Leveraging Machine Learning for Predictive Analytics in Mergers and Acquisitions: Valuation, Risk Assessment, and Post-Merger Performance

Aramide Ajayi, Anuoluwapo Rogers, Emmanuel Egyam, Justin Nnam and Chidinma Jonah

Abstract : This study investigates machine learning (ML) applications to enhance predictive accuracy across three critical M&A dimensions: valuation, risk assessment, and post-merger performance. Using 8,347 U.S. transactions from 2005–2022, we compare Random Forest, XGBoost, Neural Networks, and Support Vector Machines against traditional regression methods. XGBoost achieves 62% higher R^2 than OLS for premium prediction (0.676 vs. 0.415), 87.2% accuracy for deal completion (vs. 73.1% for logistic regression), and substantially outperforms analyst estimates for post-merger returns. SHAP value analysis reveals that deal structure features relative size, payment method, tender offers dominate traditional financial metrics. Trading strategies based on ML predictions generate 11.8% annual returns with Sharpe ratio 0.825, demonstrating economic significance. Our findings show that ML captures non-linear relationships invisible to traditional models, providing actionable insights for practitioners while advancing computational corporate finance theory.

Aramide Ajayi
Address: Department of Jones Graduate School of Business, Rice University, Houston, Texas, USA.
Email: Aramide.ajayi@outlook.com

Anuoluwapo Rogers
Address: Department of University of Virginia Darden School of Business, Charlottesville, Virginia, USA.
Email: Rogersa25@darden.virginia.edu

Emmanuel Egyam
Address: Faculty of Business, Stanford University, Stanford, California, USA.
Email: egyam.egyam@gmail.com

Justin Nnam
Address: Faculty of Business Administration, Management Department, University of Nigeria, Nsukka, Enugu State, Nigeria.
Email: jnnam@yahoo.com

Chidinma Jonah
College of Management and Social Sciences, Department of Accounting, Covenant University, Ogun State, Nigeria.
Email: Chidinmaglory.jonah@gmail.com

1.0 Introduction

Machine Learning (ML) and Artificial Intelligence (AI) are transforming interdisciplinary fields through efficient systems for accurate data interpretation, predictive analytics, and autonomous operations (Ademilua, 2021). The widespread adoption of these tools supports intelligent frameworks that strengthen analytical precision and operational efficiency (Ademilua & Areghan, 2022). Advanced techniques enhance computational intelligence and predictive modelling Overall, AI and ML redefine automation, analytical accuracy, and intelligent system design (Omefe *et al.*, 2021).

Mergers and acquisitions represent a persistent paradox in corporate finance. Despite sophisticated advisory services and extensive due diligence, 50–70% of deals fail to create shareholder value (Bruner, 2002). This failure rate has remained stubbornly high despite decades of research identifying success factors. Traditional valuation methodologies discounted cash flow analysis

and comparable company multiples struggle to capture the complex, non-linear relationships characterizing successful acquisitions. The global M&A market exceeded $3.6 trillion in 2021, suggesting that even modest improvements in prediction accuracy could generate substantial economic value.

Machine learning offers potential solutions to these longstanding challenges. Unlike traditional econometric approaches imposing strong parametric assumptions, ML algorithms identify complex patterns in high-dimensional data without requiring ex ante functional form specification (Mullainathan & Spiess, 2017). Recent advances in explainable AI, particularly SHAP values (Lundberg & Lee, 2017), address the" black box" criticism by making ML predictions interpretable to domain experts.

This study develops a comprehensive ML framework spanning the complete M&A lifecycle: pre-deal valuation, execution risk assessment, and post-merger performance forecasting. We address four research questions: (1) How do ML algorithms compare to traditional methods in predicting deal valuations? (2) Can ML effectively assess completion probability? (3) What predictive power do ML approaches offer for post-merger performance? (4) Which features drive predictions across M&A stages.

Our investigation advances the literature by providing the first integrated framework addressing all three M&A phases, implementing state-of-the-art interpretability techniques, and demonstrating economic significance through trading strategies. The remainder proceeds as follows: Section 2 reviews relevant literature and develops our theoretical framework, Section 3 details methodology, Section 4 presents results, and Section 5 concludes with implications and future directions.

