Journal of Emerging Economics and Business
Publisher by: **Suria Academic Press**
https://jurnal.suriaacademicpress.com/index.php/jeeb
e-ISSN: xxxx-xxxx
DOI: https://doi.org/xx.xxxx/jeeb.xxxx
Vol. 1, No. 1, December 2025, pp. 33–43

Predicting Cryptocurrency Price Direction Using a Hybrid LSTM Encoder and XGBoost Head: Implementation and Evaluation on Ten Major Digital Assets

Ade Kurniawan[1], Muliyono[2], Hari Suriadi[3]
[1]Universitas Putra Indonesia "YPTK" Padang, Indonesia
[23] Universitas Muhammadiyah Sumatera Barat, Padang, Indonesia
Email: ade.kurniawan.a17@gmail.com,[1] muliyono94@gmail.com[2], suriadihari6@gmail.com[3]

Open Access Article

Article History
Submission: 2025-12-1
Received: 2025-12-1
Revised: 2025-12-12
Accepted: 2025-12-30
Publish: 2025-12-30

Corresponding author:
Ade Kurniawan
ade.kurniawan.a17@gmail.com

ABSTRACT

The primary challenge in cryptocurrency price prediction lies in the market's highly volatile, non-linear, and near–random walk behavior, which makes traditional predictive models unable to achieve consistent accuracy. This study aims to develop and evaluate a hybrid model combining Long Short-Term Memory (LSTM) and XGBoost to predict price direction and returns for ten major cryptocurrencies using daily data from 2023 to 2025. Historical data were processed through feature engineering, normalization, and sliding-window sequence construction, and the models were evaluated using TimeSeriesSplit to prevent data leakage. The results show that the hybrid model consistently outperformed both LSTM and XGBoost, achieving an average directional accuracy of 58.6%, significantly higher than the baselines (51.7% for LSTM and 53.6% for XGBoost). The average RMSE of 0.0289 indicates stable return predictions without systematic bias. Statistical validation through paired t-tests and McNemar tests confirmed the significance of the improvement at $p < 0.001$. A trading simulation using a 1-day holding period produced an annualized return of 41.5% with a Sharpe ratio of 1.12, outperforming the buy-and-hold strategy. These findings highlight that integrating LSTM's temporal representation with XGBoost's non-linear learning capabilities is an effective and computationally efficient approach for cryptocurrency price forecasting, offering practical value for the development of algorithmic trading systems.

Keywords : Cryptocurrency; Price Prediction; Hybrid Model; LSTM; XGBoost; Algorithmic Trading

INTRODUCTION

The cryptocurrency market has evolved into a highly dynamic global financial ecosystem, characterized by a market capitalization exceeding two trillion US dollars and the growing participation of both institutional and retail investors (Corbet et al., 2018). Cryptocurrency price movements operate continuously for 24 hours without trading halts and largely outside strict market regulation, making them one of the most volatile classes of financial assets. These characteristics present significant challenges for analysts, researchers, and industry practitioners in producing accurate and stable price predictions.

Early studies on crypto assets suggest that cryptocurrency price changes tend to follow a random walk pattern, similar to the highly stochastic movements observed in stock markets (Nadarajah & Chu, 2017). Consequently, conventional statistical models such as ARIMA, GARCH, and VAR often struggle to capture the structural complexity and non-linear dynamics inherent in digital assets. High volatility, regime shifts, and substantial market noise make cryptocurrency price forecasting considerably more challenging than forecasting traditional financial assets.

To address the limitations of conventional statistical approaches, machine learning and deep learning techniques have increasingly been adopted in financial research, particularly for cryptocurrency price prediction. Among these, the Long Short-Term Memory (LSTM) model has become one of the most widely used approaches due to its ability to capture long-term dependencies in time-series data (Hochreiter & Schmidhuber, 1997). However, several studies indicate that LSTM still exhibits certain limitations in capturing complex non-linear patterns that are not entirely dependent on temporal structures, and it often suffers from overfitting when applied to highly volatile data.

On the other hand, tree-based models such as Extreme Gradient Boosting (XGBoost) have proven to be highly effective for tabular data and are capable of extracting non-linear interactions among variables in a more stable manner (Chen & Guestrin, 2016). This model excels in producing robust and efficient predictions; however, it is not inherently designed to learn temporal dependencies among observations directly. Consequently, the use of XGBoost as a standalone model in cryptocurrency price prediction has not yet achieved optimal performance.

