Fundamental vs. Technical Features for Stock Trend Prediction: A Multi-Horizon Data Stream Learning Analysis

José Júnior de Oliveira Silva*, Roberto Souto Maior de Barros*, Silas Garrido Teixeira de Carvalho Santos*†
**Centro de Informática, Universidade Federal de Pernambuco*, 50740-560, Recife-PE, Brazil
{jjos,roberto,sgtcs}@cin.ufpe.br
†*SiDi Institute*, 51110-160, Recife-PE, Brazil
s.garrido@sidi.org.br

***Abstract*—This paper systematically compares the predictive power of fundamental versus technical features to forecast stock price movements in a data stream learning setting. To enable this comparison, we construct the B3 Technical and Fundamental (B3TF) dataset, which provides daily fundamental ratios rigorously aligned to avoid look-ahead bias, as well as adjusted market data for companies listed on the Brazilian Stock Exchange (B3). This novel contribution solves the temporal misalignment problem. A comparative analysis was conducted using Hoeffding Trees and a delayed Prequential protocol across nine forecast horizons. To the best of our knowledge, this is the first work to make this comparison using a strictly incremental learning protocol. Our results reveal a structural divergence. While technical models maintain stable accuracy around the random baseline, their discriminative power degrades, with global kappa turning negative (reaching $\kappa = -0.0246$ at $h = 240$) in the long term. In contrast, fundamental features demonstrate greater robustness. Although the predictive signal remains marginal, it persists over longer horizons, reaching a global kappa peak ($\kappa = 0.0469$) and accuracy = 52.39% at $h = 150$ trading days. Validating these findings on 10 highly liquid assets confirms the trend: the fundamental model retained stability ($\kappa = 0.0129$), whereas the technical model collapsed ($\kappa = -0.0537$). These findings delineate a boundary of predictability: fundamental ratios provide persistent signals driven by value convergence, distinguishing them from noise even when price-based models fail.**

***Index Terms*—Data Stream Learning, Fundamental Analysis, Stock Price Prediction.**

I. Introduction

Predicting stock price movements is a highly complex problem due to the dynamic and stochastic nature of financial markets [1]. Financial time series are typically non-stationary and, therefore, susceptible to concept drift, a phenomenon characterized by shifts in the underlying probability distribution over time [2]. Several factors can drive this change, including economic shifts, regulatory updates, and variations in investor sentiment [1], [3]. Traditional batch learning models often rely on the assumption of stationarity, which leads to significant performance degradation when applied in environments with continuous data streams [2], [4]. Thus, several data stream learning techniques have been proposed to address this problem [5].

In financial forecasting, predictive methodologies can be categorized into two main areas: *technical analysis* and *fundamental analysis*. The technical analysis basis is the belief that historical price and volume data inherently reflect collective market psychology and future trends. It relies on quantitative indicators—such as moving averages and momentum oscillators—to identify recurring patterns in the asset's prices [6], [7]. Conversely, fundamental analysis seeks to determine an asset's intrinsic value by evaluating the underlying economic health of the entity that issues it. This approach integrates a broad spectrum of variables, ranging from macroeconomic indicators to company-specific financial statements, e.g., earnings reports, cash flow, and debt ratios, under the assumption that the market price will eventually converge to this intrinsic valuation over the long term [6], [8].

The feasibility of these predictive methodologies is frequently challenged by the Efficient Market Hypothesis (EMH). In its semi-strong form, EMH states that asset prices instantly reflect publicly available information; consequently, neither technical nor fundamental analysis can consistently generate abnormal returns [9]. In contrast, the Adaptive Market Hypothesis (AMH) [10] reconciles the rationale of EMH with behavioral anomalies, viewing market dynamics as a continuous process of adaptation. In fact, some works [3], [11], [12] show that market efficiency is a dynamic and time-varying attribute, often characterized by non-stationarity and long-range dependence.

Technical analysis is the main focus of the literature on financial data stream learning, relying mainly on the vast availability of price and volume data in comparison to fundamental variables [6], [13]. Such market-based features effectively capture transient sentiment and short-term volatility, but they are decoupled from the asset's intrinsic valuation. By relying exclusively on recent price and volume history, these methods face inherent limitations in deciphering the causal mechanisms that drive structural market shifts [1], [6]. The integration of fundamental data into stream learning models is limited, primarily due to the challenge of aligning low-frequency quarterly reports with continuous market streams, as well as the scarcity of publicly available datasets suitable for this task.

