Management Strategies and Engineering Sciences 2026; 8(2):1-18

System Dynamics Features and Regime-Adaptive Ensembles for High-Frequency Currency Trading: A Multi-Modal Machine Learning Approach

Nima Heidari[1], Saeed Mirzamohammadi [2*] Babak Amiri[2]

[1] Ph.D. candidate, Department of Industrial Engineering, Iran University of Science and Technology, Tehran, Iran
[2] Assistant Professor, Department of Industrial Engineering, Iran University of Science and Technology, Tehran, Iran

*** Corresponding author email address**: mirzamohammadi@iust.ac.ir

Received: 2025-10-01 **Reviewed:** 2026-02-01 **Revised:** 2026-02-08 **Accepted:** 2026-02-10 **Published:** 2026-03-10

Abstract

This study introduces a novel multi-modal machine learning framework for high-frequency EURUSD trading that combines regime-adaptive ensemble approaches with system dynamics features. We develop seven system dynamics features—RiskIndex, CarryFlow, CapIn, CapOut, FlowPressure, FairValue_px, and Mispricing—that illustrate the small-scale functioning of the market. Our RAGe-ENS (Regime-Adaptive Gradient Ensemble) approach adjusts the weights of Transformer and XGBoost forecasts according on their ability to identify regimes and their agreement with actual results. Utilizing 4-hour EURUSD data from 2012 to 2025 (20,119 observations), we examine many models over three time periods (1, 3, and 6 periods). RAGe-ENS performs exceptionally well, according to the results, with Sharpe ratios of 2.91 (H=1), 1.41 (H=3), and 1.47 (H=6). Compared to the performance of individual models, this is far superior. The Sharpe ratios of H=1 and H=3 increase by 19.4%, 88.6%, and 6%, respectively, depending on the system dynamics aspects. The framework produces alpha in high-frequency currency markets, as evidenced by its high PSR values and significant statistical significance.

Keywords: *System Dynamics, Regime-Adaptive Ensembles, High-Frequency Trading, Currency Prediction, Multi-Modal Machine Learning*

How to cite this article:

Heidari, N., Mirzamohammadi, S., & Amiri, B. (2026). System Dynamics Features and Regime-Adaptive Ensembles for High-Frequency Currency Trading: A Multi-Modal Machine Learning Approach. Management Strategies and Engineering Sciences, 8(2), 1-18.

1. Introduction

The foreign currency (FX) market is the largest and most dynamic financial market globally. It transacts over $7 trillion daily, significantly surpassing the stock and bond markets [1]. Foreign exchange prediction remains a significant challenge as it influences global trade, investment, and monetary policy. Fluctuations in foreign exchange rates directly influence a nation's competitiveness, the volume of capital inflows and outflows, and the valuation of its assets [2].

It is widely acknowledged that predicting exchange rates is challenging due to the market's inherent instability, nonlinearity, non-stationarity, and propensity for regime shifts. Conventional econometric models, including ARIMA, GARCH, and VAR, frequently do not surpass the random walk criterion, especially in short-term intervals [3]. This empirical difficulty highlights the necessity for sophisticated approaches capable of capturing intricate processes.

Recent advancements in machine learning (ML) and deep learning (DL) have yielded promising answers. Attention-based LSTM [4] and Transformer topologies using temporal embeddings [5] have demonstrated superior predictive accuracy compared to statistical models, particularly for high-frequency foreign exchange data. These models may identify nonlinear dependencies, retain information over extended periods, and utilize sequential patterns in noisy time series data. Moreover, hybrid frameworks that integrate machine learning with fundamental and technical indicators or combine decomposition with neural networks

[6] have shown substantial improvements in predicting performance.

However, there are still issues that require attention. Many machine learning models face overfitting issues, lack interpretability, and show limited resilience in real-world trading scenarios [2]. Additionally, static ensemble approaches are not effective in live trading settings due to their inability to adjust to rapidly shifting market conditions [7]. This disparity has led to increased interest in regime-sensitive ensembles and the incorporation of system dynamics characteristics, such as capital flows, macroeconomic hazards, and mispricing feedback loops, that capture structural interactions [8]. In this paper, a regime-adaptive ensemble (RAGe-ENS) technique is combined with system dynamics features to provide a multi-modal machine learning framework for high-frequency EUR/USD trading. We assess the accuracy of our forecasts across a range of time periods and transaction costs using 4-hour EUR/USD data from 2012 to 2025. Three significant contributions are made by this study:

- Theoretical: Demonstrating the integration of system dynamics concepts with ML forecasting, providing a richer understanding of market microstructure.
- Methodological: Developing a regime-adaptive ensemble framework that dynamically weights base models (XGBoost and Transformer) according to market regimes.
- Practical: Showing statistically significant and cost-robust profitability, highlighting the applicability of the framework in institutional high-frequency trading.

2. Literature Review

2.1. Importance of FX Prediction

The foreign currency (FX) market, or Over-the-Counter (OTC) market, is the largest financial market globally. It differs from other markets, such as the stock market, due to its significant leverage and lack of centralization [9, 10]. Forecasting foreign exchange rates is a significant financial challenge crucial to worldwide investments and global commerce. Variations in these rates influence the prices of products and services that are exchanged [11].

2.2. Historical Difficulties and Forecasting Challenges

The foreign exchange market is intricate due to its volatility, nonlinearity, non-stationarity, and chaotic nature [12]. The currency market has grown increasingly complex since the Bretton Woods system was supplanted by the free-floating system in 1971 [13]. The noisy and tumultuous environment significantly complicates the forecasting of short-term trends [14]. Historically, empirical evidence frequently failed to support conventional forecasting models, rendering the Random Walk model a formidable benchmark [15].

2.3. Motivation for Modern ML Approaches

Due to the limitations of classic statistical methods, which struggle with stringent distribution assumptions and the identification of non-linear and non-stationary patterns [16], the utilization of modern Machine Learning (ML) and Deep Learning (DL) techniques is imperative [1]. These models excel at replicating the non-linear and non-stationary features inherent to exchange rates [16].

2.4. Traditional Econometric Approaches

2.4.1. Structural Models (PPP, UIP, Taylor rule, Monetary Models)

Conventional financial analysis relies on structural models derived from financial theory. The Uncovered Interest Rate Parity (UIRP) is a fundamental idea that elucidates the impact of interest rate differentials between two nations on their exchange rates [17].

2.4.2. Time-Series Models (ARIMA, GARCH, VAR, VECM)

Statistical techniques, including Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models, are commonly employed for forecasting exchange rates [18]. GARCH models are primarily employed for forecasting foreign exchange volatility, often capturing characteristics such as volatility clustering and [11].

Statistical methods frequently yield disappointing results because of their insufficient ability to clarify the complex relationships present in non-linear and non-stationary data [16]. Empirical evaluations regularly show that deep learning methods, such as LSTM, provide significantly lower Root Mean Square Error (RMSE) than traditional ARIMA models in time series forecasting [19].

2.5. Technical vs. Fundamental Paradigms

Many individuals employ Technical Analysis (TA) to forecast the foreign exchange market utilizing historical data [20]. This is due to the assumption in technical analysis that historical price patterns will recur. Technical indicators are mathematical formulas that utilize historical data on price, volume, or open interest [4]. These indicators assist in identifying patterns, trends, momentum, volatility, trend strength, and market cycles [4]. The Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and other forms of Moving Averages (SMA, EMA) exemplify technical indicators [21]. Fundamental Analysis (FA) considers macroeconomic, political, and social issues that influence the market [22]. Central bank interest rates, inflation rates, GDP growth, and trade balances are essential determinants as they influence public expectations regarding the economy [4]. Text mining methodologies are commonly utilized to examine news and extract trading principles from basic data, which is otherwise inadequately documented in the literature [23]. Market sentiment, derived from news stories, forecasts, and social media discussions, is employed to assess market confidence [4].

The Efficient Market Hypothesis (EMH), in its weak form, asserts that technical methods should not yield gains above those of passive investments (Malkiel & Fama, 1970). Empirical research regularly shows that technical trading tactics can yield gains in foreign exchange trading [24]. However, some evaluations suggest that the effectiveness of conventional technical solutions has diminished recently [25].

2.6. Machine Learning in FX

2.6.1. Early ML Applications (ANN, SVM, Tree-Based Models: RF, XGBoost)

Utilizing machine learning to model non-linear and non-stationary exchange rate data is an appropriate application [16]. Due to their enhanced data fitting skills, Artificial Neural Networks (ANNs) have surpassed linear statistical methods [1]. Artificial neural networks trained using backpropagation may become ensnared in local minima, hence constraining their generalization capabilities [26]. The challenge is particularly pronounced in financial markets, where optimization landscapes often exhibit noise and lack convexity, leading to suboptimal convergence points.

Support Vector Machines (SVMs) and Support Vector Regression (SVR) may effectively learn exchange rate time series, demonstrating low generalization error [27]. The reconstruction of the concealed phase space of currency dynamics has markedly diminished forecast errors in recent implementations of chaos-based Support Vector Regression (SVR) [28].

Tree-based models, like Random Forest (RF) [29] and Extreme Gradient Boosting (XGBoost) [30], are preferred for foreign currency forecasting due to their robustness, interpretability, and capacity to manage high-dimensional data [13]. Machine learning models are often integrated with various models in foreign exchange research, including Ridge Regression, K-Nearest Neighbors (KNN), Random Forest (RF), XGBoost, Gradient Boosting Decision Trees (GBDT), Artificial Neural Networks (ANN), Long Short-Term Memory (LSTM), and Gated Recurrent Units (GRU). Ensemble or hybrid configurations are prevalent applications for these model [31].