The limitations of single-model approaches have led to the growing need for hybrid methods capable of combining the strengths of LSTM in capturing temporal information with the capability of XGBoost in handling complex non-linear structures. The concept of stacked generalization (Wolpert, 1992) provides a theoretical foundation suggesting that combining multiple heterogeneous models can produce more accurate predictions than relying on a single model. Nevertheless, such hybrid implementations remain relatively limited in comprehensive applications within cryptocurrency markets involving multiple assets with diverse characteristics.

A review of the existing literature reveals several shortcomings in previous studies. First, most hybrid LSTM–XGBoost studies have been conducted on a single asset primarily Bitcoin thus failing to test the generalizability of the model across other cryptocurrencies that exhibit different volatility characteristics (Silvanus et al., 2022). Second, many studies still employ inappropriate data validation techniques, such as random k-fold cross-validation, which can lead to data leakage and produce biased performance estimates (Bergmeir et al., 2018). Third, some studies do not include formal statistical testing, making it difficult to determine whether improvements in model performance are statistically significant from a scientific perspective.

In addition, previous studies have rarely provided a replicable end-to-end implementation framework, particularly with regard to data pipeline construction, technical feature selection, sliding window formation, and the integration of an LSTM encoder with an XGBoost meta-learner. The absence of a clearly structured implementation framework limits the reproducibility of existing research and constrains further development in both academic and industrial contexts. In particular, the availability of a well-defined framework is crucial for supporting subsequent research and practical applications, especially in the development of automated trading systems.

In response to this gap, the present study proposes a two-stage hybrid architecture that utilizes LSTM as a temporal encoder to extract latent representations from a 60-day price sequence window. These representations are then combined with several technical indicators, including Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and Bollinger %B, before being processed by XGBoost as a meta-learner. This approach is expected to address the limitations inherent in each individual model and produce more accurate and stable predictions of price direction.

Furthermore, this study implements TimeSeriesSplit validation with a forward-chaining scheme to ensure that no look-ahead bias occurs and to strengthen the robustness of predictions on unseen data. The evaluation is conducted on ten major cryptocurrencies based on market capitalization, allowing the results of this study to demonstrate not only model performance on Bitcoin but also cross-asset consistency across cryptocurrencies with varying volatility characteristics.

Accordingly, this study contributes to the literature in two primary aspects: (1) a methodological contribution, in the form of a structured hybrid model that is replicable and statistically validated; and (2) an empirical contribution, through consistent findings across multiple cryptocurrency assets. These contributions are expected to provide a solid foundation for future research and offer practical value for industry practitioners developing machine-learning-based automated trading systems.

LITERATURE REVIEW

Research on financial asset price prediction using machine learning and deep learning techniques has grown significantly in response to the increasing complexity of modern financial markets. The Long Short-Term Memory (LSTM) model, introduced by Hochreiter and Schmidhuber (1997), has become a foundational approach in time-series forecasting due to its ability to address the vanishing gradient problem and capture long-term dependencies in sequential data. In numerous financial studies, LSTM has demonstrated promising performance, particularly in modeling the non-linear price dynamics of stocks and cryptocurrencies. Nevertheless, several studies indicate that LSTM is highly sensitive to hyperparameter configurations, prone to overfitting, and may exhibit instability when applied to highly volatile data such as cryptocurrency markets.

On the other hand, research on tree-based models, particularly Extreme Gradient Boosting (XGBoost), has shown significant progress in the analysis of tabular data. Chen and Guestrin (2016) highlight that XGBoost excels in efficiently learning complex non-linear relationships through the integration of gradient boosting mechanisms and regularization techniques. In financial contexts, this model has been widely applied to identify technical patterns and market factors, largely due to its ability to capture intricate interactions among variables. However, a primary limitation of XGBoost lies in its inability to explicitly model temporal dependencies, which restricts its performance when dealing with datasets that strongly depend on time-order relationships.

Research on ensemble learning further suggests that combining multiple models with different characteristics often leads to improved predictive accuracy. Wolpert (1992), through the concept of stacked generalization, introduced a meta-learning approach in which a higher-level model (meta-learner) learns patterns from the outputs produced by base models. Subsequent studies by Kuncheva (2014) and Zhou (2012) reinforce the idea that model diversity within an ensemble, whether in terms of structure or learning mechanisms, significantly contributes to enhanced predictive performance. This theoretical perspective provides a strong foundation for developing hybrid models that integrate LSTM and XGBoost, leveraging the complementary strengths of both approaches.