To address this frequency mismatch, we constructed a dataset that integrates daily market data (prices and volume) from the Brazilian Stock Exchange (B3[1]) with fundamental ratios derived from quarterly financial reports retrieved from the Brazilian Securities and Exchange Commission (CVM[2]). By strictly indexing these accounting metrics by their public release dates, we generate a continuous stream of daily fundamental features aligned with the corresponding price and volume series. This approach transforms static quarterly data into dynamic features, enabling the model to capture real-time market re-evaluations of intrinsic value.

Building upon this dataset, we examine how fundamental features affect stock price prediction in a data stream learning setting. To this end, we conducted a systematic comparison between an online classifier trained on daily-updated fundamental features and the same classifier trained on technical (price and volume) data. Furthermore, to investigate the potential of fundamental variables for long-term forecasting, we evaluated the model's performance across nine distinct forecast horizons.

The remainder of this paper is organized as follows. Section II provides the theoretical background and reviews related work on financial stream learning. Section III details the proposed methodology, including the construction of the B3TF dataset and the experimental setup. Section IV presents the results and discusses the comparative performance of the feature domains across multiple horizons. Section V concludes the study and outlines future research directions. Code, datasets, and supplementary materials—including an ARF ablation study and Wilcoxon statistical validation—are publicly available at https://github.com/junior-oliveira/fundamental-vs-technical-streams.

II. Background and Related Work

Most research on stock price forecasting relies on technical analyses or a combination of sentiment and technical features [6]. This predominance is partly due to the greater availability of data for these approaches compared to the data required for fundamental analysis. The latter, although crucial for long-term strategies, accounts for only approximately 23% of stock price forecasting studies [13], [14].

A. Predictive Paradigms and Model Evolution

Stock price prediction techniques can be divided into two approaches: *regression*, which seeks to estimate exact future prices, and *classification*, which focuses on forecasting directional price trends (e.g., up or down) [7], [14]. While regression offers precision, it is susceptible to noise; even marginal estimation errors can be interpreted as trend reversals, compromising the integrity of trading models [15]. As a result, classification frameworks are the most common approaches in the financial machine learning community [13]. The rationale is that correctly identifying the sign of price movement is frequently the primary determinant of trading profitability, while offering substantially greater resilience to the inherent stochasticity of financial markets [16].

Historically, stock price forecasting has evolved from linear statistical models, such as ARIMA [17], to modern Deep Learning (DL) approaches. Traditional models often fail to capture the non-linear dynamics of financial markets, prompting a shift toward machine learning algorithms like ANN and SVM [7], [13], [18]. Currently, DL represents the state-of-the-art due to its automatic feature extraction capabilities [1], [19]. Within this paradigm, Long Short-Term Memory (LSTM) networks are predominant for managing temporal dependencies and vanishing gradients [15], [16]. At the same time, Convolutional Neural Networks (CNN) are increasingly used for feature extraction and processing time series as image-like data [18], [19].

B. Challenges in Adaptive Financial Modeling

Despite the popularity of deep learning approaches, these methods generally require significant amounts of data to achieve convergence. Therefore, the model is usually trained using all available assets as an aggregated dataset in a global model approach. This process addresses data scarcity, but, at the same time, tends to dilute the unique characteristics of individual assets. On the other hand, the research in [20] showed that traditional machine learning local models, trained per stock, can outperform deep learning global architectures as they can capture the specificities and idiosyncrasies of each stock.

From a methodological standpoint, the current standard practice continues to depend on batch learning techniques [21]. These techniques operate under the assumption that the data-generating process is stationary. Consequently, as non-adaptive models, they become obsolete in inherently non-stationary financial environments [4]. On the other hand, Data Stream Learning (DSL) constantly updates the model to changes in the underlying data distribution (concept drift) [5]. This capacity to learn and manage evolving concepts incrementally, while processing financial data as it arrives in a continuous stream [5], makes DSL a practical approach for predicting stock price movements [21], [22].

