−4.2 [−9.9, +1.2]
Logit (ratios)	+0.031 [−0.004, +0.066]	−0.6 [−3.6, +2.4]	+2.0 [−4.4, +7.4]
X-RBA (raw+ratios)	+0.029 [+0.010, +0.049]	+1.4 [0.0, +3.4]	+1.5 [−2.6, +4.3]
Isolation forest	+0.151 [+0.109, +0.186]	+2.4 [−0.4, +4.2]	+6.1 [+2.0, +10.1]

Capture differences are in percentage points. Bold: 95% interval excludes zero.

The budget view also reveals a result that AUC alone hides. At the 5% budget, RUSBoost captures more frauds (16.1%) than X-RBA (11.9%), although the bootstrap interval for this difference includes zero (Table IV). The two models therefore have different strengths: X-RBA ranks the full population bet-ter and has the highest precision among the top 1% of firms, while RUSBoost, which aggressively rebalances the classes at every boosting round, matches or slightly exceeds X-RBA at budgets of about 3% and above (e.g., 26.9% versus 25.5% at a 10% budget). For practitioners, this means the choice of model should depend on the review budget, and the budget-based evaluation we propose should accompany AUC in any audit-analytics study.

Adding the 14 expert ratios to the raw items did not help the gradient-boosted learner: the raw-only model was significantly better in AUC. This supports the argument of [5] that flexi-ble learners extract the relevant signals from raw items, while derived ratios, with their missing values and outliers, can add noise.

5.3. Explaining the Risk Scores

Fig. 4 summarizes the Shapley attributions for the 352 firm-years flagged in the top 1% of their test year. At the level of audit risk areas (panel a), financing and leverage contribute the most to elevated risk scores, followed by equity and market val-uation. Profitability, revenue and receivables, and the asset base form a second tier, while inventory and taxes contribute little.

At the level of individual items (panel b), the sale of common and preferred stock is the single strongest signal, followed by shares outstanding, gross PP&E, current assets, and accounts payable.

The prominence of equity issuance agrees with prior evi-dence that firms raising external capital face stronger incentives to inflate reported performance [6], and with the inclusion of an issuance indicator in the F-score. The model learned this pattern from data alone, a useful sanity check on its behavior.

Fig. 5 illustrates a firm-level explanation for one correctly flagged case: a 2004 firm-year later identified in an SEC AAER and ranked 7th of 5,934 firms in its test year. Financing and leverage account for the largest share of the elevated score (+1.81 in log-odds), driven mainly by stock issuance, followed by equity and market valuation (+0.60). Revenue and receiv-ables, the asset base, and inventory add

small positive contribu-tions, while profitability slightly lowers the score. For an en-gagement team, this suggests scrutinizing the accounting for re-cent capital raises, testing current assets and payables, and eval-uating the related controls, rather than applying uniform proce-dures to all accounts.

6. Discussion

6.1. Implications for Risk-Based Auditing and Internal Control

The results support three practical uses of X-RBA. First, at the engagement-planning stage, the risk score offers an objective, data-driven input to the assessment of the risks of material mis-statement required by AS 2110 [2], complementing the audi-tor's understanding of the entity. Second, the area-level expla-nations indicate *where* risk is concentrated, so that substantive procedures and tests of controls can be directed to the relevant accounts. This supports the evaluation of internal control over financial reporting under SOX [1]: an area that drives an ele-vated score is a natural candidate for a closer look at control de-sign and operating effectiveness. Third, the budget-based eval-uation lets audit firms, internal audit functions, and regulators set review budgets with a realistic expectation of what each will capture. Because X-RBA uses only data that U.S. registrants already file with the SEC, a regulator can apply it to all pub-lic companies at negligible cost to prioritize filing reviews, with the explanations providing an auditable record of why each fil-ing was prioritized.

6.2. Limitations

This study has several limitations. First, AAER labels capture only detected and prosecuted fraud, so undetected frauds count as non-fraud and all precision and capture figures are conserva-tive. Second, absolute precision is low (about one flagged firm-year in 44 at the top 1%), so the scores support prioritization, not accusation, and must be combined with auditor judgment. Third, the test period is 2003–2008 and the public dataset ends in 2014; results should be verified on recent filings under newer standards such as ASC 606 once their labels mature. Fourth, our RUSBoost baseline uses 500 trees without the tuning of the original study and may understate its reported performance [5].