Within the cryptocurrency domain, several studies have examined the performance of deep learning and machine learning models individually. Wang, Wang, and Zhang (2020) reported that LSTM can achieve relatively low Root Mean Square Error (RMSE) in Bitcoin price prediction; however, the directional accuracy of price movement remains limited at approximately 54%. Similarly, Permentier et al. (2024) applied a combination of several machine learning algorithms but obtained accuracy levels of only 52–56% across five cryptocurrency assets. Meanwhile, Silvanus et al. (2022) proposed a hybrid LSTM–XGBoost model for Bitcoin prediction and reported an RMSE reduction of approximately 15% compared with a standalone LSTM model. Nevertheless, their study did not include statistical significance testing and did not evaluate the model across multiple cryptocurrency assets.

The research gap becomes more evident when considering that most previous studies focus on a single asset typically Bitcoin thus limiting the ability to assess the generalizability of predictive models across other cryptocurrencies with different volatility characteristics. Furthermore, many studies fail to employ appropriate time-series validation techniques. Bergmeir, Hyndman, and Koo (2018) emphasize that the use of random k-fold cross-validation in time-series data can lead to data leakage, resulting in biased evaluation outcomes that do not accurately reflect a model's real-world predictive capability. Consequently, forward-chaining validation approaches are considered essential for time-series forecasting, although they are frequently overlooked in earlier research.

Recent developments in deep learning research also indicate a growing trend toward the use of more complex architectures, such as Transformer-based models and attention mechanisms. Zhang et al. (2023) found that deep learning ensembles could achieve prediction accuracies of approximately 60–62% in cryptocurrency forecasting, albeit with substantially higher computational costs. Similarly, Dantas et al. (2022) demonstrated that hybrid neural network ensembles can improve predictive accuracy; however, such models are often difficult to replicate and are frequently accompanied by insufficient methodological documentation. These findings highlight the importance of developing models that are not only accurate but also computationally efficient and practically implementable.

Based on the above literature review, several key limitations can be identified: (1) single models such as LSTM or XGBoost still exhibit limited predictive accuracy; (2) hybrid approaches have not been systematically applied across multiple cryptocurrency assets; (3) proper time-series validation methods are often neglected; (4) some studies do not include formal statistical testing; and (5) reproducible implementation frameworks remain scarce. Therefore, the present study aims to address these gaps by developing a hybrid LSTM–XGBoost model that is comprehensively evaluated across ten cryptocurrency assets using appropriate time-series validation methodologies and rigorous statistical analysis.

METHOD

This study adopts a quantitative approach with a predictive design to develop a machine learning model capable of projecting cryptocurrency price direction based on historical data. The selection of a predictive approach is grounded in the research objective, which focuses on the development and evaluation of a hybrid model based on LSTM and XGBoost. The study is conducted using an experimental framework in which the performance of baseline models is compared with that of the proposed hybrid model to determine its effectiveness and feasibility.

The data used in this research were obtained from Yahoo Finance, a platform that provides openly accessible cryptocurrency price data. Ten major assets were selected based on their highest market capitalization, namely Bitcoin (BTC), Ethereum (ETH), Binance Coin (BNB), Solana (SOL), Ripple (XRP), Cardano (ADA), Dogecoin (DOGE), Toncoin (TON), TRON (TRX), and Polkadot (DOT). The observation period spans November 2023 to November 2025 with a daily interval. Each asset contains approximately 730 observations, resulting in a total dataset of more than 7,000 data points. All data were carefully examined to ensure the absence of missing values or anomalies that could potentially affect model performance.

The data preprocessing stage was conducted through several key steps. First, return and log-return values were calculated to minimize the influence of price scale differences across assets. Next, several technical indicators, including Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and Bollinger %B, were computed to enrich the feature set, as these indicators have been widely recognized as relevant in previous technical analysis studies. Subsequently, all features were normalized using MinMax scaling to ensure stable training of the LSTM model. The final preprocessing step involved constructing 60-day sequential windows (sliding windows) to capture temporal patterns that serve as inputs for the model.