Time series data are frequently treated as data streams, often arriving as sequential online data [5]. Crucially, unlike conventional data streams assumed to be independent and identically distributed (i.i.d.), data points in a time series are expected to exhibit strong temporal dependence [23]. Consequently, applying traditional cross-validation (CV) methods is unsuitable for training, as these techniques produce undesired results due to serially correlated features [16], potentially mixing the temporal order of data [2]. Instead, the preferred methodology for performance monitoring in these environments is Prequential evaluation (Interleaved Test-Then-Train), which is highly advocated for data streams. This methodology is highly effective because it explicitly ensures that each incoming data instance is used first to test the model before being utilized for subsequent training, guaranteeing that predictions are constantly evaluated on unseen data [2].

[1] https://www.b3.com.br/

[2] https://www.gov.br/cvm

Despite advances in adaptive learning, integrating fundamental analysis into this paradigm remains a significant challenge. The primary obstacle is the frequency mismatch between low-frequency fundamental data, typically quarterly reports, and daily price streams [6]. While market data are naturally continuous, financial statement updates are usually static for long periods, breaking the continuity required for online learning. This has resulted in a literature heavily skewed toward market-based features, leaving the adaptive, real-time use of fundamental ratios largely unexplored.

III. Proposed Methodology

The first step in the investigation of fundamental features impact on stock price movement prediction was to construct the B3 Technical and Fundamental (B3TF) dataset, which integrates daily market data with financial ratios derived from quarterly accounting disclosures. To enable comparative analysis, the dataset is organized into two distinct feature domains: *fundamental*, which provides a daily series of five financial ratios (P/E, P/B, P/S, EV/EBIT, and ROE); and *technical*, comprising adjusted daily market data (open, high, low, and close prices, and trading volume).

Market data were processed from raw B3 flat files, while fundamental data were sourced from standardized quarterly and annual financial reports provided by CVM. Each domain comprises 58 subsets (one per asset), generated for each of the nine forecast horizons $h \in \{1, 30, 60, \ldots, 240\}$ trading days.

Raw fundamental data were systematically parsed to extract the specific accounting line items required to compute the target indicators. A critical step was the implementation of a dual-timestamp structure that records both the fiscal period to which the data refers and the date it was officially published by CVM: this distinction is essential because financial reports become official information only after public release [24]. Thus, for any trading day t, the feature vector only incorporates information published on or before day t, preventing look-ahead bias. To transform low-frequency quarterly data into continuous daily features, we calculate the ratio between the daily market price at time t and the fundamental metrics derived from the reports published before t. We use a rolling 12-month sum for flow variables (e.g., net income) to capture a full year of operations, while using the most recent reported value for stock variables (e.g., book value). This ensures the features reflect the most current financial position while fluctuating daily in response to market price movements.

The sample period starts in 2011, aligning with the availability of standardized electronic reports from the CVM. However, this initial year serves exclusively as an initialization window for aggregated calculations. Consequently, the effective experimental evaluation begins in early 2012, aligned with the publication of the 2011 annual reports, and extends through June 2023. To ensure market liquidity and data continuity, the dataset is restricted to common (ON) and preferred (PN) shares that were continuously listed on the B3's main exchange throughout this entire period. While acknowledging the resulting survivorship bias [24], this criterion ensures the longitudinal stability required for a controlled, head-to-head comparison between fundamental and technical features.

A. Problem Formulation and Data Representation

We formulate the prediction task as a binary classification problem. Formally, given a forecast horizon h, the prediction target $y_t^{(h)}$ indicates the direction of price movement, defined below, where C_t denotes the adjusted closing price at day t.

$$y_t^{(h)} = \begin{cases} 1 \quad \text{(Up)} & \text{if } C_{t+h} > C_t \\ 0 \quad \text{(Down)} & \text{otherwise} \end{cases} \tag{1}$$

To capture temporal dependencies, we construct input feature vectors using a sliding window approach. For a given trading day t, each input variable k (e.g., *close* price or *P/E* ratio) is represented by a lag-vector of length $L+1$:

$$\mathbf{v}_t^{(k)} = \left[x_t^{(k)},\, x_{t-1}^{(k)},\, \ldots,\, x_{t-L}^{(k)}\right] \tag{2}$$

The final input feature vector $\mathbf{X}_t$ is the concatenation of these lag-vectors across all K variables in the domain:

$$\mathbf{X}_t = \left[\mathbf{v}_t^{(1)},\, \ldots,\, \mathbf{v}_t^{(K)}\right]. \tag{3}$$

The historical context size L is treated as a hyperparameter and optimized empirically (see Section IV). All features are kept in their raw, unscaled form.