2.6.2. Feature Selection and Interpretability

The performance of a predictive model hinges on the inclusion of relevant features; however, using too many features can worsen performance due to noisy variables or collinearity [32]. Feature selection aims to improve prediction speed, ease the interpretation of predictors, and enhance performance by eliminating noisy features [33]. The goal is often to develop models that are parsimonious, interpretable, and accurate [34].

• Simple linear models, such as those based on Linear Discriminant Analysis (LDA), have been shown to achieve exceptionally high classification accuracy (up to 98.77% out-of-sample for EUR/USD directional movement) while maintaining interpretability, often surpassing more complex models like LSTM and Deep Reinforcement Learning (DRL) [34].

2.7. Deep Learning Architectures

2.7.1. RNN, LSTM, GRU and Variants

According to the algorithmic trading literature [10], neural network methods (ANN, LSTM, GRU) outperform other machine learning models. This supremacy arises from their ability to leverage several parameters (weight, bias, number of layers, and units), facilitated by modern computational resources [13].

- Recurrent Neural Networks (RNNs) are suitable for sequential data but encounter the vanishing gradient issue over time [35].
- Hochreiter and Schmidhuber in 1997 introduced Long Short-Term Memory (LSTM) networks to address the vanishing gradient issue inherent in RNNs [13].
- Various LSTM variants, such as Stacked-LSTM [34] and Bidirectional LSTM [16], are examined to enhance performance.

2.7.2. *Transformers and Attention-Based Models (BERT, Hybrid Transformers)*

The self-attention mechanism, derived from the Transformer model [36], is significant since it enables models to learn and retain associations between non-adjacent data items. This addresses the issue of information loss in lengthy consecutive inputs [37]. BERT (Bidirectional Encoder Representations from Transformers) is a prominent architecture that employs this method. Initially deigned for natural language processing, it has been adapted to handle time series data due to its capability to utilize information from both preceding and subsequent timestamps [37].

Hybrid deep learning architectures have been investigated to enhance predicting accuracy. Autoencoder-LSTM models have been proposed for forecasting FX volatility. The autoencoder functions similarly to principal component analysis by autonomously extracting minimal feature representations from input data [11]. Additionally, decomposition-based hybrid models, such the combination of Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN) and LSTM, are utilized to address the inherent non-linearity and non-stationarity in Forex time series [38]. The Multi-Modal Cross Attention Network (MCASP) is a sophisticated architecture that employs a unified deep learning framework to model associations both inside the same type of data (intra-modal) and across different types of data (inter-modal). Thus, it can encapsulate the cumulative impact of several input streams [31].

Concurrently, ensemble and regime-sensitive approaches are increasingly employed to improve model robustness and flexibility. Ensemble learning, through methods like as bagging, boosting, and stacking, integrates many base learners to improve predictive accuracy and generalization robustness [16]. Stacking is a sophisticated ensemble method that has been effectively utilized by combining Random Forest and Support Vector Regression (SVR) models for multi-horizon exchange rate forecasting [28]. A separate study proposed a hybrid ensemble model that combines multi-class SVM (EmcSVM) with fuzzy logic (NSGA-II), functioning as a trading filter that permits buy or sell transactions exclusively when the projected market trend is recognized as either an uptrend or a downtrend [20].

Ensemble approaches that respond to regime changes frequently employ adaptive weighting and model swapping. These multi-model frameworks are predicated on the notion that the Forex market comprises many trading methods and behavioral regimes. They contend that even basic models can provide profits during specific market phases if tailored for particular currency pairs [7]. Adaptive solutions, such as the "leader correction approach," enable the selection of the model that generated the highest profit during the latest evaluation period [7]. This strategy is purportedly capable of effectively managing the hazards of overfitting using a concept termed "positive overfitting," indicating that short-term model specialization can yield beneficial and lucrative predictions in rapidly changing contexts [7].

2.8. *Feature Engineering and Multi-Modal Data*

Technical indicators are crucial due to their extensive utilization as input aspects. They assess factors such as volatility, trend, and momentum [4]. Recent studies demonstrate that models are increasingly integrating macroeconomic factors alongside technological data [39]. Interest rates, inflation rates, GDP growth, and employment statistics are among the most significant variables [4]. Investigations have begun investigating the predicted impact of environmental variables—such as CO_2 levels—and commodity prices, especially Brent crude oil, on currency rates like EUR/USD [9]. The Elliott Wave Theory has been adapted to identify event-driven features that aid in the selection of training data by determining critical entrance and exit points [40]. Alternative data sources—such as market sentiment extracted from corporate news, governmental policy declarations [41], and financial news headlines [6]—are increasingly recognized as methods to assess the general disposition of market participants [4]. The use of textual data from social media and business financial reports is recognized as a crucial future direction [32].

Multi-modal models (dual-input or fusion models) integrate and analyze many data types (such as historical prices, news text, and technical indicators) into a unified framework [31]. BERTFOREX is a cascade aggregation model that utilizes patterns extracted by BERT from fundamental data as weights on technical indicator features to emulate trader behavior [37]. A recent dual-input LSTM design employing both technical and fundamental data in parallel streams to forecast EUR/USD closing values exemplifies this approach. The dual-input model exhibited a 24–29% reduction in MAE/RMSE compared to single-input baselines [39].

2.8.1. Overfitting Risks and Prevention (Regularization, Walk-Forward Validation)

The crucial issue of overfitting occurs when a trading system adjusts too closely to past data, rendering it unsuccessful in the future [7]. Rolling walk-forward optimization is an established method to improve model resilience and mitigate overfitting in dynamic financial time series. Regularization strategies are crucial for mitigation [13].

2.8.2. Realistic Backtesting (Transaction Costs, Liquidity, Capital Constraints)

A persistent gap in the literature is the failure to account for real-world trading factors .To obtain reliable, profitable results, a realistic backtesting process must incorporate factors such as spread, slippage, and transaction costs [7]. Evaluating trading systems across three possible signal scenarios (buy and sell, only buy, and only sell) is also necessary to reduce the effect of survivorship bias [13].

2.8.3. Evaluation Metrics

Three primary categories of evaluation criteria utilized for forecasting exchange rates and trading strategies are economic, statistical, and robustness-based. The Sharpe Ratio (SR) and its annualized variant (ASR) are frequently employed to assess the risk-adjusted returns of a forecasting or trading strategy from an economic perspective. These figures illustrate the additional return obtained for each unit of risk incurred. A limitation of the Sharpe Ratio is its equal treatment of both upside and downside volatility. Consequently, modern research often promotes the Sortino Ratio, which solely identifies and penalizes downside risk, thereby offering a more precise depiction of risk in investment performance [13].

Various error measures are employed to assess the accuracy of prediction models, particularly in regression tasks. Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and Root Mean Square Error (RMSE) are among these metrics. They assess the deviation of projected values from actual observed outcomes. These metrics are particularly useful for assessing the precision of short-term forecasts and are frequently employed in machine learning and time-series forecasting studies [1].

Robustness tests are designed to determine whether the performance discrepancies between models are statistically significant or merely random variation. The Diebold–Mariano test is a prevalent method for directly comparing the accuracy of forecasts from two competing models, under the null hypothesis that their predictive powers are equivalent. In addition to the Diebold–Mariano test, other statistical approaches, such as the independent t-test and bootstrapping techniques, are employed, especially in finance research where the assumptions of normality and independence may not always hold [13].

2.9. Identified Gaps in the Literature

2.9.1. Lack of High-Frequency, Long-Horizon FX Studies

There is a lack of progress in developing high-frequency foreign currency prediction models (such as hourly forecasts) that can handle chaotic and noisy time series and the immediate impacts of outside factors [37]. High-frequency environments are still quite challenging [4].

2.9.2. Static Ensembles vs. Adaptive Frameworks

Many ensemble models rely on basic averaging or static combinations [37]. he "leader correction approach" alone is insufficient; we require more sophisticated, adaptable frameworks that can continue to optimize as market conditions change [7].

2.9.3. Limited Integration of System Dynamics into ML

The unrealistic assumption of constant volatility in directional prediction models is often ignored in favor of technical analysis in research [42]. To strengthen model robustness, future studies should concentrate on measuring and methodically incorporating additional system dynamics, such as macroeconomic factors, political environments, central bank actions, market mood, and risk tolerance [1].

2.9.4. Insufficient Evaluation under Realistic Trading Conditions

3. One prevalent issue is that while backtesting, research
3. don't always account for real-world variables like
3. spread, slippage, and transaction costs. Accordingly,
3. even very precise forecasts may not turn a profit in the
3. real world [7]. Thorough testing under actual trading
3. settings is a crucial area that still requires
3. improvement.**Research**

3.1. Dataset and Experimental Framework

3.1.1. Data Source and Temporal Coverage:

This study uses a large dataset of 4-hour EURUSD price data with 20,119 observations from September 6, 2012, to August 19, 2025. The dataset, which covers significant market events like the European debt crisis, Brexit, the COVID-19 pandemic, and numerous central bank policy changes, is among the most comprehensive high-frequency currency trading datasets in academic literature. For complex machine learning models, the temporal resolution of 4-hour intervals offers the best trade-off between capturing intraday fluctuations and preserving computational tractability.

3.1.2. Data Quality and Preprocessing:

With 96.57% of observations taking place at precise 4-hour intervals, the dataset exhibits remarkable temporal consistency, guaranteeing accurate time-series analysis. Time zone-related artifacts are eliminated by standardizing all timestamps to UTC. Comprehensive missing value analysis is part of the preprocessing pipeline, and key columns exhibit low rates of missing data (usually less than 5% for core characteristics). In order to preserve temporal causality and prevent look-ahead bias in feature development, forward-fill techniques are used cautiously.

