Jurnal Genesis Indonesia
Vol. 5, No. 03, pp. 273-286
journal.iistr.org/index.php/JGI
DOI: 10.56741/IISTR.jgi.002762

AI-Driven Financial Forecasting in Corporate Finance

A Comparative Assessment of Machine Learning Models and Implications for Financial Decision-Making

[1]T. Venkatesan*, [1]A. Stephen[2], [1]P. Elakkiya, [2]Maria Joy C. Iglesia, [2]Roque B. Cruz II

Corresponding Author: * proftvv@gmail.com

[1] Knowledge Institute of Technology, Salem, Tamil Nadu, India

[2] Camarines Sur Polytechnic Colleges, Nabua, Philippines

ARTICLE INFO

Article history
Received 6 October 2026
Revised 13 November 2026
Accepted 16 November 2026

ABSTRACT

Financial forecasts underpin liquidity management, capital budgeting, operating plans, and risk assessment, yet conventional specifications may be challenged by nonlinear relationships and changing economic conditions. This study examines the relative predictive performance of linear regression and three machine learning approaches, namely random forest, extreme gradient boosting (XGBoost), and long short-term memory (LSTM), using a corporate-finance dataset described in the source study that covers 2015–2025. The reported design employs financial statement and market variables, preprocessing, a 70:30 training–testing allocation, and evaluation based on prediction accuracy, root mean squared error (RMSE), and coefficient of determination (R^2). The reported results indicate that XGBoost has the highest accuracy (0.90), lowest RMSE (0.11), and highest R^2 (0.88), compared with linear regression (0.72, 0.21, and 0.65, respectively). Random forest and LSTM also exceed the reported linear-regression performance. These findings support the potential usefulness of nonlinear algorithms for the financial variables considered, although the information available does not permit independent verification, statistical significance testing, or causal inference about managerial outcomes. The article connects the comparative results to forecasting theory, methodological safeguards against temporal leakage, model explainability, and the organizational conditions under which predictive improvements might inform financial decisions. It offers an evaluation framework for future transparent, reproducible

firm-level studies rather than asserting that accuracy gains alone guarantee improved financial performance.

Keywords
Artificial Intelligence
Corporate Finance
Financial Forecasting
Random Forest
XGBoost

This is an open-access article under the CC–BY-SA license.

Introduction

Corporate financial planning is inherently forward-looking. Managers must estimate revenue, operating costs, cash requirements, and exposures before allocating scarce capital to investments or operating activities. The quality of these projections affects not only a budget's numerical precision but also the timing of financing, the availability of liquidity buffers, the evaluation of projects, and the communication of uncertainty to stakeholders. Traditional time-series and regression approaches remain useful because they make assumptions relatively explicit and often perform well when patterns are stable [1]–[3]. Nevertheless, financial relationships may be nonlinear, heterogeneous across firms, and subject to shocks. A forecasting architecture therefore needs to balance predictive power, interpretability, data requirements, and operational feasibility.

Financial forecasting differs from generic prediction in several important respects. First, forecasts are produced at a particular information date; models should not rely on data that would only become available later. Second, the relevant horizon varies by decision. Daily market-risk monitoring, quarterly revenue guidance, and multi-year capital budgeting require different targets and validation designs. Third, prediction errors often have asymmetric costs: an underestimated funding shortfall can be more damaging than an equivalent overestimate of available cash. Finally, superior statistical fit is not sufficient evidence of improved decision quality unless forecasts are connected to explicit management actions [4]–[6].

Machine learning offers flexible tools for learning interactions and nonlinear associations from financial data. Random forests average predictions across decision trees to reduce variance [7], [8], whereas gradient boosting progressively fits learners to improve an objective function [9]. XGBoost extends boosted trees with regularization and efficient optimization [10]. LSTM networks incorporate gated memory structures that can represent temporal dependencies [11]. These models have attracted considerable interest in asset pricing, financial time-series analysis, and computational forecasting [12]–[16]. Their relative merits, however, are data- and task-specific, and strong performance reported for one setting should not be generalized automatically to every corporate forecasting problem.