B. Experimental Setup and Performance Evaluation Metrics

Experiments were conducted using the MOA framework (release 2021.07) [25] via a local modeling strategy, where a separate Hoeffding Tree, using default hyperparameters in MOA, was trained for each asset. This design choice isolates intrinsic dynamics and prevents global model bias, simulating a realistic, asset-specific analysis. To enforce realistic temporal constraints and prevent data leakage [24], we employ a delayed Prequential protocol. For each day t, the model predicts using $\mathbf{X}_t$ immediately, but the training update is deferred until $t+h$ when the true label becomes available.

To assess classifiers beyond simple accuracy, we adopted Cohen's kappa (κ), which corrects for chance agreement and better reflects non-random predictive power in noisy financial data. We also decompose performance into precision and recall to diagnose model behavior. High precision indicates reliable signals, while low precision with high recall suggests an overreactive strategy prone to noise overfitting. This decomposition selectively distinguishes correct models from those that merely guess frequently.

IV. Results and Discussion

Our experimental evaluation is structured in three sequential phases. First, to balance historical context against the curse of dimensionality in streaming learning, we determine the optimal historical window size (L). Second, we analyze metric dispersion to contrast the predictive stability of fundamental features against technical inputs. Third, we evaluate performance across all forecast horizons using a global micro-averaging strategy to capture systemic, market-wide trends. The identified optimal predictive horizon ($h = 150$) is then subjected to a final robustness validation using a portfolio of high-liquidity assets.

Figure 1 shows the sensitivity of predictive performance to the historical window size (L). For the fundamental domain, extending the window size leads to a consistent decoupling of metrics: accuracy declines to a local minimum around $L = 20$ and stabilizes at a lower plateau thereafter, while the kappa statistic degrades monotonically. In the technical domain, the deterioration in kappa is even more pronounced, deepening steadily into negative territory as L increases. Although technical accuracy shows a slight non-monotonic recovery for $L > 15$, this comes at the severe cost of drastically worse kappa. Based on these empirical observations, we adopt $L = 1$ for both domains to enforce parsimony and maximize predictive reliability.

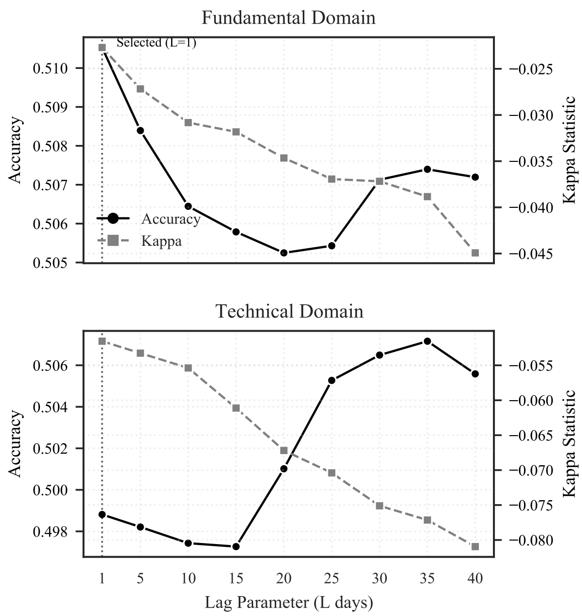

Fig. 1. Impact of historical window size (L) on predictive performance. The dual-axis plots contrast accuracy (solid line) and kappa statistic (dashed line) for the fundamental (top) and technical (bottom) domains. The vertical dotted line indicates the selected window ($L = 1$).

A. Asset Performance Distribution

This section analyzes the cross-sectional performance of the models across all 58 assets. For a stable, horizon-independent view, we compute asset-level mean performance metrics by averaging across all forecast horizons using the optimal input window ($L = 1$). Figure 2 presents the resulting distribution, enabling a direct comparison between feature domains.