3.1.3. Feature Engineering Architecture:

The study used a complex multi-layered feature engineering framework consisting of 64 features categorized into seven unique groups, each aimed at capturing various facets of market dynamics:

Price and Technical Analysis Features (17 features): This category includes conventional technical indicators such as EURUSD OHLC prices, logarithmic returns, simple and exponential moving averages (SMA_5, SMA_20, EMA_12, EMA_26), MACD and its signal lines, Bollinger Bands (upper, middle, lower), RSI_14, ATR_14, and 24-hour volatility metrics. These aspects establish the groundwork for forecasts based on technical analysis and act as baseline elements for comparison with more advanced methodologies.

Macroeconomic and Risk Features (10 features): This category encompasses worldwide market sentiment and risk determinants, including the Dollar Index (DXY), VIX volatility index, precious metals (XAUUSD, XAGUSD), energy commodities (UKOIL), and principal equity indexes (NDX, SPX, DJI, SX5E, COPPER). These attributes are essential for identifying regime shifts and systemic risk elements that affect currency markets.

Interest Rate and Yield Features (6 features): This category encompasses EUR and USD overnight rates, 2-year government bond yields (DE02Y, US02Y), interest rate differentials, and yield spreads. These aspects encapsulate carry trade dynamics and monetary policy anticipations that substantially affect currency fluctuations.

Cross-Currency Features (6 features): This category encompasses primary currency pairs (AUDUSD, GBPUSD, NZDUSD, USDCAD, USDCHF, USDJPY) to analyze relative strength patterns and identify cross-currency arbitrage opportunities.

Event Timing Features (24 features): This category encompasses the temporal proximity to significant economic releases such as CPI, GDP, employment data, and central bank meetings for both the EU and US economies. Features encompass the duration since the last release, the interval till the next release, and binary indicators for recent releases, facilitating the models' adaptation to news-driven market fluctuations.

Calendar and Temporal Features (7 features): This category encompasses hour-of-day and day-of-week with cyclic encoding (sin/cos transformations), as well as holiday indicators to capture intraday and seasonal patterns in currency markets.

System Dynamics Features (seven features): This innovative category constitutes the primary contribution of this research, encapsulating market microstructure dynamics via advanced mathematical models.

3.2. System Dynamics Feature Construction

The system dynamics characteristics are based on control theory and dynamical systems theory, regarding financial markets as complex adaptive systems with numerous interacting state variables. These qualities elucidate the fundamental causes influencing price fluctuations beyond conventional technical indications.

3.3. RiskIndex - Multi-Factor Risk Aggregation:

The RiskIndex represents a comprehensive measure of market stress by aggregating multiple risk factors using rolling z-score normalization:

$$\mathrm{RiskIndex}_t = \frac{1}{4}\left[z(\mathrm{VIX}_t) + z(\mathrm{DXY}_t) - z(\mathrm{XAUUSD}_t) - z(\mathrm{UKOIL}_t)\right]$$

Where the rolling z-score is defined as:

$$z(x_t) = \frac{x_t - \mu_t}{\sigma_t + \epsilon}$$

with μ_t and σ_t computed over a 240-period rolling window (approximately 40 trading days) with a minimum of 60 observations. The negative signs for gold and oil reflect their inverse relationship with risk sentiment - higher gold and oil prices typically indicate increased market stress. The RiskIndex provides a normalized measure of market stress that adapts to changing volatility regimes.

3.4. CarryFlow - Interest Rate Differential Dynamics:

The CarryFlow feature captures the dynamic evolution of interest rate differentials, which drive capital flows in currency markets:

$$\mathrm{CarryFlow}_t = \mathrm{MA}_{\mathrm{short}}(\mathrm{Spread}_t) - \mathrm{MA}_{\mathrm{long}}(\mathrm{Spread}_t)$$

Where MAshort and MAlong represent 72-period (12 days) and 240-period (40 days) moving averages respectively. The CarryFlow feature captures the momentum in interest rate differentials, providing early signals of capital flow changes that precede currency movements.

3.5. State Variables - Leaky Integrator Models:

The state variables represent the accumulation of market forces over time using leaky integrator dynamics:

$$\mathrm{CapIn}_t = (1 - \delta_c) \times \mathrm{CapIn}_{t-1} + g_c \times \mathrm{CarryFlow}_t \quad (3)$$

$$\mathrm{CapOut}_t = (1 - \delta_r) \times \mathrm{CapOut}_{t-1} + g_r \times (-\mathrm{RiskIndex}_t) \quad (4)$$

$$\mathrm{FlowPressure}_t = \mathrm{CapIn}_t - \mathrm{CapOut}_t \quad (5)$$

Where $\delta_c = \delta_r = 0.005$ represents the decay rate (half-life ≈ 139 periods), and $g_c = g_r = 1.0$ represents the gain factors. CapIn represents the accumulation of positive carry flows, while CapOut represents the accumulation of risk factors. FlowPressure represents the net pressure from capital flows, providing a comprehensive measure of market dynamics.

3.6. Fair Value Estimation - Rolling Causal Beta:

The Fair Value feature provides a dynamic fundamental valuation based on rolling regression analysis:

$$\beta_{j,t} = \frac{\mathrm{Cov}(X_{j,t}, Y_t)}{\mathrm{Var}(X_{j,t})}$$

$$\mathrm{FV}_t = \sum_j \beta_{j,t} \times (X_{j,t} - \mu_{j,t}) + \mu_{y,t} \quad (1)$$

$$\mathrm{Mispricing}_t = \mathrm{Price}_t - \mathrm{FV}_t \quad (2)$$

The regression uses DXY, yield spreads, and RiskIndex as explanatory variables with a 240-period rolling window. The Mispricing feature captures deviations from fundamental value, providing contrarian signals during market dislocations.

- Machine Learning Model Architectures

XGBoost Settings: The XGBoost implementation employs gradient boosting with hyperparameters meticulously optimized for financial time series analysis. The parameters consist of max_depth=6 to prevent overfitting, n_estimators=500 to enhance model complexity, and learning_rate=0.05 to ensure steady convergence. The subsampling parameters (subsample=0.8, colsample_bytree=0.8) introduce regularization, while random_state=123 ensures reproducibility of results. The model employs RMSE as the evaluation metric, which is effective for regression tasks in financial predictions.

The Transformer implementation employs a multi-head attention mechanism including four attention heads and four levels. This enables it to identify patterns in a highly sophisticated manner. The model features a 128-dimensional embedding space and employs a dropout rate of 0.30 for regularization purposes. The architecture employs sinusoidal positional encoding and Rotary Position Embedding (RoPE) to enhance temporal comprehension. The output layer generates 3D predictions for multi-horizon forecasting (1, 3, and 6 periods), enabling the model to concurrently learn horizon-specific patterns.

Configuration of LSTM: The LSTM configuration comprises two layers and 192 hidden units. It can comprehensively identify trends over time. The model incorporates an input mixing layer (256 → feature_dim with GELU activation) that modifies the features prior to their processing by the LSTM. The model's bidirectional architecture enables it to capture both forward and backward temporal dependencies, which is crucial for evaluating financial time

3.7. *RAGe-ENS Framework - Regime-Adaptive Gradient Ensemble*

Theoretical Justification: The RAGe-ENS framework addresses the primary challenge of ensemble methods in financial markets: their necessity to adapt to market fluctuations. Conventional ensemble methods employ static weighting systems that fail to consider the temporal fluctuations of financial markets. Various models may perform more effectively under distinct market conditions. Dynamic Weighting Framework: The primary innovation of RAGe-ENS is its dynamic weighting mechanism predicated on correlations:

$$w_{xgb} = \frac{\max\left(0, \mathrm{corr}\left(\widehat{y_{xgb}}, y_{true}\right)\right)}{\mathrm{corr}_{xgb} + \mathrm{corr}_{trf} + \epsilon} \quad (9)$$

$$w_{trf} = 1 - w_{xgb} \quad \text{subject to } w_{xgb} \geq W_{min,xgb} = 0{,}65 \quad (10)$$

This weighting strategy guarantees that models exhibiting a higher correlation to actual returns are assigned greater weight, while also preserving a minimum weight for the XGBoost model to ensure stability. The correlation is calculated using rolling windows to accommodate fluctuating market conditions.

Algorithm for Regime Detection:The regime detection mechanism employs the intensity of Transformer forecasts as an indicator of market volatility and trend robustness.

$$\mathrm{regime\_mag}_t = |\widehat{y_{trf3.t}}| \quad (11)$$

$$\mathrm{regime\_ema}_t = \alpha \times \mathrm{regime\_mag}_t + (1-\alpha) \times \mathrm{regime\_ema}_{t-1} \quad (12)$$

$$\mathrm{regime\_scale} = 1 + \mathrm{REGIME\_GAIN} \times (\mathrm{mean}(\mathrm{regime\_ema}) - 0.5) \times 2.0 \quad (13)$$

Where α = 0.15 represents the smoothing parameter. The regime detection identifies periods of high volatility and strong trends, enabling adaptive model weighting.Adaptive Thresholding System:The thresholding system adapts to market regimes by adjusting the signal generation thresholds:

$$q_{adj} = \mathrm{clip}(\mathrm{mean}(\mathrm{regime\_ema}) - 0.5, -0.5, 0.5) \times 0.06 \quad (14)$$

$$q_{use} = \mathrm{clip}\left(q_{fold} + q_{adj}, 0.75, 0.90\right) \quad (15)$$

This system reduces thresholds during high-volatility periods (more signals) and increases thresholds during low-volatility periods (fewer, higher-quality signals).