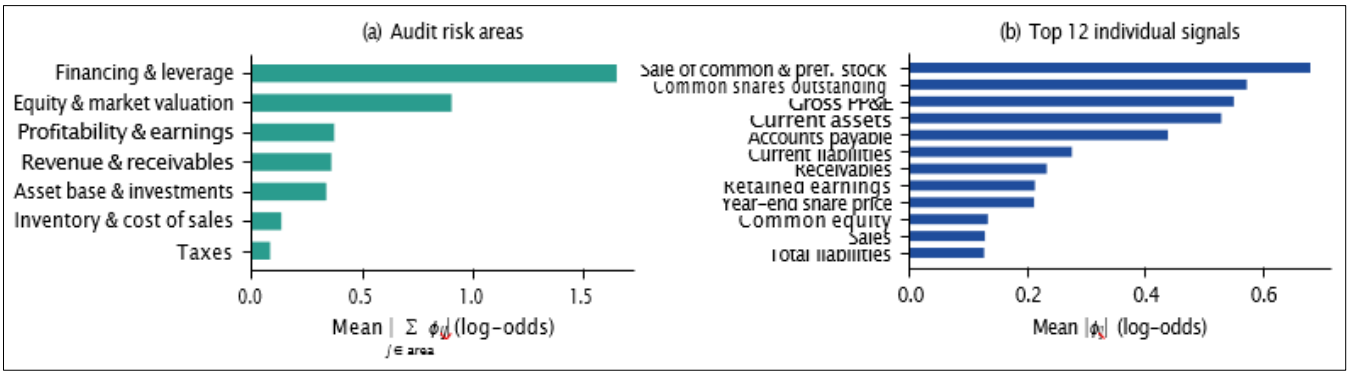


Fig 4: Global explanation of X-RBA for the 352 firm-years in the top 1% of risk scores (2003–2008). (a) Mean absolute Shapley contribution of each audit risk area to the log-odds risk score. (b) Mean absolute Shapley contribution of the 12 most influential financial-statement items.

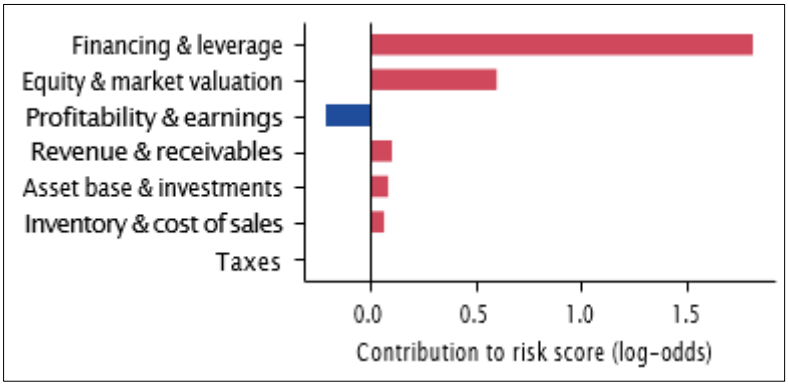


Fig 5: Firm-level explanation for a correctly flagged fraud firm-year (2004, rank 7 of 5,934). Bars show each audit risk area's Shapley contribution to the log-odds risk score; red increases risk and blue decreases it.

Finally, Shapley values explain the model, not the causes of fraud, and some items (e.g., shares outstanding) partly proxy for firm size.

7. Conclusion

This paper proposed and evaluated X-RBA, an explainable ma-chine learning framework for risk-based audit analytics. Using 146,045 public U.S. firm-years with SEC enforcement labels and a strict out-of-sample protocol, X-RBA achieved the best overall ranking of misstatement risk (mean AUC 0.713), sig-nificantly outperforming a re-implementation of the RUSBoost benchmark and a ratio-augmented variant, and roughly tripling the hit rate of random review among the riskiest 1% of firms. Budget-based evaluation showed that model rankings change with the review budget, so studies should report capture at re-alistic budgets alongside AUC. The Shapley-based explanation layer translated each risk score into audit risk areas, identifying equity issuance and

financing activity as the dominant drivers of elevated risk and showing how a flagged firm's profile can di-rect substantive procedures and internal-control evaluation. Fu-ture work will extend the framework to recent XBRL filings, add textual signals from management discussion and auditor re-ports, and evaluate it prospectively with practicing auditors.

References

1. Sarbanes–Oxley Act of 2002. Pub L No. 107-204, 116 Stat. 745. 2002.
2. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019;1(5):206-215.
3. Beneish MD. The detection of earnings manipulation. Financ Anal J. 1999;55(5):24-36.
4. Cecchini M, Aytug H, Koehler GJ, Pathak P. Detecting management fraud in public companies. Manage Sci. 2010;56(7):1146-1160.
5. Perols J. Financial statement fraud detection: An analysis of statistical and machine learning algorithms. Auditing. 2011;30(2):19-50.
6. Seiffert C, Khoshgoftaar TM, Van Hulse J, Napolitano A. RUSBoost: A hybrid approach to alleviating class imbalance. IEEE Trans Syst Man Cybern A Syst Hum. 2010;40(1):185-197.
7. Bertomeu J, Cheynel E, Floyd E, Pan W. Using machine learning to detect misstatements. Rev Account Stud. 2021;26(2):468-519.
8. Štrumbelj E, Kononenko I. Explaining prediction models and individual predictions with feature contributions. Knowl Inf Syst. 2014;41(3):647-665.
9. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems 30. 2017. p. 4765-4774.
10. Friedman JH. Greedy function approximation: A gradient boosting machine. Ann Stat. 2001;29(5):1189-1232.
11. Ke G, *et al*. LightGBM: A highly efficient gradient boosting decision tree. In: Advances in Neural Information Processing Systems 30. 2017. p. 3146-3154.
12. Liu FT, Ting KM, Zhou ZH. Isolation forest. In: Proceedings of the 8th IEEE International Conference on Data Mining (ICDM); 2008. p. 413-422.
13. Pedregosa F, *et al*. Scikit-learn: Machine learning in Python. J Mach Learn Res. 2011;12:2825-2830.
14. Järvelin K, Kekäläinen J. Cumulated gain-based evaluation of IR techniques. ACM Trans Inf Syst. 2002;20(4):422-446.

How to Cite This Article

Fuadedzi P. Advancing data-driven, risk-based audit analytics to strengthen fraud detection, internal controls, financial-reporting reliability, and regulatory compliance across U.S. organizations. International Journal of Management and Organizational Research. 2026;5(6):18–24.

Creative Commons (CC) License

This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.