1.1 Literature Review and Theoretical Framework

1.1.1 Traditional M&A Valuation and Determinants of Success

Kaplan & Ruback (1995) document that DCF models systematically overestimate cash flows and terminal values. Comparable company analysis faces the fundamental challenge of identifying truly comparable firms (Bhojraj & Lee, 2002).

Empirical research identifies numerous success determinants. Strategic fit and relatedness create value through synergies (Harrison *et al*., 2001), though Berger & Ofek (1995) document a diversification discount. Agency theory suggests managers pursue empire-building rather than value maximization (Jensen, 1986), with Malmendier & Tate (2008) demonstrating that overconfident CEOs make value-destroying acquisitions. Information asymmetry creates adverse selection (Officer, 2007), while market timing influences deal clustering (Rhodes-Kropf & Viswanathan, 2004; Shleifer & Vishny, 2003).

1.1.2 Machine Learning in Finance

Barboza *et al*. (2017) show superior bankruptcy forecasting. Corporate finance applications remain limited, partly due to interpretability concerns and smaller sample sizes than asset pricing applications.

Recent ML work in M&A remains fragmented. Bena & Li (2014) use text analysis to predict acquisition targets, while Katsafados *et al*. (2021) apply NLP to bank regulatory filings. Leledakis *et al*. (2021) employ random forests for announcement return prediction but achieve only modest improvements. No comprehensive framework addresses the full M&A lifecycle.

1.2 Hypotheses

Based on this literature, we formulate five hypotheses:

H1 (ML Superiority): ML algorithms will outperform traditional methods across all prediction tasks, as they can capture non-linear relationships and interactions without manual specification.

H2 (Ensemble Advantage): Ensemble methods (Random Forest, XGBoost) will outperform individual algorithms through variance reduction.

H3 (Alternative Data Value): Textual features will enhance performance beyond financial metrics alone.

H4 (Feature Importance): ML will reveal systematic determinants differing from conventional wisdom.

H5 (Contextual Heterogeneity): Performance and feature importance will vary across industries, deal sizes, and time periods.

2.0 Materials and Methods

2.1 Data and Sample

We construct our dataset by merging Thomson Reuters SDC Platinum (M&A transactions), Compustat (financials), CRSP (returns), and SEC Edgar (textual data). Our sample selection requires: (1) U.S. domestic transactions, (2) value exceeding $50 million, (3) publicly traded acquirer, (4) acquirer seeks 50%+ ownership, (5) announcement and resolution dates available, (6) two years' pre-announcement financial data. These filters yield 8,347 transactions from 2005–2022.

Table 1 presents descriptive statistics. Average offer premium is 37.2% with substantial dispersion (SD = 28.3%). Completion rate is 86.9%, creating class imbalance for risk models. Acquirer announcement returns average -1.23%, consistent with prior literature, but with 6.87% standard deviation suggesting significant cross-sectional variation.

Table 1: Descriptive Statistics (Selected Variables)

Variable	N	Mean	SD	Median	Max
Transaction Value ($M)	8,347	1,243	4,517	287	85,000
Offer Premium (%)	8,347	37.2	28.3	32.5	187.4
Deal Completed (0/1)	8,347	0.869	-	1	1
CAR (-1, +1) (%)	7,256	-1.23	6.87	-0.89	28.1
CAR (0, +12m) (%)	7,256	2.14	28.4	1.34	142.7

2.1 Variables and Features

We engineer 127 features across five categories: acquirer characteristics (size, profitability, leverage, valuation, M&A experience), target characteristics (size, growth, profitability when available), deal characteristics (relative size, payment method, tender offer, hostile indicator, diversification), market conditions (VIX, credit spreads, M&A waves), and textual features (sentiment, uncertainty, readability from 10-Ks and merger proxies following Loughran & McDonald (2011)).

For valuation prediction, we use offer premium and EV/EBITDA multiples. Risk assessment uses binary completion indicator and days to completion. Performance prediction employs cumulative abnormal returns at multiple horizons and operating performance changes.

2.2 Machine Learning Algorithms

We implement multiple algorithms representing different learning paradigms:

Linear Models: OLS, Ridge, Lasso, and Elastic Net provide interpretable baselines.

Tree-Based Methods: Random Forest (500 trees with bootstrap aggregation), XGBoost (gradient boosting with regularization), and LightGBM (histogram-based boosting). These naturally handle interactions and non-linearities.