The analysis confirms a structural divergence. While both models exhibit a median kappa below zero, the fundamental model shows greater robustness. Its median kappa (-0.0115) is higher than the technical model's (-0.0510), and its accuracy has a tighter, more elevated interquartile range (0.4770–0.5345) versus the technical range (0.4418–0.5385). This indicates a more consistent predictive strategy.

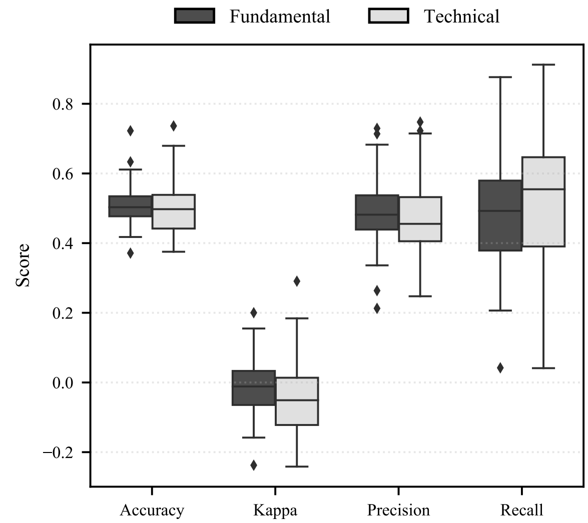

Fig. 2. Distribution of asset-level mean performance metrics ($L = 1$). Each boxplot summarizes cross-sectional variation across the 58 assets. The dark grey (fundamental) and light grey (technical) boxes illustrate the comparative reliability of the domains.

Conversely, the technical model shows a significant decline in discriminative performance. Its lower quartile kappa reaches -0.1220, signaling predictive failure for many assets. This is explained by an imbalanced trade-off: the technical model achieves higher median recall (0.5546 vs. 0.4925) but lower median precision (0.4553 vs. 0.4817). This pattern confirms the technical classifier is overreactive, generating a higher volume of less reliable signals due to stochastic noise.

B. Forecast Horizon Analysis

We next analyze the temporal evolution of predictive power across forecast horizons. To mitigate the high variance inherent in long-term forecasting and reveal market-wide trends, we employ a global micro-averaging strategy. This involves aggregating the confusion matrices from all 58 assets at each horizon ($h = 1$ to $h = 240$) before computing performance metrics. This approach attenuates asset-specific noise, allowing the structural divergence between the fundamental and technical domains to emerge more clearly.

Figure 3 presents the evolution of global accuracy, kappa, precision, and recall. While the cross-sectional analysis in Section IV-A revealed negative median performance across assets and horizons, this temporal perspective uncovers a more nuanced pattern with phase-dependent behavior.

Our analysis delineates three temporal phases. In the initial transitional phase ($1 \leq h \leq 60$), performance decays rapidly from the immediate horizon ($h = 1$), with both domains converging near the random baseline (kappa ≈ 0 at $h = 60$). Subsequently, a domain-specific pattern emerges. The fundamental model shows a relative improvement, reaching a local maximum at $h = 150$ (kappa $= 0.0469$, accuracy = 52.39%), which aligns with the latency period of fundamental price discovery. However, it is crucial to note that even this peak

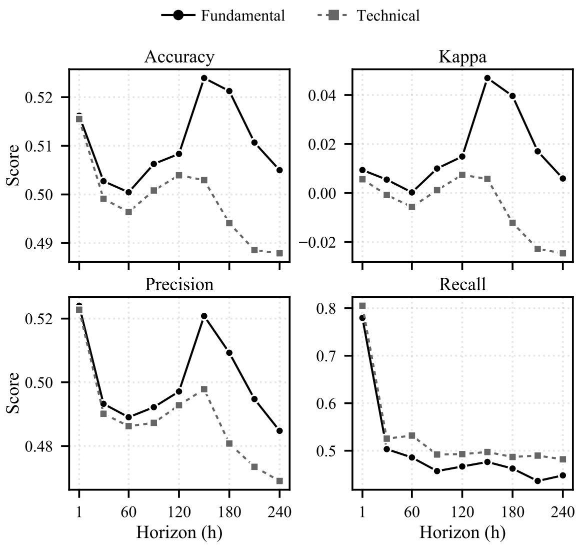

Fig. 3. Global performance evolution (micro-average). The solid line (fundamental) and dashed line (technical) reveal domain-specific temporal patterns. A relative peak for the fundamental domain emerges around $h = 150$, where global kappa reaches 0.0469, and accuracy is 52.39%, while technical performance trends downward.