The present study focuses on the corporate-finance interpretation of a reported comparison among linear regression, random forest, XGBoost, and LSTM. The source manuscript describes secondary financial observations spanning 2015–2025 and presents three summary tables containing descriptive statistics, model metrics, and a qualitative comparison of forecasting methods. Its main numerical finding is that XGBoost ranks first on three reported performance measures. The source does not identify the number of firms, observations, exact data providers, target-specific evaluation procedures, or hyperparameter settings. This revised article retains its reported results without creating additional empirical observations, while making the corresponding limitations explicit.

The contribution is therefore twofold. Empirically, it organizes the original reported comparisons in a coherent evaluation narrative. Methodologically, it distinguishes what these aggregate results establish from what remains untested: calibration, uncertainty, reproducibility, temporal generalization, and realized managerial benefits. Specifically, the article addresses three research questions. How do the reported predictive metrics compare across conventional and machine learning specifications? What can reasonably be inferred about applications in capital budgeting, working-capital management, and risk assessment? What design and governance requirements must be satisfied before the findings can support stronger claims about deployment?

Materials and methods

A. Research design and empirical scope

The source study describes a quantitative comparative forecasting design based on secondary corporate financial information. Its stated observation window is 2015–2025, and the intended setting includes firms from multiple industries. Data categories comprise corporate financial statements, stock market records, and publicly available financial datasets. Revenue, operating profit, cash flow, and stock-price-related measures are reported as relevant variables. The specific sampling frame, firm identifiers, country coverage, frequency, number of records, and database names are not stated in the supplied material; accordingly, these characteristics cannot be reconstructed responsibly and no new sample composition is claimed here.

A useful distinction is between prediction targets and explanatory features. A firm-level forecast may target next-quarter revenue growth, cash flow measured in monetary units, an operating-profit ratio, or a stock return at a stated horizon. Variables observed at the forecast origin may then serve as predictors. Each target needs its own definition, units, horizon, and baseline because an error of 0.11 can have different interpretations for standardized ratios and currency-denominated cash flows. The source presents model-level metrics but does not specify whether these correspond to one target, an average across targets, or a transformed

response. This article therefore reproduces the metrics as reported without assigning them a previously unspecified unit or outcome.

B. Data processing and feature construction

The original methodology lists duplicate removal, treatment of inconsistent entries, mean imputation or interpolation for missing values, normalization, and feature selection. These procedures are plausible preparatory steps but require careful sequencing in a forecasting exercise. To avoid leakage, any fitted imputation parameters, feature-selection rules, and scaling constants should be estimated only within each training window and applied unchanged to its validation or test observations [17], [18]. Interpolation should not use values from dates later than the relevant forecast origin. Similarly, financial statement values should be assigned by public release date rather than fiscal-period end date if the model is intended to reproduce an investor's or manager's real information set.

Potential predictors include lagged revenue growth, operating margins, lagged cash-flow ratios, leverage, working capital, volatility, and macroeconomic conditions, provided their timestamps and measurement conventions are known. Revenue growth and stock returns are conceptually different from cash-flow levels and require distinct treatment of zero values, outliers, and scale. Where multiple firms appear in the panel, entity identification and industry affiliation can improve specification, but training and evaluation should reflect whether the objective is forecasting existing firms or generalizing to unseen companies. The source does not report a final feature list or missingness diagnostics, so the discussion above sets out reproducibility requirements rather than additional steps verified to have occurred.

C. Forecasting models and comparison logic

Linear regression is the conventional comparator. With predictors x and target y, the conditional mean is specified as $E(y|x)=\beta_0+\Sigma\beta_j x_j$, with coefficients estimated from training data. Its value lies in parsimony, interpretability, and usefulness as a benchmark. A limited linear specification is not equivalent to the full range of statistical forecasting alternatives: autoregressive models, exponential smoothing, and dynamic regression can accommodate temporal structure and provide important additional baselines [1]–[3].