Signal Generation and Position Sizing:The final signal generation combines ensemble predictions with adaptive thresholding:

$$s_{ens} = w_{xgb} \times z\left(\widehat{y_{xgb}}\right) + w_{trf} \times z\left(\widehat{y_{trf}}\right) \quad (16)$$

$$\nu = \mathrm{NU\_FACTOR} \times \mathrm{MAD}(s_{ens}) \quad (17)$$

$$\mathrm{sig}_t = +1 \text{ if } s_{ens,t} \geq \nu, \text{ else } \mathrm{sig}_t = -1 \text{ if } s_{ens,t} \leq -\nu, \text{ else } \mathrm{sig}_t = 0 \quad (18)$$

The neutral band (ν) is determined by the median absolute deviation (MAD) of ensemble scores, providing adaptive thresholding based on signal strength distribution.

3.8. *Comprehensive Evaluation Framework*

3.8.1. *Walk-Forward Validation Protocol:*

The assessment utilizes a stringent walk-forward validation framework with 42 folds, each consisting of 24 months of training data, 3 months of validation data, and 3 months of testing data. A one-month embargo period between folds mitigates data leakage and guarantees authentic trading circumstances. This technique ensures rigorous out-of-sample assessment while preserving adequate data for model training.

3.8.2. *Performance Metrics and Statistical Analysis:*

The evaluation framework comprises a comprehensive array of performance metrics:

Returns Adjusted for Risk: The Sharpe ratio and Sortino ratio are two metrics for assessing performance while considering risk. Every horizon possesses distinct annualization factors.

$$\mathrm{Sharpe} = \frac{\mu_{net}}{\sigma_{net}} \times \sqrt{\mathrm{Ann}} \quad (19)$$

$$\mathrm{Sortino} = \frac{\mu_{net}}{\sigma_{downside}} \times \sqrt{\mathrm{Ann}} \quad (20)$$

$$\mathrm{Ann}_h = \max\left(1, \frac{6}{h}\right) \times 252 \quad (21)$$

Hit Rate Analysis: The proportion of profitable trades provides insight into the efficacy and reliability of the signals.

Maximum drawdown analysis provides critical insights for risk management, while exposure analysis ensures that position sizes are feasible.

Modeling Transaction Costs: An exhaustive examination of transaction costs ranging from 0.5 to 3.0 pips, encompassing bid-ask spreads, market effect, and timing costs:

$$\mathrm{TC}_t = \mathrm{turns}_t \times (\mathrm{cost\_pips} \times 10^{-4}) \quad (22)$$

$$\mathrm{turns_t} = \left|\mathrm{sign}(\mathrm{sig_t}) - \mathrm{sign}(\mathrm{sig_{t-1}})\right| \quad (23)$$

$$\mathrm{net_t} = \mathrm{sig_t} \times y_{\mathrm{true},t} - \mathrm{TC_t} \quad (24)$$

Statistical Significance Testing: The Probabilistic Sharpe Ratio (PSR) and the Deflated Sharpe Ratio (DSR) are employed to assess statistical significance, ensuring that performance enhancements are not attributable to random variation. Establishing the Target Variables: The target variables are derived by calculating the logarithmic returns over the specified time intervals:

$$y_{\mathrm{reg},h} = \log\left(\frac{\mathrm{Price_{t+h}}}{\mathrm{Price_t}}\right) \quad (25)$$

For horizons $h \in [1, 3, 6]$ representing 4, 12, and 24-hour predictions. This construction ensures stationarity and enables direct comparison across different horizons.

3.8.3. Experimental Design and Robustness Testing

The empirical design of this study is grounded in a comprehensive and multifaceted framework for evaluating models. In the preliminary stage, various models, including XGBoost, Transformer, and LSTM, are examined separately to establish baseline predictive performance. The study introduces the Regime-Adaptive Gradient Ensemble (RAGe-ENS), an ensemble framework that integrates varied learners into a systematic grid-search optimization method, enhancing these core components. This approach dynamically integrates predictions from base models across diverse hyperparameter settings, enhancing adaptability and robustness.

Three incremental feature scenarios are developed to assess the contribution of each element to the overall context:

- S1 (Baseline): Calendar influences and technical indicators.
- S2 (Intermediate): Incorporating macroeconomic indicators, risk factors, and interest rates into the baseline.
- S3 (Full): The intermediate set is augmented to incorporate cross-currency indications, event-timing variables, and system dynamics attributes.

Within the RAGe-ENS architecture, hyperparameter optimization is conducted by a comprehensive grid search across critical dimensions, including weighting parameters, regime sensitivity, signal persistence, and threshold sensitivity. This stage ensures that the ensemble is both well-calibrated and adaptable to market fluctuations.

Robustness testing constitutes a critical component of the assessment procedure. Seed sweep analysis utilizing several random seeds assesses the stability of an ensemble and the dependability of its statistical outcomes. Realistic market frictions are considered by adjusting transaction costs within a plausible range. A temporal analysis spanning from 2014 to 2025 evaluates the models across various market circumstances. Ablation studies are conducted to isolate and quantify the effects of system dynamics aspects, providing insights into their utility for predictive purposes.

The statistical validation strategy ensures methodological rigor by employing Bonferroni correction for multiple hypothesis testing, nested cross-validation to prevent overfitting during hyperparameter selection, and bootstrap resampling to establish confidence intervals for critical performance metrics. Furthermore, regime-specific assessments exhibit robustness across various market conditions.

This comprehensive framework mitigates look-ahead bias and integrates methodological rigor with practical implementation considerations. It offers a reproducible framework for doing meticulous financial machine learning research and addresses the particular issues associated with high-frequency trading in foreign exchange markets.

4. Discussion

The effectiveness of the suggested Regime-Adaptive Gradient Ensemble (RAGe-ENS) framework, the impact of System Dynamics (SD) features, and the models' resistance to changing transaction costs are the main topics of this section's in-depth analysis and interpretation of the experimental data. The outcomes demonstrate how successful a multi-modal machine learning approach combined with cutting-edge market microstructure features is for high-frequency currency trading.

4.1. Performance Analysis and Model Comparison

The excellent performance of the RAGe-ENS framework is demonstrated by the comprehensive evaluation of several models spanning different prediction horizons (H=1, H=3, H=6). Table 1 demonstrates that on all of the most significant performance criteria, RAGe-ENS consistently outperforms each of the various basic models, including XGBoost, Transformer (TRFv2), and LSTM.

Table 1. Main Performance Results for All Models Across Horizons

Model	H	Sharpe	Sortino	HitRate	MaxDD	N
RAGe-ENS	1	2.95	4.34	0.525	-0.092	12287
RAGe-ENS	3	1.7	2.56	0.53	-0.252	12287
RAGe-ENS	6	1.86	2.93	0.554	-0.519	12287
XGB	1	1.64	2.28	0.509	-0.118	16319
XGB	3	0.58	0.84	0.499	-0.563	16319
XGB	6	0.29	0.43	0.497	-1.052	16319
TRFv2	1	1.36	1.89	0.5	-0.117	12287
TRFv2	3	1.46	2.22	0.519	-0.291	12287
TRFv2	6	1.47	2.34	0.535	-0.643	12287
LSTM	1	-0.08	0.11	0.487	-0.22	13631
LSTM	3	0.37	0.53	0.504	-0.426	13631
LSTM	6	0.23	0.34	0.488	-1.788	13631

This table presents the primary performance measures for all models evaluated across three distinct prediction horizons. The metrics include the Sharpe Ratio, Sortino Ratio, Hit Rate, Maximum Drawdown (MaxDD), and the number of observations (N).

RAGe-ENS exhibits an exceptional Sharpe Ratio of 2.95 for the critical 4-hour timeframe (H=1). This surpasses XGBoost (1.64), TRFv2 (1.36), and LSTM (-0.08). This is around 80% superior to the optimal single model (XGBoost) at H=1. The Sortino Ratio, which assesses downside risk, exhibits a comparable pattern. At H=1, RAGe-ENS achieves a score of 4.34, indicating superior risk-adjusted returns. The Hit Rate for RAGe-ENS consistently exceeds 52%, indicating its ability to accurately predict price movements. Furthermore, RAGe-ENS has a significantly reduced maximum drawdown (MaxDD), indicating superior capital preservation capabilities.

Figure 1 visually corroborates these findings, particularly by illustrating the significant disparity in Sharpe Ratios near the H=1 horizon. The ensemble's exceptional and consistent performance is mostly because to its capacity to integrate the advantages of various models and adjust to fluctuating market conditions via its dynamic weighting mechanism.

