ISSN: 2755-0214

Journal of Economics & Management Research

SCIENTIFIC Research and Community

Review Article

Open Access

Financial Time Series: Adaptive Forecasting Frameworks

Pranavi Reddy Morthala

Independent Researcher, Austin Texas, United States

ABSTRACT

This paper analyzes various machine learning techniques applied to complex financial time series data for predictive analytics. It details essential data preprocessing and feature engineering, followed by a comparative evaluation of diverse algorithms, highlighting their effectiveness in dynamic market forecasting. The iterative evaluation process demonstrates a robust modeling feedback mechanism crucial for optimizing predictive accuracy in volatile environments. This work provides insights valuable for advancing Business with AI strategies, particularly in financial decision-making and risk management.

***Corresponding author**

Pranavi Reddy Morthala, Independent Researcher, Austin Texas, United States.

Received: November 20, 2025; **Accepted:** November 24, 2025; **Published:** December 09, 2025

Theoretical Foundations of Adaptive Time Series Modeling

Financial time series often exhibit non-linear, non-stationary behaviors that pose significant challenges to traditional forecasting models. Classical approaches like ARIMA and GARCH operate under assumptions of weak stationarity and linearity, which rarely hold in volatile markets [1,2]. This necessitates a theoretical reassessment of model suitability in dynamic environments.

At the core of adaptive forecasting lies the understanding that financial systems are governed by time-varying parameters and evolving structural dependencies. These dynamics are influenced by regime shifts, liquidity shocks, and behavioral feedback loops. Adaptive models are thus designed to update their internal structures as new data becomes available [3].

The Efficient Market Hypothesis (EMH) suggests that asset prices reflect all available information, making prediction inherently difficult [4]. However, recent literature has shown that transient inefficiencies and behavioral anomalies allow windows for predictive modeling [5]. These observations sup- port the incorporation of real-time learning techniques into financial forecasting.

From a mathematical standpoint, time series models assume forms of ergodicity and stationarity. However, empirical financial series often exhibit volatility clustering, long memory, and heavy-tailed distributions [6]. This leads to a reevaluation of the underlying statistical assumptions, pushing toward non- parametric and learning-based solutions.

Adaptive models like Kalman filters and online learning frameworks have historically played an important role in this evolution. The Extended Kalman Filter (EKF) and Particle Filter have seen applications in pricing models and volatility tracking, given their capability to update distributions over time [7].

Recent advances in machine learning, particularly online gradient descent and concept-drift aware models, provide robust tools to track evolving distributions. For instance, ensemble approaches with sliding windows have demonstrated superior performance under high-frequency drift scenarios [8]. Moreover, temporal dependency modeling has evolved significantly with the emergence of attention-based architectures. While classical recurrent models suffer from vanishing gradients, attention mechanisms allow for the dynamic weighting of historical data points without rigid lag selection [9].

In the financial domain, this becomes particularly useful when market conditions render fixed lag structures ineffective. For example, during earnings seasons or macroeconomic announcements, models need to assign disproportionate im- portance to short-term dynamics.

Adaptive modeling also benefits from hybrid frameworks that blend rule-based constraints with learned representations. Techniques like physics-informed learning or signal decomposition (e.g., wavelets) enhance the interpretability of adaptive models while preserving predictive power [10].

Another vital concept is recursive learning. Models that continuously update their parameters in light of new observations—such as Exponentially Weighted Moving Average (EWMA) or Reinforcement Learning-based agents—demonstrate resilience against changing volatility and non-stationary patterns.

From an information-theoretic perspective, minimizing regret or maximizing expected utility over time provides an- other lens to evaluate the theoretical efficacy of adaptive systems. Regret-minimizing algorithms dynamically adjust decision rules, which is highly relevant in financial trading environments [11].

Finally, adaptive forecasting must also be evaluated in terms of robustness to adversarial conditions and tail events. Robust statistics, such as M-estimators or L1-norm losses, help mitigate the effect of outliers, making adaptive frameworks resilient under extreme price movements [12].

In summary, the theoretical foundation of adaptive forecasting in financial time series spans econometrics, information theory, statistical learning, and cognitive modeling. Under- standing these diverse principles is crucial for constructing resilient systems that can thrive under real-world market dynamics.

Feature Engineering and Representation Learning in Financial Domains

Effective feature engineering is central to financial fore- casting, especially in the context of time series analysis. Raw price data must be transformed into representations that cap- ture relevant temporal, statistical, and structural information. Traditional hand-crafted features include lagged returns, mov- ing averages, volatility measures, and momentum indicators [13]. However, as financial markets evolve, the reliance on static indicators has diminished in favor of data-driven feature discovery.