Random forest constructs multiple trees from resampled observations and candidate predictors, then aggregates their predictions [7], [8]. This strategy can capture nonlinear responses and complex interactions without imposing a single global functional form. XGBoost, by contrast, builds an additive ensemble in which successive trees optimize a regularized loss [9], [10]. The learning rate, tree depth, subsampling, number of boosting iterations, and early-stopping criterion can markedly affect its forecasting behavior. LSTM applies recurrent network gates to update memory states across an ordered sequence [11]. Its use is most

defensible when sufficient observations and meaningful temporal patterns are available; when data are sparse or largely tabular, tree ensembles may be competitive or more stable [14]–[16].

The original source states that models were allocated 70% of observations for training and 30% for testing, with hyperparameter tuning and cross-validation. It does not indicate whether that division respected time order. For genuine forward-looking evaluation, a chronological split or rolling-origin validation is normally preferable to random shuffling of time-dependent observations [17], [19]. Model selection should take place within training data, leaving a final untouched period for assessment. Because no executable workflow is supplied, the study's stated split is described here without implying that these additional safeguards have been documented.

D. Performance indicators and interpretation

The original study reports accuracy, RMSE, and R^2, and mentions MAE in the methodology. RMSE is the square root of the mean squared error, $RMSE = \sqrt{[(1/n)\Sigma(y_i-\hat{y}_i)^2]}$, and penalizes large errors disproportionately. $MAE = (1/n)\Sigma|y_i-\hat{y}_i|$ has a more direct absolute-error interpretation. The conventional R^2 metric describes explained variation relative to a reference mean within a particular evaluation sample, although out-of-sample definitions must be explicitly reported [1], [20]. The reported accuracy index is not mathematically defined in the source. Since financial targets are described primarily as continuous quantities, it cannot safely be equated with classification accuracy; it may represent a separate normalized score, but the source provides no formula.

Meaningful performance comparison requires equivalent targets, common holdout observations, and consistent preprocessing. RMSE cannot be directly pooled across cash flow in millions of rupees, percentages, and market-price series without normalization or an explicit aggregation rule. Confidence intervals, paired forecast tests, and evaluation across repeated forecast origins would permit stronger evidence than a single summary table. The Diebold–Mariano framework provides one approach to assessing differences in predictive loss when its assumptions are met [21]. Because the source supplies neither residuals nor repeated forecast errors, no statistical test is recalculated in this revision.

E. Decision-support and governance framework

The decision-support interpretation links forecast quantities to decisions without assuming that correlation demonstrates causality. Capital budgeting requires scenario-specific free cash flow projections and sensitivity to discount rates; treasury operations depend on near-term cash inflows and payment obligations; and risk committees need downside distributions, limits, and stress conditions [4]–[6]. A deployment-ready evaluation should therefore assess both forecast loss and decision loss, for example financing costs, liquidity breaches, or investment selection under uncertainty. Appropriate safeguards include model

versioning, monitoring for drift, audit trails, privacy protection, and interpretable diagnostics [22]–[26].

Results

A. Financial characteristics of the reported sample

Table 1 reproduces the descriptive indicators from the source manuscript. Average revenue growth is positive at 12.50%, with observations extending from −5.20% to 25.80%. Operating profit averages 10.80%, while stock returns average 10.20%. Cash flow is reported in millions of Indian rupees, with a mean of 8.40 and a minimum of −2.50. Negative observed minima demonstrate that unfavorable outcomes are included in the summarized material, although they do not by themselves reveal sample composition, forecast difficulty, or the frequency of adverse events.