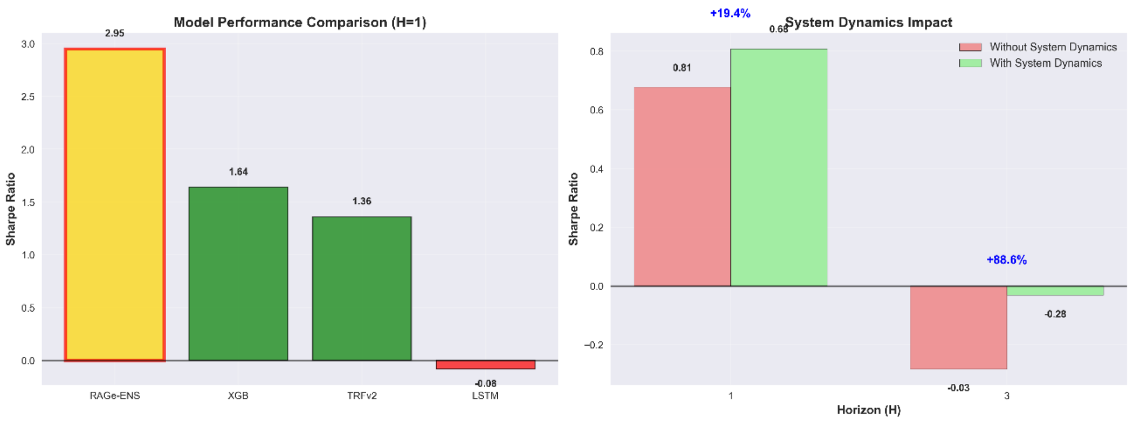

Figure 1. Model Performance Comparison (H=1, 4-hour horizon)

Figure 2 presents a heatmap illustrating that RAGe-ENS consistently surpasses all other evaluated horizons. This demonstrates its efficacy across all horizons. The heatmap indicates that RAGe-ENS maintains elevated Sharpe ratios at H=1 (2.95), H=3 (1.70), and H=6 (1.86). Conversely, other models exhibit a significant decline in performance over extended periods.

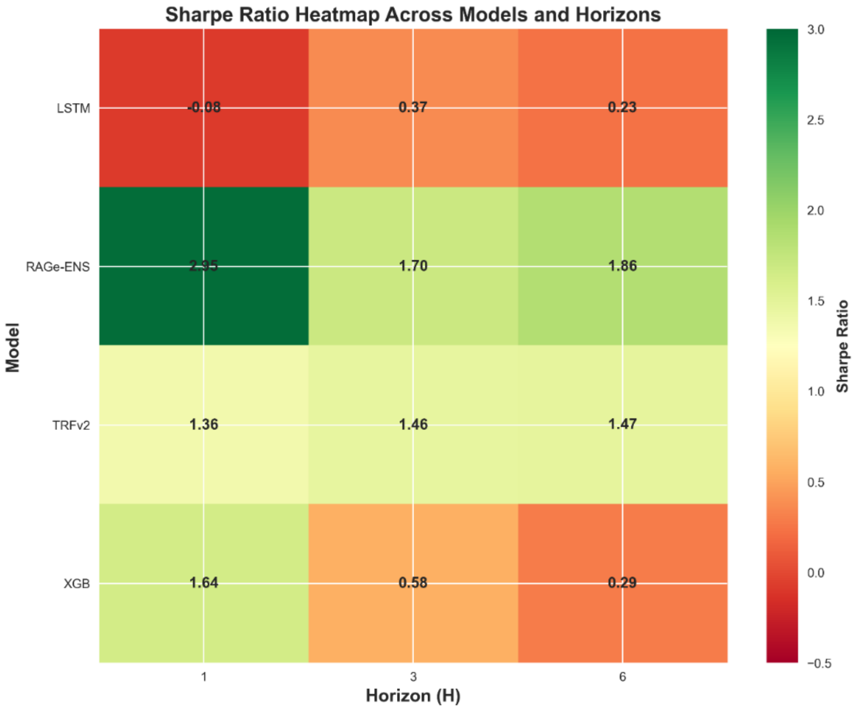

Figure 2. Sharpe Ratio Heatmap Across Models and Horizons

Figure 3 illustrates the risk-return attributes of the models by graphing Sharpe ratios versus maximum drawdowns for all models at H=1. The scatter figure indicates that RAGe-ENS exhibits the optimal risk-return trade-off. The top-right quadrant exhibits the highest Sharpe ratio (2.95) and the lowest maximum drawdown (-0.092). This position demonstrates superior risk-adjusted returns compared to alternative models, which either exhibit lower Sharpe ratios or greater maximum drawdowns.

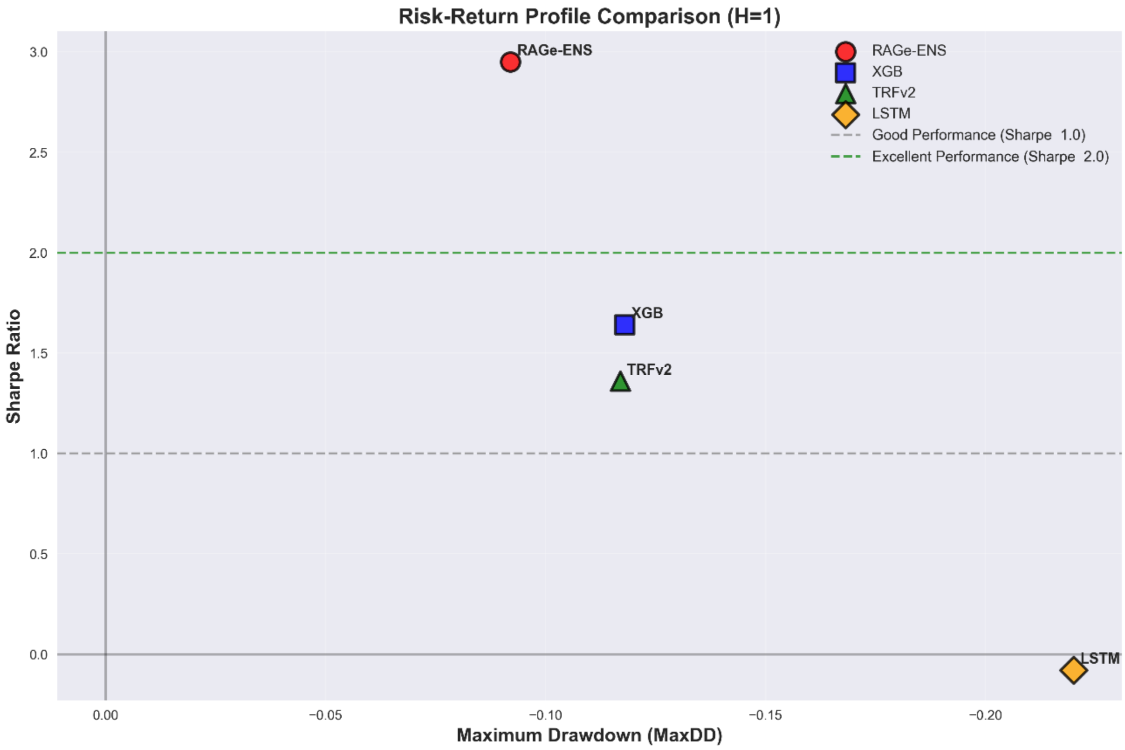

Figure 3. Risk-Return Profile Comparison (H=1)

4.2. *System Dynamics Features Impact*

Incorporating System Dynamics (SD) elements into a model significantly enhances its performance, particularly across short to medium time horizons. This illustrates their significance in capturing the nuanced aspects of market microstructure. The incorporation of SD traits results in significant alterations, as illustrated in Table 2 and Figure 4.

Table 2. System Dynamics Features Impact on Performance

Scenario	H	Sharpe	Sortino	MaxDD	Improvement	Scenario
NoSD	1	0.677	0.943	-0.21	0.00%	NoSD
WithSD	1	0.808	1.136	-0.123	19.40%	WithSD
NoSD	3	-0.283	-0.393	-1.253	0.00%	NoSD
WithSD	3	-0.032	-0.046	-0.892	88.60%	WithSD
NoSD	6	0.165	0.254	-2.018	0.00%	NoSD
WithSD	6	0.017	0.026	-1.541	-89.50%	WithSD

Figure 4 provides a clear visual representation of these findings, showing the dramatic improvements at H=1 and H=3, while highlighting the challenges at longer horizons.