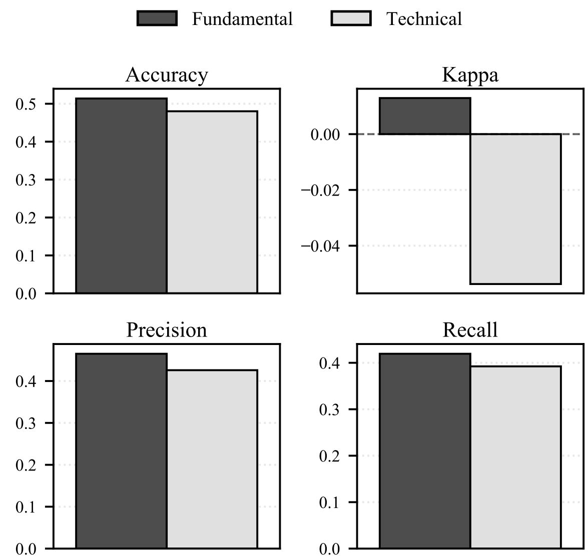

Fig. 4. Micro-average performance on a sector-diversified portfolio of high-liquidity assets ($h = 150$). The fundamental model maintains a positive kappa, while the technical model deteriorates significantly below the random baseline. This confirms the fundamental data robustness to illiquidity, though its magnitude is attenuated in this high-liquidity environment.

represents only a marginal improvement over the baseline and is inconsistent across assets, as evidenced by the negative cross-sectional median reported earlier. Finally, in the long-term phase ($h > 180$), the technical model's performance systematically deteriorates to significantly negative kappa values ($\kappa = -0.0246$ at $h = 240$), confirming its lack of predictive consistency over extended horizons.

C. High-Liquidity Asset Analysis

To validate the robustness of the performance peak at $h = 150$ days against potential illiquidity effects, we constructed a diversified portfolio of 10 highly-liquid assets representing distinct market subsectors. Assets were first ranked by their 200-day average trading volume preceding the dataset cutoff (June 30, 2023). From this ranking, we selected the single most liquid asset from each unique subsector to ensure broad economic coverage and avoid concentration bias.

The resulting portfolio is: VALE3 (Mining), PETR4 (Oil, Gas & Biofuels), ITUB4 (Financial Intermediaries), RENT3 (Car Rental), ELET3 (Electric Energy), LREN3 (Retail Trade), GGBR4 (Steel & Metallurgy), HYPE3 (Healthcare Commerce & Distribution), BRFS3 (Processed Foods), and CIEL3 (Diverse Financial Services).

Next, we aggregated the confusion matrices of these assets using a global micro-average to assess predictive performance at the optimal horizon ($h = 150$). The results, presented in Figure 4, demonstrate the resilience of the fundamental signal in a high-liquidity environment.

Note the fundamental model retains a positive, stable performance (kappa = 0.0129, accuracy = 51.38%). In contrast, the technical model suffers a significant deterioration, performing worse than a random classifier (kappa = -0.0537, accuracy = 48.02%). This clear divergence confirms that the predictive edge of fundamental features is a structural property related to information latency, not an artifact of market illiquidity.

V. Conclusion and Future Work

The primary objective of this paper was to conduct a systematic comparison between fundamental and technical feature domains within a data stream learning framework. To enable this investigation, we first addressed the critical challenge of frequency mismatch by constructing the B3TF dataset. This step provided the necessary condition—a strict leakage-free environment—to evaluate the true predictive utility of daily-updated accounting ratios against market data without look-ahead bias.

Our multi-horizon analysis revealed a distinct dichotomy in feature behavior. We demonstrated that technical features, while prevalent in the literature, lack the signal persistence required for long-term forecasting in streaming environments. The technical model exhibited rapid performance degradation, with its discriminative power regressing to or below the random baseline over extended horizons, as evidenced by its negative cross-sectional median kappa.