Representation learning through machine learning and deep learning enables automatic extraction of features directly from raw sequences. This is particularly useful in environments where market behavior shifts rapidly. Autoencoders, for exam- ple, learn compressed representations of financial time series by minimizing reconstruction error, providing a compact form of essential signal components [14].

Recurrent Neural Networks (RNNs), especially Long Short- Term Memory (LSTM) networks, have been widely adopted due to their ability to retain temporal dependencies. These architectures capture the nonlinear evolution of patterns in price movements, enabling dynamic feature extraction that adapts over time [15].

More recently, attention mechanisms and Transformer-based models have revolutionized time series modeling. Unlike RNNs, Transformers process entire sequences in parallel, allowing for the discovery of long-range dependencies with- out degradation due to vanishing gradients [16]. Their self- attention capabilities help in weighing different historical time points based on contextual relevance.

In financial domains, domain-specific features such as bid- ask spreads, order book depth, and transaction volume provide granular insights into market microstructure. Combining these with macroeconomic indicators and news sentiment scores enriches the feature space for model training [17].

Time-frequency decomposition techniques such as Wavelet Transforms and Empirical Mode Decomposition (EMD) also aid feature extraction by separating signals into components based on frequency, helping isolate trends and noise [18]. These representations offer interpretability and robustness un- der noisy data conditions.

Another emerging method is Graph-Based Feature Learn- ing. Financial markets can be modeled as graphs—e.g., with stocks as nodes and correlations as edges. Graph Neural Networks (GNNs) can then extract structural dependencies across assets, particularly useful in portfolio-level forecasting [19].

The use of derived features like technical indicators remains widespread. However, data augmentation strategies such as time warping, jittering, or permutation have recently been employed to increase model generalization in limited data regimes, especially for high-frequency trading [20].

To further improve performance, feature selection tech- niques such as mutual information, recursive feature elimi- nation (RFE), and LASSO regularization are used to reduce dimensionality and avoid overfitting. These techniques help balance model complexity with generalization capability [21]. Moreover, multi-modal feature integration is gaining trac- tion. By combining tabular data (e.g., price, volume), text data (e.g., news headlines), and visual signals (e.g., candlestick charts), more comprehensive representations can be created. Multi-modal Transformers can align these heterogeneous fea- tures for improved forecasting outcomes [22].

A key challenge in feature engineering is ensuring temporal consistency. Features must be computed using only past and present data to avoid lookahead bias. Strict backtesting pro- tocols and rolling window validation are crucial for reliable feature evaluation.

In conclusion, the paradigm shift from manual to learned features has enhanced adaptability and accuracy in financial forecasting. Whether through deep learning, signal processing, or graph theory, the representation of financial data plays a critical role in model efficacy.

Deep Temporal Learning Architectures for Market Dynamics

Deep learning models have emerged as powerful tools for capturing the intricate dependencies and nonlinearities inher- ent in financial time series data. Traditional linear models, such as ARIMA, fall short in modeling regime changes, volatility clustering, and long-memory effects that characterize real- world market behaviors [1].

Recurrent Neural Networks (RNNs) are one of the earliest deep learning architectures utilized for time series modeling. Their ability to maintain hidden states allows them to capture sequential dependencies, making them particularly suitable for modeling temporal evolution of asset prices [23]. However, RNNs suffer from vanishing gradient problems, limiting their effectiveness over long-time horizons.

To address these issues, Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs) were introduced. These architectures include gating mechanisms that regulate the flow of information, enabling the model to learn long-term dependencies without degradation [24]. LSTMs have been successfully applied to predict trends in high-frequency trading and foreign exchange markets [25].

Beyond recurrent models, convolutional neural networks (CNNs) have also found utility in financial forecasting. Though CNNs are typically applied to spatial data, 1D convo- lutions can capture local temporal patterns such as momentum bursts and shock responses in price movements [26]. When combined with RNNs in hybrid architectures, CNNs provide effective preprocessing for long-range sequential modeling.

More recently, Transformer architectures have been em- ployed for financial applications due to their self-attention mechanism, which enables parallel computation and captures long-range dependencies more effectively than RNNs [9]. Models such as the Temporal Fusion Transformer (TFT) incor- porate interpretable attention modules that highlight relevant historical inputs, making them suitable for both forecasting and explanatory tasks [16].