Table 1. Descriptive statistics of reported financial variables

Variable	Mean	Standard deviation	Minimum	Maximum
Revenue growth (%)	12.50	6.20	−5.20	25.80
Cash flow (₹ million)	8.40	4.10	−2.50	15.60
Operating profit (%)	10.80	5.50	−3.40	20.30
Stock return (%)	10.20	7.30	−8.10	22.40

The largest stated standard deviation among the percentage measures is associated with stock returns (7.30 percentage points), followed by revenue growth (6.20) and operating profit (5.50). Such comparison describes dispersion on the reported percentage scale, not a formal volatility ranking after adjusting for timing or firm-level dependencies. The source does not indicate whether observations were monthly, quarterly, or annual, and this ambiguity limits conclusions about the persistence of fluctuations. The numerical table should therefore be viewed as a summary rather than a complete characterization of the dataset.

B. Comparative predictive performance

Table 2 contains the central model comparison. The reported XGBoost results are accuracy 0.90, RMSE 0.11, and R^2 0.88. Random forest follows at 0.85, 0.14, and 0.82, respectively; LSTM is reported at 0.83, 0.16, and 0.80. Linear regression has the lowest reported accuracy (0.72) and R^2 (0.65), and the highest RMSE (0.21). Accordingly, XGBoost ranks first under all three tabulated indicators, provided the metrics were calculated on equivalent evaluation tasks.

Table 2. Reported model performance

Model	Reported accuracy	RMSE	R^2
Linear regression	0.72	0.21	0.65
Random forest	0.85	0.14	0.82
XGBoost	0.90	0.11	0.88
LSTM	0.83	0.16	0.80

Relative to the linear-regression baseline, the reported reduction in RMSE for XGBoost is (0.21−0.11)/0.21, approximately 47.6%. The corresponding difference in R^2 is 0.23 points and

in the undefined accuracy indicator is 0.18 points. These are arithmetic contrasts between the published summary values, not independent experimental estimates or indicators of statistical significance. Similarly, the RMSE values for random forest and LSTM are lower than the reported linear-regression value, yet the comparison lacks target-specific error distributions, uncertainty bounds, and details of preprocessing. Consequently, the evidence supports a descriptive ranking of reported scores, but not a universal claim that one architecture is optimal in all financial forecasting settings.

C. Qualitative comparison of forecasting approaches

Table 3 reproduces the source's categorical comparison of conventional and AI-based methods. It characterizes conventional forecasting as having moderate predictive accuracy, high error, and high risk, whereas AI-based methods are described as having high accuracy, low error, and low risk. These labels are interpretive rather than separately measured outcomes. In particular, the risk categories are not defined as financial loss, predictive uncertainty, model risk, or another observable measure. The categories should not be treated as additional empirical estimates or evidence of realized risk reduction.

Table 3. Qualitative comparison reported in the original study

Forecasting method	Prediction accuracy	Error rate	Risk level
Traditional methods	Moderate	High	High
AI-based models	High	Low	Low

D. Implications for corporate financial decisions

The source identifies four corporate functions that could benefit from higher-quality forecasts: capital budgeting, risk management, financial planning, and working-capital management. These are relevant applications rather than measured treatment effects in the dataset. In capital budgeting, improved cash-flow forecasts may reduce error in projected project values, but the impact depends on forecast horizon, discounting, and the distribution of downside outcomes. In risk management, predictive scores may help identify potential exposures, yet decision makers also need reliable uncertainty estimates and stress tests. In financial planning, timely forecast updates may help revise expense envelopes and funding assumptions. In working-capital management, short-horizon forecasts could assist payment scheduling, credit facilities, and cash buffers.

No table of actual investment decisions, realized financing costs, avoided liquidity deficits, or changes in corporate performance appears in the source. Thus, the current evidence establishes a model-level comparison as reported by the authors, whereas downstream management benefits remain plausible mechanisms requiring future verification. The distinction is important because a change in predicted values does not automatically change the actions taken by managers, and a change in actions does not necessarily improve realized financial outcomes.