In contrast, fundamental features demonstrated superior robustness and a more stable predictive profile. The identification of a relative performance peak at approximately 150 trading days (global $\kappa = 0.0469$) suggests that fundamental signals can exhibit horizon-dependent persistence, potentially related to the latency of price discovery. Crucially, our robustness analysis on a high-liquidity portfolio confirmed the structural nature of this divergence. Even though the technical model's performance collapsed ($\kappa = -0.0537$) on liquid assets, falling significantly below the random baseline, the fundamental model retained a marginal but stable predictive edge ($\kappa = 0.0129$). This contrast indicates that fundamental ratios possess a resilience to short-term market noise that purely technical approaches lack over longer horizons.

Future work will expand on these findings through four key directions: (i) investigating hybrid feature strategies that combine technical and fundamental indicators to leverage both short-term momentum and long-term value signals; (ii) evaluating the robustness of this approach using advanced stream learning ensembles (e.g., Adaptive Random Forests [26]) and other state-of-the-art classifiers designed for non-stationary environments [27], [28]; (iii) performing cross-market validation by applying this methodology to international exchanges (e.g., NYSE, NASDAQ) to assess the universality of the fundamental signal latency; and (iv) exploring an expanded feature space, incorporating alternative valuation multiples and macroeconomic variables to capture a broader range of market drivers.

Acknowledgment

José Junior O. Silva is on leave from Instituto Federal de Alagoas (IFAL). Roberto S. M. Barros is supported by research grant number 305539/2022-1 from Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq). Silas Santos was previously supported by post-doctorate grant number 88887.374884/2019-00 from Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES).

Disclaimer

The views and opinions expressed in this research are solely those of the authors and do not necessarily reflect the official policy or position of the company.

References

[1] Z. Hu, W. Liu, J. Bian, X. Liu, and T.-Y. Liu, "Listening to chaotic whispers: A deep learning framework for news-oriented stock trend prediction," in *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, ser. WSDM '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 261–269. [Online]. Available: https://doi.org/10.1145/3159652.3159690

[2] J. a. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Comput. Surv.*, vol. 46, no. 4, Mar. 2014. [Online]. Available: https://doi.org/10.1145/2523813

[3] K.-P. Lim and R. Brooks, "The Evolution of Stock Market Efficiency Over Time: A Survey of the Empirical Literature," *Journal of Economic Surveys*, vol. 25, no. 1, pp. 69–108, 2011. [Online]. Available: https://doi.org/10.1111/j.1467-6419.2009.00611.x

[4] G. Ditzler, M. Roveri, C. Alippi, and R. Polikar, "Learning in non-stationary environments: A survey," *IEEE Computational Intelligence Magazine*, vol. 10, no. 4, pp. 12–25, 2015.

[5] H. M. Gomes, J. Read, A. Bifet, J. P. Barddal, and J. a. Gama, "Machine learning for streaming data: state of the art, challenges, and opportunities," *SIGKDD Explor. Newsl.*, vol. 21, no. 2, p. 6–22, Nov. 2019. [Online]. Available: https://doi.org/10.1145/3373464.3373470

[6] Y. Huang, L. F. Capretz, and D. Ho, "Machine learning for stock prediction based on fundamental analysis," in *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2021, pp. 01–10.

[7] O. Bustos and A. Pomares-Quimbaya, "Stock market movement forecast: A systematic review," *Expert Systems with Applications*, vol. 156, p. 113464, 2020.

[8] R. Frankel and C. M. Lee, "Accounting valuation, market expectation, and cross-sectional stock returns," *Journal of Accounting and Economics*, vol. 25, no. 3, pp. 283–319, 1998. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0165410198000263

[9] E. F. Fama, "Efficient capital markets: A review of theory and empirical work," *The Journal of Finance*, vol. 25, no. 2, pp. 383–417, 1970. [Online]. Available: http://www.jstor.org/stable/2325486

[10] A. W. Lo, "The adaptive markets hypothesis: Market efficiency from an evolutionary perspective," *The Journal of Portfolio Management*, 2004. [Online]. Available: https://papers.ssrn.com/abstract=602222

[11] C. Los, "Nonparametric efficiency testing of asian stock markets using weekly data," *SSRN Electronic Journal*, vol. 14, 10 2004.