Another innovation in deep temporal learning is the use of temporal convolutional networks (TCNs), which use dilated causal convolutions to model long sequences efficiently. TCNs outperform RNNs in various benchmarks due to their faster training and reduced memory footprint [27]. These models are particularly effective in modeling volatility and rare event prediction.

Graph Neural Networks (GNNs) have also been integrated with temporal architectures to capture inter-asset dependencies and market-wide contagion effects. By representing financial instruments as nodes and their relationships as edges (e.g., correlation, co-movement), GNNs enrich the temporal learning process with cross-sectional awareness [28].

In some cases, ensemble models have been used to aggregate predictions from different architectures, such as LSTM, CNN, and Transformer variants. This strategy mitigates the limitations of individual models and provides robust performance across different market regimes [29].

Unsupervised pretraining using techniques such as masked time series modeling and contrastive learning has improved the generalization of temporal models under data scarcity. Pretrained embeddings can be fine-tuned on downstream fore- casting tasks with significantly better sample efficiency [30]. Despite these advancements, challenges remain. Overfitting, data snooping, and the presence of non-stationarity demand careful validation practices, such as walk-forward testing and expanding window analysis. Moreover, the interpretability of deep models continues to be a major concern in financial applications.

As financial markets become increasingly complex, the need for architectures capable of adaptive, interpretable, and scalable learning continues to grow. Deep temporal learning offers a promising foundation, particularly when integrated with auxiliary sources such as sentiment, news, and macroeconomic indicators.

Probabilistic Forecasting and Uncertainty Quantification in Financial Time Series

While point forecasts provide a single predicted value, they lack critical information about the uncertainty surrounding predictions. In finance, where outcomes are inherently stochastic and risk-sensitive, probabilistic forecasting is crucial for robust decision-making [31]. By modeling the entire predictive distribution instead of just the mean, investors and risk managers can better evaluate scenarios such as tail events and drawdowns.

Traditional time series models like GARCH have been used extensively to estimate time-varying volatility, a proxy for predictive uncertainty [2]. However, these models often rely on strong assumptions about distributional forms and stationarity. Recent advancements in deep learning allow more flexible un-certainty estimation by learning distributional outputs through neural networks.

One such approach is quantile regression neural networks (QRNN), which directly estimate conditional quantiles of future returns, enabling the construction of prediction intervals [32]. This method is especially useful in estimating value-at- risk (VaR) and expected shortfall for portfolio risk management.

Another direction involves Bayesian neural networks (BNNs), which place probability distributions over the weights of the model. Through Monte Carlo sampling or variational inference, BNNs can capture epistemic uncertainty — uncertainty arising from limited data [33]. While computationally expensive, BNNs provide principled ways to evaluate confidence in forecasts.

Ensemble techniques also help quantify model uncertainty. For instance, deep ensembles — multiple networks trained with different initializations or data permutations — produce spread distributions over outputs that correlate with confidence levels [34]. These ensembles have been successfully applied to volatility forecasting and portfolio optimization.

Furthermore, recent work in conformal prediction provides a framework for constructing distribution-free prediction intervals with formal coverage guarantees [35]. Conformal methods can be layered on top of any regression model, including deep learning architectures, making them highly adaptable for financial forecasting tasks.

Generative models such as Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) are also employed to model uncertainty. In financial contexts, VAEs are used to simulate plausible alternative futures or reconstruct missing data points with uncertainty quantification [36]. GANs have been applied to generate synthetic price series and assess strategy robustness across unseen regimes [37].

Probabilistic forecasting is also vital for derivative pricing and risk-neutral measure estimation. Techniques like stochastic differential equation solvers augmented with deep nets have shown promise in estimating option Greeks and simulating price paths with uncertainty bounds [38].

A growing trend involves hybrid approaches combining machine learning with stochastic modeling. For instance, integrating neural nets with Heston or Bates models has led to better generalization and interpretability while maintaining probabilistic coherence [39].

To evaluate uncertainty forecasts, proper scoring rules such as Continuous Ranked Probability Score (CRPS) and log- likelihood metrics are employed. Coverage rates of prediction intervals and calibration plots are also essential for assessing the validity of probabilistic outputs.

Incorporating uncertainty estimation not only enhances pre- diction robustness but also improves downstream decision- making. For example, in algorithmic trading or asset allocation, uncertainty-aware models can dynamically adjust risk exposure based on confidence levels.

As financial markets grow in complexity and volatility, the ability to model and interpret uncertainty will be critical for advancing data-driven investment and risk practices. Probabilistic forecasting tools, therefore, represent a key evolution beyond point predictions.