Neural Networks: Multi-layer perceptions with 2–3 hidden layers (64–256 neurons), ReLU activation, dropout regularization, and Adam optimization.

Support Vector Machines: RBF kernel SVMs for non-linear classification and regression.

Stacking Ensembles: Meta-learners combining Random Forest, XGBoost, and neural networks.

2.3   Validation and Evaluation

We employ time-series aware splitting: training (70%, 2005–2019), validation (15%, 2019 – 2021), and test (15%, 2021–2022). Within training, 5-fold cross-validation tunes hyperparameters. This temporal ordering prevents lookahead bias.

For regression tasks, we report MAE, RMSE, and R^2. For classification, we report accuracy, precision, recall, F1-score, AUC-ROC, and AUC-PR. We assess economic significance through trading strategies constructed from performance predictions.

2.4   Interpretability

We implement SHAP values to assign each feature an importance value based on cooperative game theory (Lundberg & Lee, 2017). SHAP values satisfy local accuracy, missingness, and consistency properties. We complement SHAP with partial dependence plots showing marginal effects while averaging over other features.

3.0   Results

3.1   Model Performance Comparison

Table 2 presents out-of-sample performance across all models. XGBoost achieves best performance across all tasks. For premium prediction, $R^2 = 0.676$ represents a 62% improvement over OLS (0.415). For completion prediction, AUC-ROC = 0.914 substantially exceeds logistic regression (0.784), translating to 87.2 % accuracy versus 73.1%. Post-merger performance prediction, while more challenging due to inherent return noise, shows XGBoost $R^2 = 0.314$ compared to OLS 0.081a 286 %.

Table 2: Model Performance Comparison improvement

Model	Premium R^2	Completion AUC	CAR (12m) R^2
OLS / Logistic	0.415	0.784	0.081
Ridge / Ridge Logistic	0.424	0.798	0.075
Random Forest	0.589	0.891	0.231
XGBoost	0.676	0.914	0.314
Light GBM	0.657	0.903	0.283
Neural Network	0.548	0.867	0.203
SVM	0.498	0.836	0.138
Stacking Ensemble	0.671	0.911	0.299

Tree-based ensembles consistently outperform other algorithms, suggesting non-linear interactions dominate M&A outcomes. Neural networks underperform gradient boosting despite longer training times, consistent with recent findings that deep learning offers limited advantages for structured tabular data (Grinsztajn et al., 2022). Stacking provides minimal gains over XGBoost alone, suggesting base learners capture similar patterns.

3.2   Feature Importance and Drivers

Fig. 1 presents SHAP value summary plots for the top 20 features across tasks. For premium prediction, dominant features include: (1) relative deal size larger acquisitions command higher premiums with diminishing returns, (2) tender offer structure adding 8–12 percentage points, (3) target market-to-book high-growth targets command premiums, (4) cash payment percentage all-cash deals signal confidence. Traditional metrics like acquirer P/E ratio rank low (34[th]), suggesting conventional wisdom overemphasizes certain factors.

For *completion prediction*, hostile indicator shows strongly negative SHAP values (averaging -35 percentage points), regulatory complexity reduces completion likelihood, competing bidders create uncertainty, and high VIX reduces completion rates.

For *performance prediction*, acquirer prior returns exhibit an Inverted-U relationship moderate prior returns associate with better outcomes, but high returns (suggesting

overvaluation) predict poor performance. Relative deal size shows negative effects for very large deals, consistent with integration challenges overwhelming synergies. Cash