[12] D. O. Cajueiro and B. M. Tabak, "The hurst exponent over time: testing the assertion that emerging markets are becoming more efficient," *Physica A: Statistical Mechanics and its Applications*, vol. 336, no. 3, pp. 521–537, 2004. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378437103011828

[13] I. K. Nti, A. F. Adekoya, and B. A. Weyori, "A systematic review of fundamental and technical analysis of stock market predictions," *Artificial Intelligence Review*, vol. 53, no. 4, pp. 3007–3057, 2020.

[14] H. Rezaei, M. Hooshmand, and K. Abdi, "Machine learning based classification and regression approaches for return and risk prediction in the stock market using fundamental feature engineering," *Discover Analytics*, vol. 3, no. 1, p. 17, 2025. [Online]. Available: https://doi.org/10.1007/s44257-025-00042-5

[15] Y. Liu, Z. Su, H. Li, and Y. Zhang, "An lstm based classification method for time series trend forecasting," in *2019 14th IEEE Conference on Industrial Electronics and Applications (ICIEA)*, 2019, pp. 402–406.

[16] A. Dezhkam, M. T. Manzuri, A. Aghapour, A. Karimi, A. Rabiee, and S. M. Shalmani, "A Bayesian-based classification framework for financial time series trend prediction," *The Journal of Supercomputing*, vol. 79, no. 4, pp. 4622–4659, 2023.

[17] G. Box, G. Jenkins, G. Reinsel, and G. Ljung, *Time Series Analysis: Forecasting and Control*. San Francisco: Holden-Day, 1970.

[18] S. P. Panda, B. Kumar Sahoo, B. Pati, T. Mukherjee, and S. Bebortta, "State-of-the-art analysis of stock market forecasting: Techniques, challenges, and future scopes," in *Proceedings of 2nd International Conference on Signal Processing, Communication, Power and Embedded System (SCOPES)*, 2024, pp. 1–5.

[19] 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, p. 106181, 2020.

[20] J. J. O. Silva, R. S. M. Barros, and S. G. T. C. Santos, "Stock price movement prediction based on optimized traditional machine learning models," in *Proceedings of IEEE Symposium Series on Computational Intelligence (SSCI)*, Mexico City, Mexico, 2023, pp. 711–716.

[21] A. L. Suárez-Cetrulo, D. Quintana, and A. Cervantes, "A survey on machine learning for recurring concept drifting data streams," *Expert Systems with Applications*, vol. 213, p. 118934, 2023.

[22] H. Shahparast, S. Hamzeloo, and E. Safari, "An incremental type-2 fuzzy classifier for stock trend prediction," *Expert Systems with Applications*, vol. 212, p. 118787, 2023.

[23] I. Žliobaitė, A. Bifet, J. Read, B. Pfahringer, and G. Holmes, "Evaluation methods and decision theory for classification of streaming data with temporal dependence," *Machine Learning*, vol. 98, no. 3, pp. 455–482, 2015.

[24] M. L. Prado, *Advances in Financial Machine Learning*. John Wiley & Sons, 2018.

[25] A. Bifet, G. Holmes, R. Kirkby, and B. Pfahringer, "MOA: Massive online analysis," *Journal of Machine Learning Research*, vol. 11, p. 1601–1604, 2010.

[26] H. M. Gomes, A. Bifet, J. Read, J. P. Barddal, F. Enembreck, B. Pfahringer, G. Holmes, and T. Abdessalem, "Adaptive random forests for evolving data stream classification," *Machine Learning*, vol. 106, no. 9, pp. 1469–1495, 2017.

[27] A. Bifet, G. Holmes, and B. Pfahringer, "Leveraging bagging for evolving data streams," in *Machine Learning and Knowledge Discovery in Databases*, ser. LNCS. Springer, 2010, vol. 6321, pp. 135–150.

[28] R. S. M. Barros, S. G. T. C. Santos, and P. M. Gonçalves Júnior, "A boosting-like online learning ensemble," in *Proceedings of IEEE International Joint Conference on Neural Networks (IJCNN)*, Vancouver, Canada, 2016, pp. 1871–1878.