The hybrid model architecture consists of two main components. The LSTM network functions as a temporal encoder that generates latent representations of price patterns within a specific time window. These temporal representations are then combined with technical indicators and trading volume features, and subsequently processed by XGBoost as a meta-learner to predict the direction of the next-day price movement. This approach integrates the strengths of LSTM in capturing temporal dependencies with the advantages of XGBoost in extracting non-linear relationships among variables, in line with the concept of stacked generalization introduced by Wolpert (1992).

Model validation is performed using the TimeSeriesSplit technique with five folds, ensuring that the test data always occur chronologically after the training data. This validation approach is widely recommended in time-series analysis to prevent data leakage and improve the model's generalization capability (Bergmeir et al., 2018). Model performance is evaluated using several metrics, including accuracy, precision, recall, F1-score, and Root Mean Square Error (RMSE) to assess the quality of return prediction. To determine whether performance improvements are statistically significant, paired t-tests and McNemar tests are applied.

This study does not involve human participants or personal data and therefore does not require specific ethical approval. All research procedures were conducted in accordance with principles of scientific transparency, reproducibility, and open data utilization, allowing the findings of this study to be independently verified and replicated by other researchers.

RESULTS AND DISCUSSION

To obtain a comprehensive understanding of the effectiveness of the developed models, this study evaluates the performance of LSTM, XGBoost, and the hybrid model across ten cryptocurrency assets using test data under a forward-chaining validation approach. The evaluation focuses on several performance indicators, including directional prediction accuracy, Root Mean Square Error (RMSE) of return prediction, as well as classification metrics such as precision, recall, and F1-score, which are widely recognized as important measures for validating the performance of predictive models (Chollet, 2018; Goodfellow et al., 2016).

A summary of the evaluation results is presented in Table 1.

Table 1. Comprehensive Model Performance for Each Cryptocurrency (Test Set)

Cryptocurrency	LSTM Acc (%)	XGB Acc (%)	Hybrid Acc (%)	RMSE Reg	CV Mean (%)	Precision	Recall	F1-Score
BTC-USD	51.2	53.8	58.6	0.0261	57.1	0.61	0.65	0.63
ETH-USD	52.1	54.2	59.3	0.0287	58.9	0.63	0.67	0.65
BNB-USD	50.8	52.6	56.4	0.0342	55.2	0.57	0.61	0.59
SOL-USD	53.0	54.9	61.2	0.0198	61.0	0.65	0.71	0.68
XRP-USD	51.5	53.1	57.8	0.0305	56.9	0.59	0.63	0.61
ADA-USD	50.6	51.9	55.1	0.0368	53.8	0.56	0.59	0.57
DOGE-USD	52.3	54.5	60.1	0.0275	59.5	0.64	0.68	0.66
TON-USD	51.9	53.7	58.9	0.0298	57.4	0.60	0.64	0.62
TRX-USD	52.6	55.1	62.3	0.0225	61.8	0.66	0.72	0.69
DOT-USD	50.9	52.4	56.7	0.0331	55.3	0.58	0.62	0.60
Mean	**51.7**	**53.6**	**58.6**	**0.0289**	**57.7**	**0.61**	**0.65**	**0.63**

The results presented in Table 1 indicate that the hybrid model consistently outperforms the baseline models in predictive performance. On average, the hybrid model achieves an accuracy of 58.6%, representing a substantial improvement over LSTM (51.7%) and XGBoost (53.6%). This increase of 6.9 percentage points compared with LSTM and 5 percentage points compared with XGBoost highlights the importance of integrating temporal deep learning architectures with gradient-boosted decision trees. These findings are consistent with previous studies suggesting that combining deep learning models with tree-based learners can enhance generalization capability and reduce prediction errors in complex time-series data (Zhang et al., 2023; Kim & Kang, 2020).

An accuracy level exceeding 58% can be considered relatively high in cryptocurrency markets, given that numerous studies have shown that cryptocurrency price movements often approximate a random walk process (Urquhart, 2016; Corbet et al., 2019). Despite this inherent stochasticity, the hybrid model demonstrates the ability to capture non-linear patterns and temporal dependencies that are not fully captured by traditional models. The superiority of the hybrid approach becomes particularly evident in highly volatile assets such as TRX (62.3%) and SOL (61.2%), where volatility clustering patterns are more pronounced. In such cases, the LSTM network can extract richer temporal signals from the data. Conversely, model performance tends to decline for assets with relatively lower volatility, such as ADA and DOT. This finding suggests that volatility plays a crucial role in determining the effectiveness of predictive models, consistent with prior research highlighting the relationship between market volatility and predictive performance (Atsalakis & Valavanis, 2009).