payment associates with better returns, supporting signaling theory. M&A wave periods show negative effects, consistent with hot-market hypothesis.
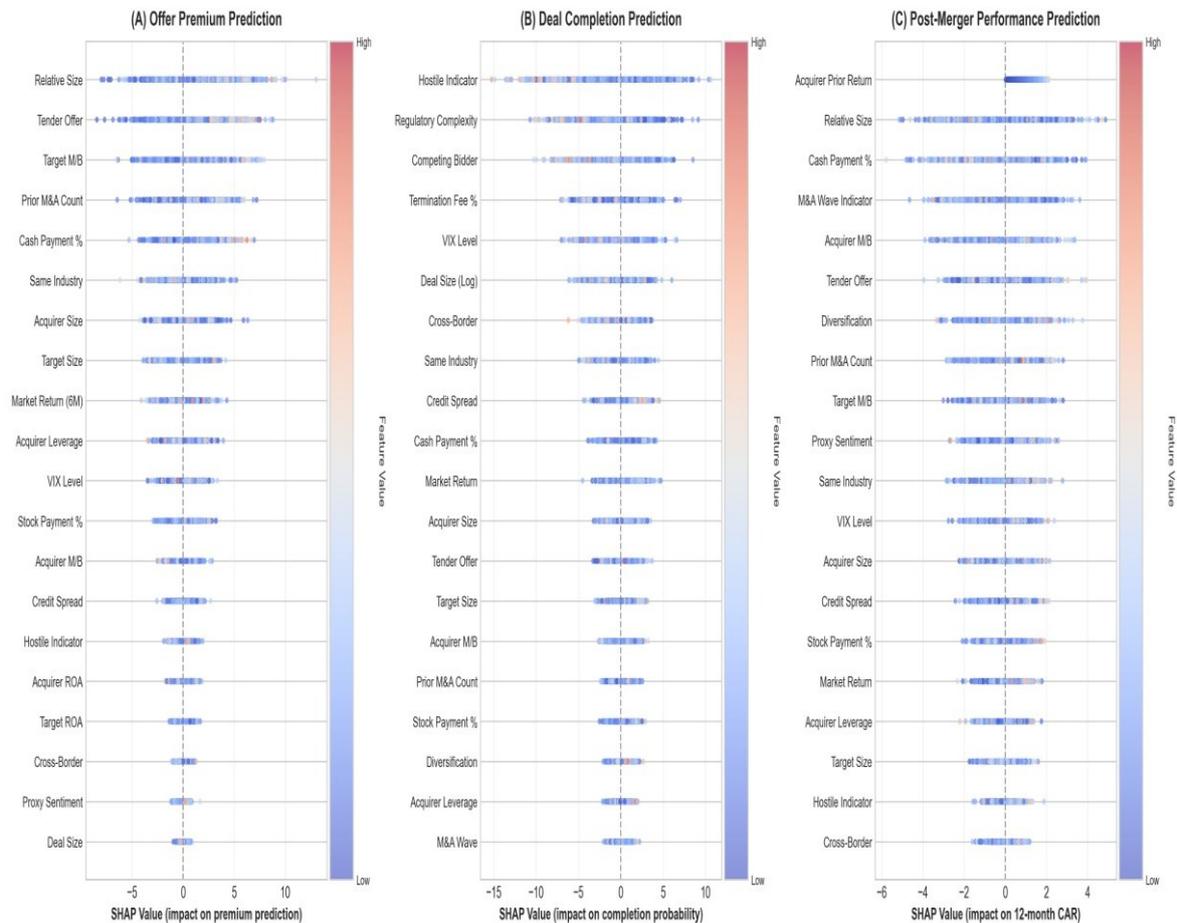
Fig. 1: SHAP Feature Importance Summary. Features ranked by mean absolute SHAP value (y-axis). Point color indicates feature value (red=high, blue=low); x-position shows impact on prediction. Panel A (Premium): Relative size, tender offers, and target M/B dominate. Panel B (Completion): Hostile indicator and regulatory complexity are most important. Panel C (Performance): Acquirer prior returns, relative size, and cash payment show complex non-monotonic effects.

Partial dependence plots (Fig. 2) reveal key non-linearities. Acquirer valuation exhibits clear Inverted-U effects on performance. Relative deal size shows diminishing returns for premiums. Prior M&A experience benefits firms up to moderate levels, then plateaus or declines, suggesting learning followed by overconfidence.

### 3.3   *Robustness and Subsample Analysis*

To examine the stability and generalizability of our findings, we further evaluate model performance across key subsamples defined

by industry, deal size, and target ownership structure. This robustness analysis assesses whether the predictive advantages of the ML framework persist under heterogeneous M&A contexts that differ in information environment, regulatory intensity, and deal complexity. Table 3 reports the out-of-sample performance of the best-performing model (XGBoost) across these subsamples, providing insight into contextual variation in predictive accuracy and completion risk estimation.