Discussion

A. Explaining the observed model ordering

The reported advantage of XGBoost is consistent with the practical strengths of gradient-boosted trees for structured predictors. Boosting can flexibly represent threshold effects, interactions, and uneven signal across variables while permitting regularization [9], [10]. Corporate balance-sheet and income-statement datasets are frequently tabular and may contain relationships such as nonlinear leverage effects or interactions between operating margins and liquidity. A tree ensemble can accommodate these patterns without requiring every functional form to be specified in advance. This explanation, however, is a theoretical interpretation of the observed ranking; the source does not provide feature effects, ablation studies, or diagnostics that would confirm why XGBoost performed best in its particular sample.

The competitive performance of random forest is similarly compatible with its variance-reduction rationale [7], [8]. Averaging across trees can produce robust predictions when individual trees are unstable, but its effectiveness depends on predictor structure, data quantity, and dependence among observations. LSTM's third-place position does not establish that sequential neural methods are intrinsically weaker. Recurrent networks may be useful when sufficient longitudinal information is available and the temporal ordering contains predictive signal [11], [14]. Conversely, a short panel with many firms and comparatively few observations per firm may favor boosted trees. Without the sequence length, sample size, architecture, and tuning details, the present findings cannot resolve that question.

Comparative studies and major forecasting competitions have repeatedly demonstrated that forecasting performance varies across datasets and that sophisticated methods do not inevitably dominate well-configured statistical benchmarks [19], [27]–[30]. The present study's use of a linear-regression comparator is informative but insufficient to conclude that machine learning exceeds the broader class of traditional forecasting methods. Additional baselines should include seasonal naïve forecasts, autoregressive integrated moving average models, exponential smoothing, and regularized linear regression, where appropriate. Such comparisons would clarify whether the gain arises from nonlinear function approximation, superior handling of serial dependence, richer predictors, or differences in tuning effort.

B. Forecast quality versus managerial value

For corporate finance, the relevant outcome is often a decision rather than a point prediction. A projected revenue increase informs capacity planning only when translated into hiring, inventory, or financing choices. A forecast cash deficit matters because it can alter credit-line utilization or the timing of receivables and payables. Capital budgeting similarly depends on discounted cash-flow distributions, opportunity costs, strategic options, and risk

appetite [4]–[6]. Therefore, an improvement in RMSE is a potentially useful intermediate outcome, not a direct estimate of increased firm value.

The link between predictive and decision performance also depends on how the model treats tail events. RMSE disproportionately penalizes large deviations, which can be useful where severe misses are consequential, but it does not describe forecast interval coverage or the probability of crossing a liquidity threshold. Probabilistic forecasting, quantile regression, and calibrated uncertainty estimates may be necessary for managers to distinguish expected performance from downside exposure [30]–[33]. Relevant evaluations could measure the cost of a cash shortfall, number of covenant breaches anticipated in advance, average capital tied up in precautionary liquidity, or net present value of investments selected using model-supported forecasts.

Implementation additionally requires organizational processes that turn forecasts into auditable recommendations. Finance teams should document who approves forecast use, how exceptions are handled, and how disagreements between model output and managerial judgment are resolved. Stable metrics on a historical test set cannot substitute for live model surveillance. A model may deteriorate after accounting-policy changes, shifts in market structure, alterations in customer demand, or extraordinary macroeconomic events. Regular backtesting, drift monitoring, and retraining criteria are needed to prevent initially favorable results from being mistaken for enduring performance [22]–[25].

C. Interpretability, responsibility, and technical assurance

Financial forecasting creates a tension between flexible prediction and transparent justification. A financial manager often needs to explain why expected receipts decline or why a funding gap expands. For tree ensembles, global and local feature-attribution approaches such as SHAP can help characterize how predictors contribute to model outputs [34], while local surrogate explanations provide another perspective [35]. These techniques should be used diagnostically, because an explanation of a model prediction is not equivalent to a causal explanation of the economic process. Correlated predictors, feature engineering, and changes in baseline distributions can also affect interpretive narratives.