From the perspective of return prediction, the average RMSE value of 0.0289 indicates that the hybrid model produces relatively stable forecasts without evidence of systematic drift. An RMSE value below 0.03 suggests that the model successfully captures short-term patterns in price dynamics without suffering from overfitting. This stability is partly attributed to the regularization mechanisms embedded in XGBoost, which control tree complexity and improve generalization (Chen & Guestrin, 2016). These findings support the view that combining sequential models with boosted tree models can produce more reliable quantitative estimates in highly dynamic markets (Fischer & Krauss, 2018).

To validate the statistical significance of the performance improvement, this study applies both paired t-tests and McNemar tests. The paired t-test yields a value of $t = 18.2$ with $p < 0.001$, indicating that the accuracy improvement of the hybrid model is statistically significant and not attributable to random variation. The McNemar test produces consistent results, with Hybrid vs. LSTM ($\chi^2 = 127.4$; $p < 0.001$) and Hybrid vs. XGBoost ($\chi^2 = 89.3$; $p < 0.001$). These outcomes reinforce the empirical evidence that the hybrid model not only demonstrates superior performance but also achieves statistically significant improvements, in line with recommended validation practices for financial prediction models (De Prado, 2018).

The feature importance analysis reveals that the latent vectors generated by the LSTM encoder contribute approximately 45% of the total feature weight, indicating that short-term temporal patterns play a critical role in the success of the prediction model. This finding aligns with theoretical perspectives suggesting that financial markets exhibit time-dependent structures, particularly in the short term, where momentum dynamics and trend reversals exert substantial influence (Lo, 2004). Technical indicators such as RSI contribute approximately 18%, supporting the literature that emphasizes the effectiveness of momentum indicators in volatile markets (Bollinger, 2002). Meanwhile, trading volume contributes around 11.7%, highlighting the continued importance of market liquidity variables, consistent with insights from financial market microstructure theory (Kyle, 1985).

To examine the practical relevance of the proposed model, this study conducts a trading strategy simulation based on a one-day holding period. The hybrid model achieves an annual return of 41.5% with a Sharpe ratio of 1.12, significantly outperforming the buy-and-hold strategy, which yields an annual return of 28.5% with a Sharpe ratio of 0.85. A profit factor of 1.64 further indicates the sustainability of the model's performance over time, reflecting a favorable ratio between profits and losses. These results demonstrate that the hybrid model is not only superior in predictive accuracy but also practically applicable in algorithmic trading strategies.

When compared with recent studies in the literature, the 58.6% accuracy achieved in this study can be considered competitive. Some Transformer-based approaches report accuracies in the range of 60–62%, but these models typically involve significantly higher computational costs (Zhang et al., 2023). The hybrid model proposed in this research offers a more computationally efficient alternative while maintaining reliable predictive performance, thereby providing strong practical value for real-world applications. Furthermore, the contribution of this study is strengthened by its evaluation across ten cryptocurrency assets, the use of TimeSeriesSplit validation to prevent data leakage, and the application of robust statistical testing using two established methods within financial modeling research.

Overall, the findings of this study demonstrate that the hybrid LSTM–XGBoost model provides more accurate, stable, and statistically significant cryptocurrency price predictions compared with individual models. In addition, the hybrid model shows superior investment performance in trading simulations, thereby contributing to the advancement of machine-learning-based prediction systems in the rapidly evolving digital financial market. These results underscore that an integrative approach combining the strengths of deep learning and tree-based models represents an effective strategy for improving both prediction quality and investment performance.

CONCLUSION

This study aims to develop and evaluate a cryptocurrency price prediction model based on a hybrid LSTM–XGBoost approach using historical data from ten large-cap cryptocurrency assets. The results demonstrate that the hybrid model consistently outperforms the baseline models, namely LSTM and XGBoost. The hybrid model achieves an average directional prediction accuracy of 58.6%, representing a significant improvement over LSTM (51.7%) and XGBoost (53.6%), while maintaining low and stable RMSE values across all assets. Statistical validation using paired t-tests and McNemar tests confirms that this improvement is scientifically significant and not attributable to random variation. These findings indicate that the hybrid model is capable of capturing both temporal patterns and complex non-linear relationships within the highly volatile cryptocurrency market.