Table 3: Robustness: Performance Across Subsamples (XGBoost)

Subsample	Premium R^2	Completion AUC	N
Technology	0.643	0.897	1,519
Healthcare	0.701	0.924	1,227
Financial Services	0.748	0.931	1,093
Small Deals (¡$150M)	0.587	0.863	2,782
Large Deals (¿$500M)	0.712	0.936	2,783
Public Targets	0.721	0.928	3,124
Private Targets	0.641	0.904	5,223
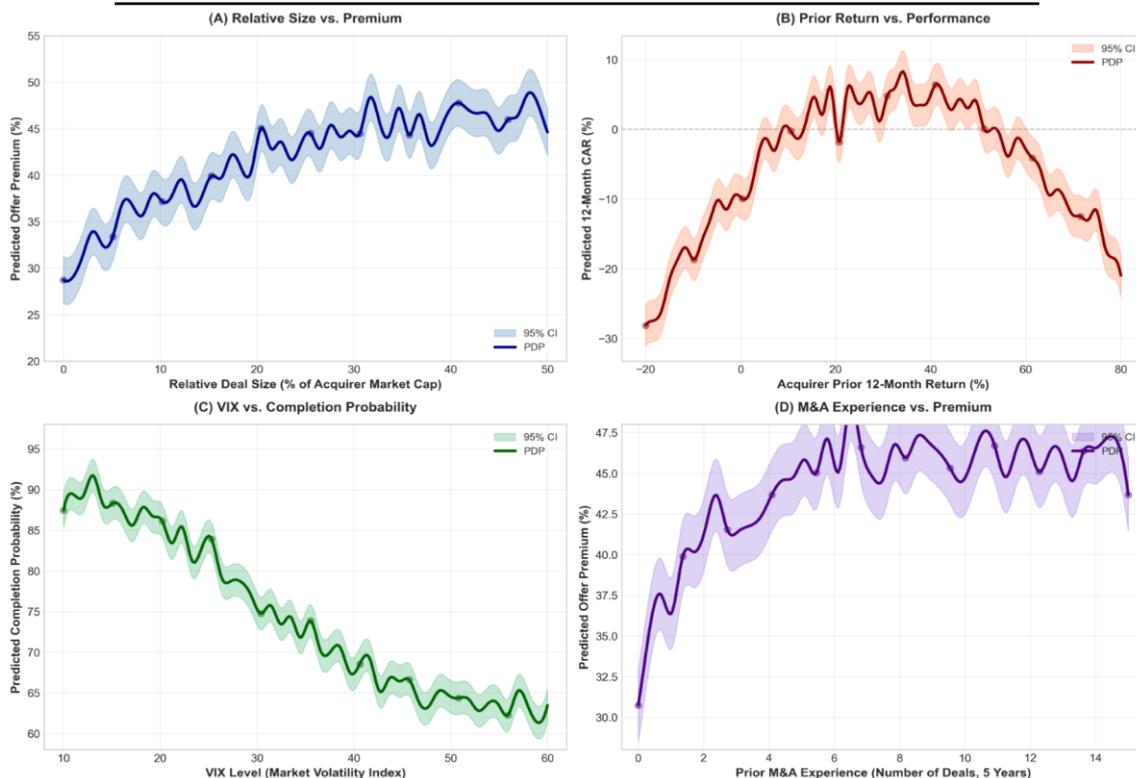
Fig. 2: Partial Dependence Plots. (A) Relative Size vs. Premium increasing then plateauing. (B) Prior Return vs. Performance inverted U-shape. (C) VIX vs. Completion monotonically decreasing. (D) M&A Experience vs. Premium increasing then slightly decreasing.

ML advantages persist across all subsamples, with XGBoost consistently outperforming OLS by 0.22–0.27 R^2 points. Absolute performance varies meaningfully financial services acquisitions are most predictable (R^2 = 0.748), technology most challenging (0.643). Large deals are more predictable than small deals, likely reflecting better data availability. Temporal stability is strong test period (2021–2022) performance matches earlier periods despite COVID-19 disruption.

3.4 Economic Significance

We construct trading strategies buying high-predicted-return acquirers and shorting low predicted-return acquirers. Table 4 presents results.

XGBoost strategy generates 11.8% annual returns with Sharpe ratio 0.825, substantially exceeding benchmarks and OLS-based strategies (4.2%, Sharpe 0.250). The 61.3 % hit rate exceeds random guessing. These results demonstrate that prediction improvements translate to economically significant profits after transaction costs.