Model development should be supported by reproducible data lineage, clear target definitions, controlled access, and review of potential biases. Models trained on publicly available financial statements may still suffer from coverage bias if observations disproportionately represent large listed firms. Survivorship bias can occur if only firms remaining in a database at the end of the sample are retained. Look-ahead bias arises when data publication timing is ignored. More generally, financial data are nonstationary, and historical relationships may fail under structural change [18], [24], [36]. These concerns are

particularly important for models that combine financial indicators with sentiment or other alternative data.

Responsible governance does not mean defaulting to the simplest algorithm. Instead, model complexity should be justified in relation to the incremental forecast benefit, decision materiality, monitoring capacity, and available expertise. Formal documentation of purpose, performance boundaries, ownership, and oversight supports scrutiny by finance managers, internal auditors, and other stakeholders [25], [26], [37]. Incorporating uncertainty and explanation into a decision process may be more valuable than a small improvement in one summary accuracy index if it improves understanding and responses to foreseeable failure modes.

Conclusions

The comparative results supplied in the original study favor machine learning over the reported linear-regression comparator. XGBoost is the leading specification in the original performance table, with reported accuracy of 0.90, RMSE of 0.11, and R^2 of 0.88. Random forest and LSTM also show higher reported scores than linear regression. These findings indicate that nonlinear algorithms warrant consideration in corporate-finance forecasting, particularly for tabular data containing potentially complex relationships. They do not, however, prove that machine learning is uniformly superior to every statistical forecasting method or that it necessarily improves realized corporate decisions.

The practical implication is to treat advanced forecasting as part of a governed financial decision system. Appropriate implementation combines clearly specified targets, time-consistent data pipelines, credible benchmarks, explainable predictions, uncertainty estimates, ongoing monitoring, and a transparent mapping from forecasts to managerial actions. Before submission to a research journal, the underlying study should disclose sample design, data provenance, outcome definitions, test-period construction, tuning, and statistical uncertainty. With these elements in place, future research can investigate whether reported predictive gains persist out of sample and whether they produce measurable benefits for budgeting, liquidity management, investment appraisal, and risk governance.

Conflict of Interest

The authors declare no conflict of interest.

References

[1] R. J. Hyndman and G. Athanasopoulos, Forecasting: Principles and Practice, 3rd ed. Melbourne, Australia: OTexts, 2021. Available: https://otexts.com/fpp3/

[2] R. S. Tsay, Analysis of Financial Time Series, 3rd ed. Hoboken, NJ, USA: Wiley, 2010.