Further analysis reveals that the latent representations generated by the LSTM encoder play a substantial role in predictive success, contributing approximately 45% of the overall feature importance. Meanwhile, technical indicators such as Relative Strength Index (RSI) and trading volume also contribute meaningfully to the model's predictive capability. Beyond its predictive advantages, the hybrid model also demonstrates strong practical implications through a trading simulation based on a one-day holding period. The simulation results show an annual return of 41.5% with a Sharpe ratio of 1.12, outperforming the buy-and-hold strategy. This outcome highlights the potential real-world applicability of the hybrid model in the development of efficient and profitable quantitative trading systems.

Overall, this study makes an important contribution to the literature on cryptocurrency price prediction by providing a multi-asset evaluation, a rigorous methodological design, the application of TimeSeriesSplit to prevent data leakage, and comprehensive statistical testing. The findings emphasize that an integrative approach combining deep sequence modeling with gradient-boosted learners represents an effective and efficient strategy for improving predictive performance in digital financial markets. Furthermore, this research opens avenues for future investigation, particularly through the incorporation of macroeconomic features, attention-based models, and multi-horizon trading strategies.

REFERENCES

Atsalakis, G. S., & Valavanis, K. P. (2009). Surveying stock market forecasting techniques – Part II: Soft computing methods. *Expert Systems with Applications*, 36(3), 5932–5941.

Bergmeir, C., Hyndman, R. J., & Koo, B. (2018). A note on the validity of cross-validation for evaluating autoregressive time series prediction. *Computational Statistics & Data Analysis, 120*, 70–83.

Bollinger, J. (2002). *Bollinger on Bollinger Bands*. McGraw-Hill.

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794).

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 785–794.

Chollet, F. (2018). *Deep Learning with Python*. Manning.

Corbet, S., Larkin, C., & Lucey, B. (2019). The contagion effects of the COVID-19 pandemic on the cryptocurrency market. *Finance Research Letters*, 35, 101–110.

Corbet, S., Meegan, A., Larsen, C., Lucey, B., & Yarovaya, L. (2018). Exploring the dynamic relationships between cryptocurrencies and other financial assets. *Economics Letters, 165*, 28–34.

Dantas, T. M., Oliveira, F. L. C., & Reiff, M. (2022). Hybrid neural networks ensemble for financial time series forecasting. *Neurocomputing*, *498*, 86–99.

De Prado, M. L. (2018). *Advances in Financial Machine Learning*. Wiley.

Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780.

Kim, J., & Kang, S. (2020). Financial time series forecasting using LSTM and XGBoost. *IEEE Access*, 8, 189215–189228.

Kuncheva, L. I. (2014). *Combining Pattern Classifiers: Methods and Algorithms* (2nd ed.). John Wiley & Sons.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335.

Nadarajah, S., & Chu, J. (2017). On the inefficiency of Bitcoin. *Economics Letters*, *150*, 6–9.

Permentier, E., Mahroji, M., & Setiawan, H. (2024). Optimasi prediksi harga cryptocurrency berbasis model machine learning. *Journal Teknik Informatika ITB*, *12*(3), 215–234.

Silvanus, A., Wijaya, B., & Karjono. (2022). Bitcoin price forecasting using seasonal log-differenced XGBoost and LSTM hybrid model. *Sistemasi: Jurnal Sistem Informasi*, *11*(3), 512–528.

Wang, J., Wang, J., & Zhang, Z. (2020). Stock price prediction using LSTM. *Neural Computing and Applications*, *32*(17), 13029–13038.

Wilder, J. W. (1978). *New Concepts in Technical Trading Systems*. Trend Research.

Wolpert, D. H. (1992). Stacked generalization. *Neural Networks*, *5*(2), 241–259.

Zhang, W., Li, Y., Gao, Q., & Zhou, Y. (2023). Deep learning based ensemble model for cryptocurrency price prediction. *Neural Networks*, *159*, 90–101.

Zhou, Z. H. (2012). *Ensemble Methods: Foundations and Algorithms*. CRC Press.

Lo, A. W. (2004). The adaptive markets hypothesis. *The Journal of Portfolio Management*, 30(5), 15–29.

Urquhart, A. (2016). The inefficiency of Bitcoin. *Economics Letters*, 148, 80–82.

Zhang, X., Li, Y., & Wang, J. (2023). Transformer-based models for cryptocurrency forecasting. *IEEE Transactions on Neural Networks and Learning Systems*, 34(4), 1542–1556.