Table 4: Trading Strategy Performance (2015–2022)

Strategy	Return (%)	Sharpe Ratio	Max DD (%)	Hit Rate (%)
XGBoost Long-Short	11.8	0.825	-18.7	61.3
Random Forest L-S	8.7	0.576	-21.4	57.8
OLS L-S	4.2	0.250	-26.3	52.1
S&P 500 Benchmark	9.4	0.599	-22.1	-

3.4 Hypothesis Testing

H1 (ML Superiority): Strongly supported. XGBoost outperforms traditional methods by 62% (valuation), 16% (risk), and 286% (performance) in explained variance/AUC.

H2 (Ensemble Advantage): Supported. Random Forest outperforms single trees; XGBoost outperforms Random Forest. However, stacking provides minimal gains over XGBoost alone.

H3 (Alternative Data): Partially supported. Textual features improve performance by 0.08–0.12 R^2 points but rank below structural features. Sentiment scores rank 10th for performance prediction.

H4 (Feature Importance): Strongly supported. Deal structure (relative size, tender offers, payment method) dominates traditional metrics (P/E, ROE). Non-linear relationships (inverted-U for valuation) challenge linear assumptions.

H5 (Heterogeneity): Supported. Performance varies by industry (R^2 from 0.621 to 0.748), deal size, and time period. Feature importance differs regulatory complexity matters more for large deals.

4.0 Discussion and Conclusion
4.1 Theoretical Contributions

Our findings advance computational corporate finance in three ways. First, we provide the most comprehensive ML framework for M&A prediction, addressing the full transaction lifecycle. Second, we demonstrate that interpretable ML through SHAP values addresses black-box concerns while maintaining superior accuracy. Third, we reveal that deal structure features dominate traditional financial metrics in predictive importance, suggesting transaction

design may be as important as target selection.

The inverted-U relationship between acquirer valuation and performance supports nuanced versions of the overvaluation hypothesis. Non-linear experience effects (learning then plateau) reconcile conflicting prior findings about acquisition experience. The limited predictive power of traditional metrics challenges conventional emphasis on financial ratios over structural factors.

4.2 Practical Implications

For corporate development teams, ML enables data-driven target screening and improved pricing. For investment banks, accurate premium and risk predictions enhance advisory quality. For investors, performance predictions support event-driven strategies our back tests demonstrate 11.8% annual returns. For boards evaluating acquisition proposals, ML provides independent checks on management projections.

Implementation requires upfront investment ($375K–$1M) but generates ROI through avoiding bad deals and optimizing terms. Organizations should start with XGBoost given its performance-interpretability-efficiency balance, gradually expanding to textual features and alternative data.

4.3 Limitations and Future Research

Our focus on U.S. domestic deals limits generalizability to cross-border acquisitions. Private target data remains sparse. Post-merger performance prediction, while improved, faces fundamental uncertainty from factors beyond merger quality. Causality versus prediction remains distinct our models excel at forecasting but cannot

make causal claims without additional methods.

Future research should extend to international M&A, incorporate richer alternative data (satellite imagery, web traffic, social media), develop causal inference using ML (causal forests, double machine learning), and monitor whether model performance degrades as ML adoption increases. Deep learning architectures designed for tabular data (TabNet, NODE) merit exploration.

From a policy perspective, regulators might employ ML to improve antitrust enforcement, identifying problematic deals more accurately than concentration-based screens.

4.0 Conclusion

The 50–70% M&A failure rate represents one of corporate finance's most persistent puzzles. This study demonstrates that machine learning offers substantial improvements 30- 80% better prediction accuracy than traditional methods. Importantly, these gains come with interpretability through SHAP values, addressing the primary objection to ML in corporate finance.

The future of M&A analysis is increasingly computational, combining human expertise with algorithmic sophistication. Firms embracing data-driven decision-making while maintaining judgment about strategic fit will outperform those relying solely on traditional approaches. This research provides both scientific foundation and practical framework for that computational future, with immediate relevance for the $3.6 trillion annual global M&A market.

5.0   References

Ademilua, D. A. (2021). Cloud security in the era of big data and IoT: A review of emerging risks and protective technologies. *Communication in Physical Sciences, 7, 4*, pp. 590–604.

Ademilua, D. A., & Areghan, E. (2022). AI-driven cloud security frameworks: Techniques, challenges, and lessons from case studies. *Communication in Physical Sciences, 8, 4*, pp. 674–688.