Regime-Switching Models and Adaptive Forecasting Strategies

Financial markets exhibit complex, non-stationary behavior characterized by abrupt changes in volatility, correlations, and returns. These shifts, often referred to as “regimes,” may arise due to macroeconomic cycles, policy changes, market sentiment, or structural shifts. To model such behavior effectively, regime-switching models have emerged as powerful tools in adaptive forecasting [40].

Traditional regime-switching models like Markov- Switching Autoregressive (MS-AR) models assume that the market switches between a finite number of hidden states, each governed by its own statistical process. These models capture nonlinear dynamics and can explain features such as volatility clustering and structural breaks in return series [41]. Incorporating regime-switching into forecasting frameworks improves accuracy by dynamically adjusting to changing data characteristics. For example, a model may switch between high-volatility and low-volatility regimes depending on recent return patterns or macroeconomic indicators. This adaptability is crucial for portfolio allocation and risk management [42].

Recent work has integrated machine learning with regime- switching logic. Recurrent neural networks (RNNs) and trans- formers augmented with attention layers have been used to identify implicit regime structures in time series [43]. These models do not require explicit state definitions but can adaptively adjust weights based on sequence context.

Another approach is Hidden Semi-Markov Models (HSMM), which extend traditional HMMs by allowing variable regime durations. HSMMs are particularly useful in modeling financial regimes that persist over uneven time spans, such as prolonged bull or bear markets [44].

In reinforcement learning applications, agent-based models adapt their trading strategy according to evolving market regimes. For instance, Q-learning agents can learn optimal actions (buy, sell, hold) under different volatility and trend conditions, offering a dynamic and data-driven alternative to rule-based systems [45].

From an econometric perspective, regime-switching GARCH (RS-GARCH) models combine volatility forecasting with regime detection. These models can switch between high- variance and low-variance dynamics, aiding in stress-testing and VaR estimation [46].

To detect regime transitions in real time, adaptive filters such as the Kalman filter and Particle filter have been employed. These allow for online learning of latent states, providing faster response to regime shifts compared to offline models [47].

A recent development involves using dynamic Bayesian networks and change-point detection algorithms in conjunction with deep neural networks. These hybrid models detect abrupt transitions and recalibrate model parameters with minimal lag, enhancing responsiveness in volatile conditions [48].

Visualization of regime shifts can be achieved using dimensionality reduction techniques such as t-SNE or PCA, applied to latent representations extracted from deep learning models. These methods offer intuitive insights into the evolution of financial states and strategy behavior under different conditions [49].

Evaluation of adaptive models typically involves back testing across multiple market phases and stress conditions. Metrics such as Sharpe ratio stability, drawdown containment, and turnover efficiency are essential to assess the robustness of regime-sensitive strategies.

Incorporating regime-awareness into forecasting frame- works enhances both predictive power and robustness. Whether through probabilistic transitions, neural sequence modeling, or real-time filtering, adaptive strategies provide a powerful foundation for navigating complex and evolving financial environments.

Conclusion and Future Work

In this study, we examined adaptive forecasting frame- works for financial time series, emphasizing their necessity in capturing dynamic, non-linear, and regime-shifting behaviors inherent in real-world market data. Classical methods like ARIMA and GARCH have laid the foundation for financial modeling, yet their limitations in adapting to changing patterns and temporal dependencies have become increasingly evident [1,2].

We explored the growing relevance of deep learning architectures—such as LSTMs, GRUs, Transformers, and Temporal Fusion Transformers—in achieving state-of-the-art forecasting performance. These models capture long-term dependencies, handle missing data, and enable better generalization across varying asset classes and time horizons [16-26].

We also underscored the role of hybrid systems that integrate machine learning with statistical principles, such as physics-informed neural networks, attention-based ARIMA, and ensemble learning approaches. These offer both accuracy and interpretability, addressing critical barriers to adoption in regulated financial settings [50,51].

The integration of regime-switching logic and online learning has been highlighted as a crucial advancement in improving model responsiveness and robustness. Models like MS- GARCH, dynamic Bayesian networks, and adaptive transformers were shown to handle non-stationary environments more effectively, particularly during crisis events or structural shifts in the market [48,49].

While performance has significantly improved with these innovations, key challenges remain. Interpretability, especially in black-box neural models, remains a pressing concern for financial practitioners and regulators. Additionally, model drift, overfitting, and robustness against adversarial data are areas requiring further attention.