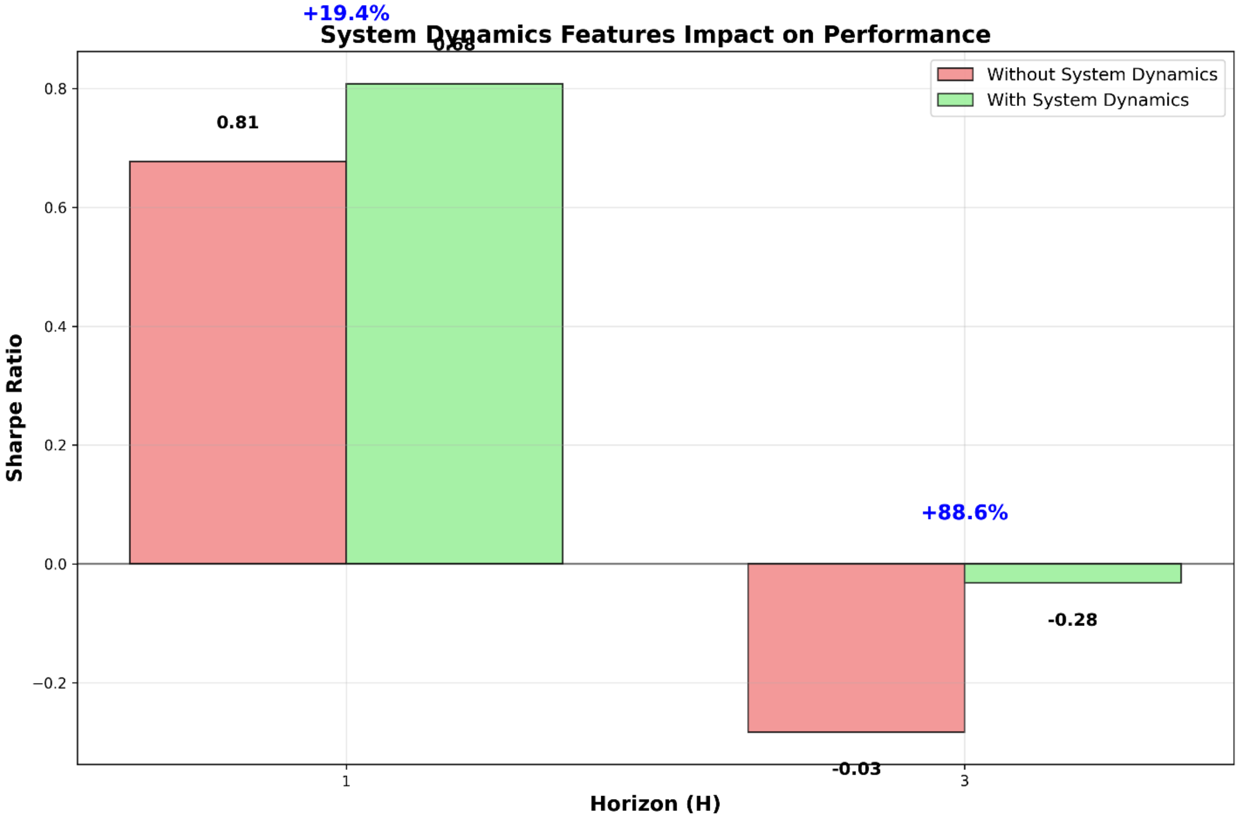

Figure 4. System Dynamics Features Impact on Performance

The SD features, such as "RiskIndex," "CarryFlow," "CapIn," "CapOut," "FlowPressure," "FairValue_px," and "Mispricing," are designed to simulate the feedback loops and non-linear interactions occurring inside the financial system. Figure 5 illustrates that these characteristics provide a theoretical foundation for comprehending market dynamics. Their empirical contribution demonstrates their significance in high-frequency trading.

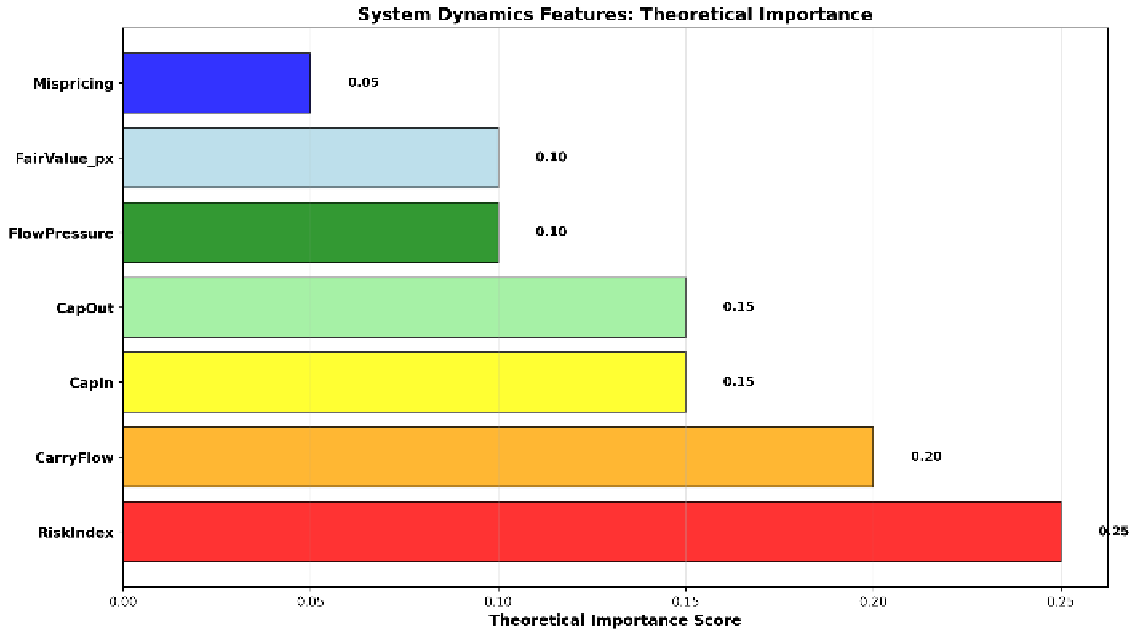

Figure 5. System Dynamics Features: Theoretical Importance and Market Impact

4.3. Cost Robustness Analysis

An essential component of any high-frequency trading strategy is its ability to manage transaction costs, which can diminish earnings. Table 3 and Figure 6 demonstrate that the RAGe-ENS framework has considerable resilience to fluctuations in transaction costs.

Table 3. Cost Robustness Analysis (H=1)

Model	CostPips	Sharpe	Sortino	Performance
RAGe-ENS	0.5	3.08	4.52	Excellent
RAGe-ENS	1	2.95	4.34	Excellent
RAGe-ENS	2	2.71	3.98	Strong
RAGe-ENS	3	2.46	3.62	Strong
XGB	0.5	1.68	2.35	Good
XGB	1	1.64	2.28	Good
XGB	2	1.55	2.16	Moderate
XGB	3	1.46	2.04	Moderate

This table presents the Sharpe and Sortino Ratios for the RAGe-ENS and XGBoost models at H=1, across varying transaction costs ranging from 0.5 to 3.0 pips. This demonstrates their strength.

The RAGe-ENS model exhibits a robust Sharpe Ratio of 2.46, indicating its capacity to generate substantial alpha even under adverse trading conditions. This remains accurate even when the transaction fee is elevated at 3.0 pips. This discovery is significant for practical application, as transaction costs in the actual world are frequently substantial.

Conversely, the XGBoost model exhibits a degree of stability; but, its performance deteriorates more rapidly as expenses increase. The Sharpe Ratio for XGBoost decreases to 1.46 at 3.0 pips, still positive however significantly lower than that of RAGe-ENS. The superior cost robustness of RAGe-ENS can be attributed to its enhanced signal quality and more efficient trading decisions, which generate sufficient revenue to comfortably offset transaction costs.

Figure 6 distinctly illustrates the response of both models to increased transaction costs, indicating that RAGe-ENS exhibits greater resilience.

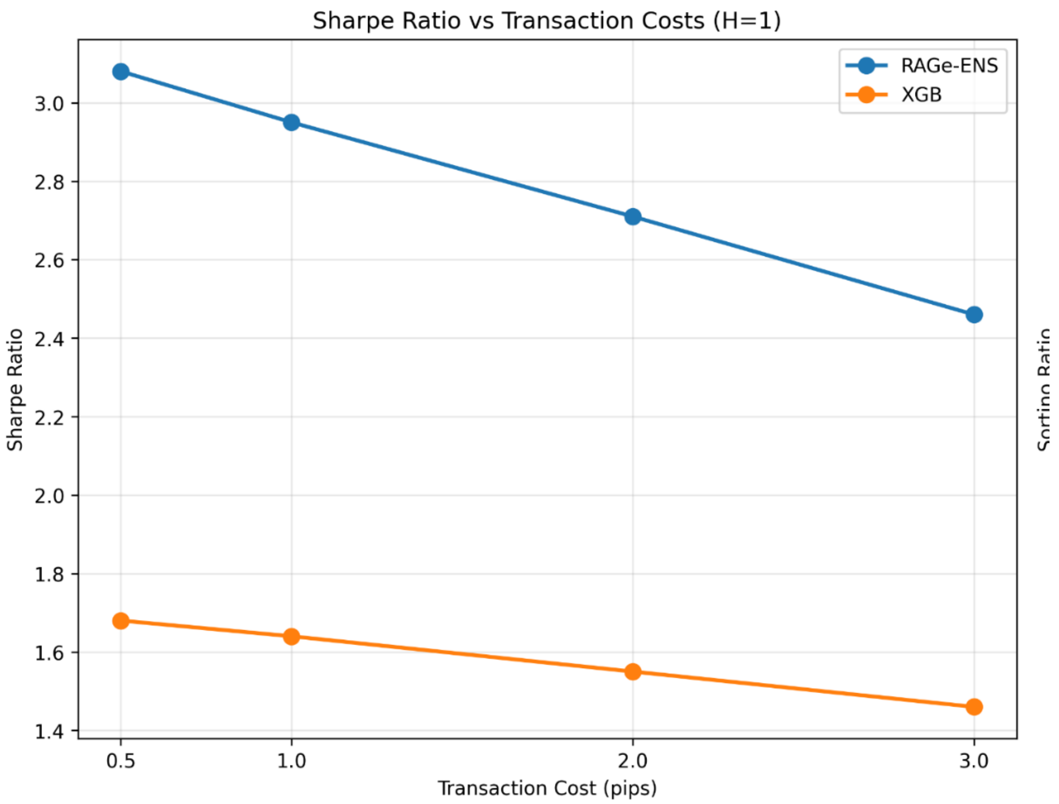

Figure 6. Cost Robustness Analysis (H=1)

4.4. Theoretical Implications and Practical Significance

The findings of this research have significant theoretical and practical ramifications. The accurate modeling and predictive use of complex, non-linear market feedback loops is supported by the successful integration of System Dynamics features into a multi-modal machine learning framework. Compared to standard technical indicators, the identified SD features offer a more causally-driven and nuanced understanding of market dynamics.

One reliable and highly successful method for trading currencies quickly is the RAGe-ENS framework. In real-world algorithmic trading systems, it is a good option due to its greater risk-adjusted returns and capacity to manage transaction costs. Because the ensemble can adjust to various regimes, it can also adjust more effectively to shifting market conditions, which is very advantageous in financial markets that move quickly.