[3] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, Time Series Analysis: Forecasting and Control, 5th ed. Hoboken, NJ, USA: Wiley, 2015.
[4] R. A. Brealey, S. C. Myers, and F. Allen, Principles of Corporate Finance, 13th ed. New York, NY, USA: McGraw-Hill, 2020.
[5] J. C. Hull, Risk Management and Financial Institutions, 5th ed. Hoboken, NJ, USA: Wiley, 2018.
[6] A. Damodaran, Investment Valuation: Tools and Techniques for Determining the Value of Any Asset, 3rd ed. Hoboken, NJ, USA: Wiley, 2012.
[7] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001, doi: 10.1023/A:1010933404324.
[8] G. Biau and E. Scornet, "A random forest guided tour," TEST, vol. 25, pp. 197–227, 2016, doi: 10.1007/s11749-016-0481-7.
[9] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," Annals of Statistics, vol. 29, no. 5, pp. 1189–1232, 2001, doi: 10.1214/aos/1013203451.
[10] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785–794, doi: 10.1145/2939672.2939785.
[11] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997, doi: 10.1162/neco.1997.9.8.1735.
[12] S. Gu, B. Kelly, and D. Xiu, "Empirical asset pricing via machine learning," Review of Financial Studies, vol. 33, no. 5, pp. 2223–2273, 2020, doi: 10.1093/rfs/hhaa009.
[13] K. J. Kim, "Financial time series forecasting using support vector machines," Neurocomputing, vol. 55, nos. 1–2, pp. 307–319, 2003, doi: 10.1016/S0925-2312(03)00372-2.
[14] T. Fischer and C. Krauss, "Deep learning with long short-term memory networks for financial market predictions," European Journal of Operational Research, vol. 270, no. 2, pp. 654–669, 2018, doi: 10.1016/j.ejor.2017.11.054.
[15] O. B. Sezer, M. U. Gudelek, and A. M. Ozbayoglu, "Financial time series forecasting with deep learning: A systematic literature review: 2005–2019," Applied Soft Computing, vol. 90, art. 106181, 2020, doi: 10.1016/j.asoc.2020.106181.
[16] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.
[17] C. Bergmeir, R. J. Hyndman, and B. Koo, "A note on the validity of cross-validation for evaluating autoregressive time series prediction," Computational Statistics & Data Analysis, vol. 120, pp. 70–83, 2018, doi: 10.1016/j.csda.2017.11.003.
[18] M. López de Prado, Advances in Financial Machine Learning. Hoboken, NJ, USA: Wiley, 2018.
[19] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "The M4 Competition: 100,000 time series and 61 forecasting methods," International Journal of Forecasting, vol. 36, no. 1, pp. 54–74, 2020, doi: 10.1016/j.ijforecast.2019.04.014.
[20] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," International Journal of Forecasting, vol. 22, no. 4, pp. 679–688, 2006, doi: 10.1016/j.ijforecast.2006.03.001.
[21] F. X. Diebold and R. S. Mariano, "Comparing predictive accuracy," Journal of Business & Economic Statistics, vol. 13, no. 3, pp. 253–263, 1995, doi: 10.1080/07350015.1995.10524599.
[22] D. Sculley et al., "Hidden technical debt in machine learning systems," in Advances in Neural Information Processing Systems, vol. 28, 2015, pp. 2503–2511.
[23] H. R. Varian, "Big data: New tricks for econometrics," Journal of Economic Perspectives, vol. 28, no. 2, pp. 3–28, 2014, doi: 10.1257/jep.28.2.3.
[24] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," ACM Computing Surveys, vol. 46, no. 4, art. 44, 2014, doi: 10.1145/2523813.
[25] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, 2023, doi: 10.6028/NIST.AI.100-1.
[26] OECD, OECD Principles on Artificial Intelligence. Paris, France: OECD, 2019. Available: https://oecd.ai/en/ai-principles
[27] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "M5 accuracy competition: Results, findings, and conclusions," International Journal of Forecasting, vol. 38, no. 4, pp. 1346–1364, 2022, doi: 10.1016/j.ijforecast.2021.11.013.
[28] S. Smyl, "A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting," International Journal of Forecasting, vol. 36, no. 1, pp. 75–85, 2020, doi: 10.1016/j.ijforecast.2019.03.017.
[29] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "Statistical and Machine Learning forecasting methods: Concerns and ways forward," PLoS ONE, vol. 13, no. 3, art. e0194889, 2018, doi: 10.1371/journal.pone.0194889.