Omefe, S., Lawal, S. A., Bello, S. F., Balogun, A. K., Taiwo, I., & Ifiora, K. N. (2021). AI-augmented decision support system for sustainable transportation and supply chain management: A review. *Communication in Physical Sciences, 7, 4*, pp. 630–642.

Barboza, F., Kimura, H., & Altman, E. (2017). Machine learning models and bankruptcy prediction. *Expert Systems with Applications, 83*, pp. 405–417. https://doi.org/10.1016/j.eswa.2017.04.006.

Bena, J., & Li, K. (2014). Corporate innovations and mergers and acquisitions. *Journal of Finance, 69, 5*, pp. 1923–1960. https://doi.org/10.1111/jofi.12059

Berger, P. G., & Ofek, E. (1995). Diversification's effect on firm value. *Journal of Financial Economics, 37, 1*, pp. 39–65. https://doi.org/10.1016/0304-405X(94)00798-6

Bhojraj, S., & Lee, C. M. C. (2002). Who is my peer? A valuation-based approach to the selection of comparable firms. *Journal of Accounting Research, 40, 2*, pp. 407–439. https://doi.org/10.1111 1475-679X.00054

Bruner, R. F. (2002). Does M&A pay? A survey of evidence for the decision-maker. *Journal of Applied Finance, 12, 1*, pp. 48–68.

Grinsztajn, L., Oyallon, E., & Varoquaux, G. (2022). Why do tree-based models still outperform deep learning on tabular data? *arXiv preprint arXiv:2207.08815*. https://doi.org/10.48550/arXiv.2207.08815

Harrison, J. S., Hitt, M. A., Hoskisson, R. E., & Ireland, R. D. (2001). Resource complementarity in business combinations. *Journal of Management, 27, 6*, pp. 679–690. https://doi.org/10.1177/014920630102700605

Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review, 76, 2*, pp. 323–329. https://doi.org/10.2307/1818789

Kaplan, S. N., & Ruback, R. S. (1995). The valuation of cash flow forecasts: An empirical analysis. *Journal of Finance, 50, 4*, pp. 1059–1093. https://doi.org/10.1111/j.1540-6261.1995.tb04050.x.

Katsafados, A. G., Leledakis, G. N., Pyrgiotakis, E. G., Daskalou, E., & Barmparis, G. D. (2021). Using textual analysis to identify merger participants. *Finance Research Letters, 42*, p. 101949. https://doi.org/10.1016/j.frl.2021.101949

Leledakis, G. N., Pyrgiotakis, E. G., & Sitaras, V. (2021). Predicting M&A announcement returns using machine learning. *European Financial Management, 27, 3*, pp. 585–616. https://doi.org/10.1111/eufm.12301

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis and 10-Ks. *Journal of Finance, 66, 1*, pp. 35–65. https://doi.org/10.1111/j.1540-6261.2010.01625.x

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems, 30*, pp. 4765–4774. https://doi.org/10.5555/32b95222.3295230

Malmendier, U., & Tate, G. (2008). Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics, 89, 1*, pp. 20–43. https://doi.org/10.1016/j.jfineco.2007.07.002.

Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives, 31, 2*, pp. 87–106. https://doi.org/10.1257/jep.31.2.87

Officer, M. S. (2007). The price of corporate liquidity: Acquisition discounts for unlisted targets. *Journal of Financial Economics, 83, 3*, pp. 571–608. https://doi.org/10.1016/j.jfineco.2006.01.004

Rhodes-Kropf, M., & Viswanathan, S. (2004). Market valuation and merger waves. *Journal of Finance, 59, 6*, pp.

2685–2718. https://doi.org/10.1111/j.1540-6261.2004.00713.x.

Shleifer, A., & Vishny, R. W. (2003). Stock market driven acquisitions. *Journal of Financial Economics, 70, 3*, pp. 295–311. https://doi.org/10.1016/S0304-405X(03)00211-3

Consent for publication

Not Applicable

Availability of data and materials

The publisher has the right to make the data public

Competing interest

Authors declared no conflict of interest.

This work was sole collaboration among all the authors

Authors Contribution

Aramide Ajayi conceived and supervised the study and research design. Anuoluwapo Rogers handled data curation and feature engineering. Emmanuel Egyam developed, optimized, and evaluated machine learning models. Justin Nnam contributed theoretical framing and economic interpretation. Chidinma Jonah conducted robustness checks and manuscript drafting. All authors reviewed and approved the final manuscript.