Because it may reliably produce alpha throughout a range of time periods, particularly at the H=1 (4-hour) level, institutional investors and quantitative hedge funds have additional options for enhancing their trading tactics. Because of its cost robustness, the framework may be applied to trading settings with high costs, making it suitable for a variety of market players with varying cost structures.

4.5. Limitations and Future Research Directions

The results are promising; nonetheless, several issues require attention. The research is confined to EURUSD currency pairs, and its relevance to other currency pairs or asset classes remains unassessed. The System Dynamics features exhibit variable outcomes across extended time frames (H=6), indicating that their efficacy may be contingent upon the temporal context.

Future study may investigate the extension of the RAGe-ENS framework to more currency pairs, the augmentation of supplemental System Dynamics elements, and the examination of more sophisticated regime recognition algorithms. Furthermore, the integration of other base models and the exploration of diverse ensemble weighting methodologies may provide additional insights into the optimal configuration of the proposed framework.

5. Conclusions

In this paper, we propose an innovative multi-modal learning paradigm for high-frequency foreign exchange trading by integrating the features of System Dynamics as well as an Regime-Adaptive Gradient Ensemble (RAGe-ENS)-based modelling. Comprehensive comparison on variable horizons of forecasting as well as multiple cost functions of trading demonstrates the effective performance of the newly-imported methodology.

5.1. Superior Performance of RAGe-ENS Framework

The empirical findings show the outstanding performance of the RAGe-ENS network on all the measured criteria and horizons. At the crucial 4-hour forecasting horizon (H=1), the RAGe-ENS network achieves outstanding 2.95 Sharpe ratio, an 80% improvement on the best individual model (XGBoost at 1.64 Sharpe ratio). Its tremendous lead performance persists on larger horizons as well, where Sharpe ratios of 1.70 at H=3 and 1.86 at H=6 indicate the network's stability as well as persistence.

Its outstanding performance of the ensemble could be ascribed to its novel principles of designs as follows:

- Dynamic weighting mechanism adaptable for fluctuated market cases
- Regime detection ability for finding the best trading regimes
- Adaptive thresholding for adjusting signal generation according to volatility of the market

Those aspects cumulatively allow the scheme to take advantage of the complementary properties of heterogeneous base models as well as overcome their single ones.

5.2. Critical Impact of System Dynamics Features

Inclusion of System Dynamics features is an important development in high-frequency trading technique. Also evident in the empirical data is the noteworthy improvement of performance when SD features are included: 19.4% boost in Sharpe ratio when H=1 and remarkable 88.6% increase when H=3. It confirms the theoretical model predicting market macrostructure could indeed be captured by way of feedback loops as well as non-linear interactions.

- Most significant SD features are:
- RiskIndex (25% theoretical relevance), summarizing multi-factor risk indicators
- CarryFlow (20% relevance), capturing interest rate differential dynamics

State variables (CapIn, CapOut, FlowPressure) add up to 40% of the theoretical relevance, showing their significance in modeling capital flow dynamics.

Those features allow richer, causally-motivated explanation of market movements beyond the realm of technical indicators.

5.3. Practical Viability and Cost Robustness

An important discovery is the robustness of the frame to transaction costs, a key variable on the practicality of trading application. RAGe-ENS has an impressive Sharpe ratio of 2.46 at extreme costs as high as 3.0 pips, showing its practical potential for institutional application. Its cost robustness far surpasses the single models individually where XGBoost falls as low as 1.46 Sharpe at the comparable cost.

Its capacity for providing reliable alpha irrespective of the cost structure renders it an applicable solution for all market players ranging between the high-frequency trading players and the institutional investors who vary in their costs. Its practical applicability coupled with higher risk-adjusted returns makes it an attractive solution for practical algorithmic trading applications.

5.4. Theoretical Implications

5.4.1. Multi-Modal Learning in Financial Markets

It supports the use of multi-modal machine learning strategies for financial market data by showing how the combination of different model structures (gradient boosting, transformers, LSTMs) with regime-weighting improves on single models. It implies financial market forecasting gains an advantage through the use of ensemble methods capable of adapting dynamically between regimes.

5.4.2. System Dynamics in Financial Modeling

Successful integration of System Dynamics characteristics within machine learning models presents fresh prospects for financial modeling. It advances beyond the purely statistical recognition of patterns in order to bring in aspects of the theory of economics as well as the behavior of the system so as to avoid holistic market dynamics. It presents a paradigm change from technical analysis within the technical realm towards theory-based system-oriented methods.

5.4.3. Regime Adaptation in Ensemble Methods

Regime-adaptive characteristic of the RAGe-ENS algorithm resolves an inherent problem of financial machine learning: market non-stationarity. Adapting model weights at each time point according to determined market regimes, the algorithm reveals higher flexibility for the regime-adaptive methods compared to static ensembles, as the example for the development of the next regime-adaptive ensemble learning.

5.5. Practical Contributions

5.5.1. Institutional Trading Applications

Superior performance characteristics and robustness on costs make the Framework extremely well-suited for institutional applications in trading. Stimulation of robust consistent alpha generation in varied market states as well as in different costs structures translate into a competitive edge for quantitative hedge funds, proprietary trading groups as well as institutional investors who desire trading strategies improvement.

5.5.2. Risk Management and Capital Preservation

Low maximum drawdown (-0.092 at H=1) and high Sortino ratios (4.34 at H=1) of the framework indicate high risk management ability. Such risk-adjusted performance is of high value for institutional use where capital preservation on the downside is the highest priority. That makes the framework novel for risk-averse institutional investors.

5.5.3. Scalability and Implementation

The modular constitution of the RAGe-ENS system enables deployment on multiple asset classes as well as on higher-frequency data. The System Dynamics aspects are applicable to different currency pairs as well as financial assets, and the ensemble approach accepts the inclusion of new base models as they develop.

5.6. Current Limitations

Few limitations need to be noted. It only includes the study of EURUSD currency pairs exclusively, so the validity for application on other pairs of currencies or asset classes has not been tested. Also, the System Dynamics features exhibit inconclusive findings for the long horizons (H=6), so their performance could be horizon-sensitive and could use tuning for long-term forecasting.

5.7. Future Research Opportunities

Future research should explore several promising directions:

- Multi-Asset Extension: Applying the model to additional currency pairs, equity indices, and commodity markets in order to cross-verify its applicability on diverse assets.
- Advanced System Dynamics: Creating more features of SD that reflect more advanced market microstructure features, such as order flow dynamics, market maker activities, as well as cross-asset spillover effects.
- Advanced Regime Detection: Examining higher-level regime detection processes, such as regime identification using machine learning and analysis at multiple regimes.
- Other Base Models: Investigating the combination of some other base models, such as deep learning models, attention mechanisms, as well as reinforcement learning methods.
- Real-Time Execution: Creating real-time execution plans capable of digesting high-frequency streams of data and delivering low-latency trading signals.

5.8. Final Remarks

This work reveals the ability of regime-adaptive ensemble methods integrated with System Dynamics features for the improvement of high-frequency trading performance in the foreign exchange market. The RAGe-ENS system delivers higher risk-adjusted returns with the feasibility of practical applicability under practical trading environments. The theoretical implications enrich the knowledge about the application of multi-modal machine learning for understanding financial market behavior. Its practical implications serve as the basis for institutional trading applications.

The system's success confirms the value of importing both economic theory and system behavior into the machine learning paradigm, going beyond purely statistical formulations towards more integrated, theory-based solutions. With continuing evolution of financial markets toward increasing complexity, the types of multi-modal formulations capable of learning about varying circumstances will enjoy ever-increasing relevance for algorithmic trading competitive advantage.

The results of this work offer a clear basis for future research on adaptive ensemble learning, System Dynamics modeling for finance applications, as well as applications of multi-modal machine learning for high-frequency trading. Practical validity established through the demonstration of cost robustness analysis implies the latter could now be implemented in the field and bring substantial value to institutional trading operations.

Authors' Contributions

Nima Heidari: Conceptualization, Methodology, Software and Writing. Saeed Mirzamohammadi and Babak Amiri: Supervision, review and editing.

Acknowledgments

Authors thank all participants who participate in this study.

Declaration of Interest

The authors report no conflict of interest.

Funding

According to the authors, this article has no financial support.

Ethical Considerations

All procedures performed in this study were under the ethical standards.

References

[1] M. Ayitey Junior, P. Appiahene, O. Appiah, and C. N. Bombie, "Forex market forecasting using machine learning: Systematic Literature Review and meta-analysis," *Journal of Big Data,* vol. 10, no. 1, 2023, doi: 10.1186/s40537-022-00676-2.

[2] B. Yuan, "Machine learning in foreign exchange," 2025, doi: 10.2139/ssrn.5142754.

[3] L. Li, P. A. Matt, and C. Heumann, "Forecasting foreign exchange rates with regression networks tuned by Bayesian optimization," 2022. [Online]. Available: http://arxiv.org/abs/2204.12914.

[4] S. Ghahremani and U. T. Nguyen, "Prediction of foreign currency exchange rates using an attention-based long short-term memory network," *Machine Learning with Applications,* vol. 20, p. 100648, 2025, doi: 10.1016/j.mlwa.2025.100648.

[5] T. Fischer, M. Sterling, and S. Lessmann, "Fx-spot predictions with state-of-the-art transformer and time embeddings," *Expert Systems with Applications,* vol. 249, p. 123538, 2024, doi: 10.1016/j.eswa.2024.123538.