[30] S. Makridakis et al., "The M5 uncertainty competition: Results, findings and conclusions," International Journal of Forecasting, vol. 38, no. 4, pp. 1365–1385, 2022, doi: 10.1016/j.ijforecast.2021.10.009.
[31] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," Journal of the American Statistical Association, vol. 102, no. 477, pp. 359–378, 2007, doi: 10.1198/016214506000001437.
[32] R. Koenker and G. Bassett, Jr., "Regression quantiles," Econometrica, vol. 46, no. 1, pp. 33–50, 1978, doi: 10.2307/1913643.
[33] C. Chatfield, "Calculating interval forecasts," Journal of Business & Economic Statistics, vol. 11, no. 2, pp. 121–135, 1993, doi: 10.1080/07350015.1993.10509938.
[34] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems, vol. 30, 2017, pp. 4765–4774.
[35] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?: Explaining the predictions of any classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.
[36] B. M. Barber, R. Lehavy, M. McNichols, and B. Trueman, "Can investors profit from the prophets? Security analyst recommendations and stock returns," Journal of Finance, vol. 56, no. 2, pp. 531–563, 2001, doi: 10.1111/0022-1082.00336.
[37] M. Mitchell et al., "Model cards for model reporting," in Proc. Conf. Fairness, Accountability, and Transparency, 2019, pp. 220–229, doi: 10.1145/3287560.3287596.
[38] J. H. Heaton, N. G. Polson, and J. H. Witte, "Deep learning in finance," Annual Review of Financial Economics, vol. 11, pp. 145–181, 2019, doi: 10.1146/annurev-financial-110217-022753.
[39] G. James, D. Witten, T. Hastie, and R. Tibshirani, An Introduction to Statistical Learning: With Applications in R, 2nd ed. New York, NY, USA: Springer, 2021.
[40] S. Arlot and A. Celisse, "A survey of cross-validation procedures for model selection," Statistics Surveys, vol. 4, pp. 40–79, 2010, doi: 10.1214/09-SS054.

Authors



Dr. T. Venkatesan is Vice Principal in the Department of Management Studies. He obtained his M.A. (Public Administration), M.Com, M.Phil, and MBA from Annamalai University and a Ph.D. degree in Management Studies, specializing in Finance from Bharathiar University, Tamil Nadu. He has published more than 40 articles in national and international journals. He has organized four international conferences. His areas of interest are Finance and Management Accounting. He has taught a diverse set of subjects like Management Accounting, Financial Accounting, Advanced Corporate Finance, Derivatives, and Investments. (email: proftvv@gmail.com).



Dr. A. Stephen is the Director of the Knowledge Institute of Technology, Salem, Tamil Nadu, India. His research focuses on human resource management, financial literacy, and consumer behavior. He has published studies on credit card promotion schemes, challenges faced by women entrepreneurs, and financial literacy among young people in Tamil Nadu, contributing to social science and management research. (email: directorkbs@kiot.ac.in).



P. Elakkiya is pursuing an MBA in Finance and Business Analytics at Knowledge Institute of Technology, Salem. She holds a B.Com in Computer Applications from Sona College of Arts and Science. Her interests include financial management, business analytics, and emerging technologies. An internship at Jayanti Herbs and Spices provided exposure to finance and accounting operations. She is familiar with Excel, Power BI, Tally, and Microsoft Office and values continuous learning. (email: elakkiya2005.slm@gmail.com).



Dr. Maria Joy C. Iglesia has a doctoral degree in Business Administration and is currently designated as the Dean of the College of Tourism, Hospitality, and Business Management in Camarines Sur Polytechnic Colleges, Philippines. She has several published research studies and is working on research related to market assessment, gender and development awareness in partnered entrepreneurs, the readiness of tertiary students to business concepts, theories, and applications, and the business management of slaughterhouses in the Philippines. She teaches business courses with expertise in microeconomics, macroeconomics, and other related subjects. (email: joyidian@cspc.edu.ph).



Roque B. Cruz II is a faculty member of the College of Tourism, Hospitality, and Business Management and the Program Chair of the BSBA-Financial Management Program at Camarines Sur Polytechnic Colleges (CSPC), Philippines. Since 2019, he has taught courses in financial management, entrepreneurship, business administration, strategic management, and innovation. His academic interests include financial management, financial literacy, entrepreneurship, and business development. He holds a BSBA in Financial Management from the University of the East–Manila and an MBA from San Beda University–Manila. (email: rcruz@cspc.edu.ph).