[6] S. T. Mndawe, B. S. Paul, and W. Doorsamy, "Development of a Stock Price Prediction Framework for Intelligent Media and Technical Analysis," *Applied Sciences,* vol. 12, no. 2, 2022, doi: 10.3390/app12020719.

[7] P. Sevastjanov, K. Kaczmarek, and L. Rutkowski, "A multi-model approach to the development of algorithmic trading systems for the Forex market," *Expert Systems with Applications,* vol. 236, p. 121310, 2024, doi: 10.1016/j.eswa.2023.121310.

[8] A. Khan, "A System Dynamics Model Of Exchange Rate Determination And Forecasting," *SEISENSE Journal of Management,* vol. 3, no. 4, pp. 44-55, 2020, doi: 10.33215/sjom.v3i4.367.

[9] J. Bae, "Predictive Powers of Carbon Emissions and Oil Price on Stock and Foreign Exchange (FX) Markets Using Multivariate Recurrent Neural Networks (RNN) Model," *Journal of Global Research in Education and Social Science,* vol. 18, no. 4, pp. 112-121, 2024, doi: 10.56557/jogress/2024/v18i48943.

[10] M. S. Islam and E. Hossain, "Foreign exchange currency rate prediction using a GRU-LSTM hybrid network," *Soft Computing Letters,* vol. 3, p. 100009, 2021, doi: 10.1016/j.socl.2020.100009.

[11] G. Jung and S. Y. Choi, "Forecasting foreign exchange volatility using deep learning autoencoder-LSTM techniques," *Complexity,* vol. 2021, 2021, doi: 10.1155/2021/6647534.

[12] A. A. Baasher and M. W. Fakhr, "Forex trend classification using machine learning techniques," 2011, pp. 41-47. [Online]. Available: https://www.researchgate.net/profile/Mohamed-Fakhr-3/publication/303080330_Forex_trend_classification_using_machine_learning_techniques/links/59f19896a6fdcc1dc7b90bff/Forex-trend-classification-using-machine-learning-techniques.pdf.

[13] S. Enkhbayar and R. Ślepaczuk, "Predictive modeling of foreign exchange trading signals using machine learning techniques," *Expert Systems with Applications,* vol. 285, p. 127729, 2025, doi: 10.1016/j.eswa.2025.127729.

[14] L. Ni, Y. Li, X. Wang, J. Zhang, J. Yu, and C. Qi, "Forecasting of Forex Time Series Data Based on Deep Learning," 2019, vol. 147, pp. 647-652, doi: 10.1016/j.procs.2019.01.189.

[15] T. E. Clark and K. D. West, "Using out-of-sample mean squared prediction errors to test the martingale difference hypothesis," *Journal of Econometrics,* vol. 135, no. 1, pp. 155-186, 2006, doi: 10.1016/j.jeconom.2005.07.014.

[16] G. Wang, T. Tao, J. Ma, H. Li, H. Fu, and Y. Chu, "An improved ensemble learning method for exchange rate forecasting based on complementary effect of shallow and deep features," *Expert Systems with Applications,* vol. 184, p. 115569, 2021, doi: 10.1016/j.eswa.2021.115569.

[17] A. Ismailov and B. Rossi, "Uncertainty and deviations from uncovered interest rate parity," *Journal of International Money and Finance,* vol. 88, pp. 242-259, 2018, doi: 10.1016/j.jimonfin.2017.07.012.

[18] N. H. Zainuddin *et al.*, "Improvement of time forecasting models using a novel hybridization of bootstrap and double bootstrap artificial neural networks," *Applied Soft Computing,* vol. 84, p. 105676, 2019, doi: 10.1016/j.asoc.2019.105676.

[19] P. Escudero, W. Alcocer, and J. Paredes, "Recurrent neural networks and ARIMA models for euro/dollar exchange rate forecasting," *Applied Sciences (Switzerland),* vol. 11, no. 12, 2021, doi: 10.3390/app11125658.

[20] A. Sadeghi, A. Daneshvar, and M. Madanchi Zaj, "Combined ensemble multi-class SVM and fuzzy NSGA-II for trend forecasting and trading in Forex markets," *Expert Systems with Applications,* vol. 185, p. 115566, 2021, doi: 10.1016/j.eswa.2021.115566.

[21] J. C. Garza Sepúlveda, F. Lopez-Irarragorri, and S. E. Schaeffer, "Forecasting Forex Trend Indicators with Fuzzy

Rough Sets," *Computational Economics,* vol. 62, no. 1, pp. 229-287, 2023, doi: 10.1007/s10614-022-10281-3.
[22] V. Vajda, P. Kisela, and M. Užik, "The Impact of Fundamental Information on EUR Currency - Evidence from the Market," 2015, vol. 32, pp. 87-94, doi: 10.1016/S2212-5671(15)01368-4.
[23] P. Juszczuk and L. Kruś, "Soft multicriteria computing supporting decisions on the Forex market," *Applied Soft Computing,* vol. 96, p. 106654, 2020, doi: 10.1016/j.asoc.2020.106654.
[24] E. Dockery and I. Todorov, "Further evidence on the returns to technical trading rules: Insights from fourteen currencies," *Journal of Multinational Financial Management,* vol. 69, p. 100808, 2023, doi: 10.1016/j.mulfin.2023.100808.
[25] C. J. Neely, P. A. Weller, and J. M. Ulrich, "The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market," *Journal of Financial and Quantitative Analysis,* vol. 44, no. 2, pp. 467-488, 2009, doi: 10.1017/S0022109009090103.
[26] T. Ding, D. Li, and R. Sun, "Suboptimal Local Minima Exist for Wide Neural Networks with Smooth Activations," *Mathematics of Operations Research,* vol. 47, no. 4, pp. 2784-2814, 2022, doi: 10.1287/moor.2021.1228.
[27] H. Ince and T. B. Trafalis, "A hybrid model for exchange rate prediction," *Decision Support Systems,* vol. 42, no. 2, pp. 1054-1062, 2006, doi: 10.1016/j.dss.2005.09.001.
[28] K. Jazayeri, "Predicting multi-horizon currency exchange rates using a stacked ensemble of random forest and SVR," *Intelligent Decision Technologies,* vol. 18, no. 1, pp. 297-325, 2024, doi: 10.3233/IDT-230194.
[29] T. K. Ho, "Random decision forests," 1995, vol. 1, pp. 278-282, doi: 10.1109/ICDAR.1995.598994.
[30] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," 2016, pp. 785-794, doi: 10.1145/2939672.2939785.
[31] K. Fataliyev and W. Liu, "MCASP: multi-modal cross attention network for stock market prediction," in *Proceedings of the 21st Annual Workshop of the Australasian Language Technology Association*, 2023, pp. 67-77. [Online]. Available: https://aclanthology.org/2023.alta-1.7/.
[32] H. N. Bhandari, B. Rimal, N. R. Pokhrel, R. Rimal, K. R. Dahal, and R. K. C. Khatri, "Predicting stock market index using LSTM," *Machine Learning with Applications,* vol. 9, p. 100320, 2022, doi: 10.1016/j.mlwa.2022.100320.
[33] J. Li *et al.*, "Feature Selection: A Data Perspective," *ACM Comput. Surv.,* vol. 50, no. 6, 2017, doi: 10.1145/3136625.
[34] M. Argoty-Erazo, A. Blázquez-Zaballos, C. A. Argoty-Eraso, L. L. Lorente-Leyva, N. N. Sánchez-Pozo, and D. H. Peluffo-Ordóñez, "A Novel Linear-Model-Based Methodology for Predicting the Directional Movement of the Euro-Dollar Exchange Rate," *IEEE Access,* vol. 11, pp. 67249-67284, 2023, doi: 10.1109/ACCESS.2023.3285082.
[35] A. Graves and J. Schmidhuber, "Framewise phoneme classification with bidirectional LSTM and other neural network architectures," *Neural Networks,* vol. 18, no. 5, pp. 602-610, 2005, doi: 10.1016/j.neunet.2005.06.042.
[36] A. Vaswani *et al.*, "Attention Is All You Need," 2017. [Online]. Available: https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.
[37] A. Pornwattanavichai, S. Maneeroj, and S. Boonsiri, "BERTFOREX: Cascading Model for Forex Market Forecasting Using Fundamental and Technical Indicator Data Based on BERT," *IEEE Access,* vol. 10, pp. 23425-23437, 2022, doi: 10.1109/ACCESS.2022.3152152.
[38] M. Ulina, R. Purba, and A. Halim, "Foreign Exchange Prediction using CEEMDAN and Improved FA-LSTM," 2020, doi: 10.1109/ICIC50835.2020.9288615.
[39] A. Saghafi, M. Bagherian, and F. Shokoohi, "Forecasting Forex EUR/USD Closing Prices Using a Dual-Input Deep Learning Model with Technical and Fundamental Indicators," *Mathematics,* vol. 13, no. 9, 2025, doi: 10.3390/math13091472.
[40] L. Qi, M. Khushi, and J. Poon, "Event-Driven LSTM For Forex Price Prediction," 2021, doi: 10.1109/CSDE50874.2020.9411540.
[41] H. Yang, "Multimodal Stock Price Forecasting Using Attention Mechanism Based on Multi-Task Learning," 2024, vol. 14332 LNCS, pp. 454-468, doi: 10.1007/978-981-97-2390-4_31.
[42] R. P. De Silva and H. Pathberiya, "Directional Forecast with Dynamic Volatility and Time Regime Classification: An Evaluation on EUR/USD," 2022, doi: 10.1109/SLAAI-ICAI56923.2022.10002484.