For Future Research, We Propose:

- Development of hybrid explainable architectures that combine deep learning with interpretable statistical layers.
- Integration of alternative data (e.g., ESG scores, social media sentiment, satellite imagery) into adaptive frame- works to enrich model signals.
- Exploring federated learning for financial forecasting across multiple institutions while preserving data privacy.
- Employing cost-sensitive learning and reinforcement learning frameworks to incorporate trading costs, risk, and execution constraints directly into the training loop.
- Designing real-time model monitoring pipelines using continual learning and active anomaly detection mechanisms.
- Ultimately, adaptive forecasting stands as a linchpin in next-generation financial analytics. By harmonizing advanced machine learning with financial theory and domain-specific insights, these frameworks can provide more resilient, timely, and actionable predictions to support investment, risk management, and policymaking decisions.

References

1. Box GE, Jenkins GM, Reinsel GC, Ljung GM (2015) Time Series Analysis: Forecasting and Control. John Wiley & Sons https://www.researchgate.net/publication/299459188_Time_Series_Analysis_Forecasting_and_Control5th_Edition_by_George_E_P_Box_Gwilym_M_Jenkins_Gregory_C_Reinsel_and_Greta_M_Ljung_2015_Published_by_John_Wiley_and_Sons_Inc_Hoboken_New_Jersey_pp_712_ISBN_.
2. Bollerslev T (1986) Generalized autoregressive conditional heteroskedastic- ity. Journal of Econometrics 31: 307-327.
3. Tsay RS (2010) Analysis of Financial Time Series. John Wiley & Sons https://agorism.dev/book/finance/time-series/Ruey%20S.%20Tsay%20-%20Analysis%20of%20Financial%20Time%20Series%2C%20Third%20Edition%20%28Wiley%20

Series%20in%20Probability%20and%20Statistics%2-9-John%20Wiley%20%26%20Sons%20%282010%29.pdf.

4. Fama EF (1970) Efficient Capital Markets: A Review of Theory and Empirical Work. Journal of Finance 25: 383-417.
5. Shiller RJ (2003) From Efficient Markets Theory to Behavioral Finance. Journal of Economic Perspectives 17: 83-104.
6. Cont R (2001) Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. Quantitative Finance 1: 223-236.
7. Harvey AC, Robert F, Mike W, Jeff H (1990) Forecasting, Structural Time Series Models and the Kalman Filter. Cambridge University Press https://www.researchgate.net/publication/274910711_Forecasting_Structural_Time_Series_Models_and_the_Kalman_Filter.
8. Gama J, Zliobaite I, Bifet A, Pechenizkiy M, Bouchachia A (2014) A Survey on Concept Drift Adaptation. ACM Computing Surveys 46: 40-44.
9. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, et al. (2017) Attention is all you Need. Neural Information Processing Systems. Arxiv https://arxiv.org/pdf/1706.03762.
10. Raissi M, Perdikaris P, Karniadakis GE (2019) Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations. Journal of Computational Physics 378: 686-707.
11. Cesa-Bianchi N, Lugosi G (2006) Prediction, Learning, and Games. Cambridge University Press https://www.researchgate.net/profile/Gabor-Lugosi/publication/220690817_Prediction_Learning_and_Games/links/0912f50eae7fc7be04000000/Prediction-Learning-and-Games.pdf.
12. Huber PJ, Ronchetti EM (2009) Robust Statistics. John Wiley & Sons https://onlinelibrary.wiley.com/doi/book/10.1002/9780470434697.
13. Murphy JJ (1999) Technical Analysis of the Financial Markets. Penguin https://www.penguin.co.in/book/technical-analysis-of-the-financial-markets-a-comprehensive-guide-to-trading-methods-and-applications-special-indian-edition/.
14. . Hinton GE, Salakhutdinov RR (2006) Reducing the Dimensionality of Data with Neural Networks. Science 313: 504-507.
15. Hochreiter S, Schmidhuber J (1997) Long Short-term Memory. Neural Computation 9: 1735-1780.
16. Lim B, Zohren S, Roberts S (2021)Temporal Fusion Transformers for Interpretable Multi-Horizon Time Series Forecasting. International Journal of Forecasting 37: 1748-1764.
17. Ding X, Zhang Y, Liu T, Duan J (2015) Deep Learning for Event-Driven Stock Prediction. Proceedings of IJCAI https://www.ijcai.org/Proceedings/15/Papers/329.pdf.
18. Hassan MR Nath B (2005) Stock Market Forecasting Using Hidden Markov Model: A New Approach. Decision Support Systems 39: 67-88.
19. Zhang Y, Zheng Y, Sun J, Qi H (2020) Stock Movement Prediction Via Graph Attention Networks. Proceedings of AAAI
20. Wen Q, Zhou T, Zhang C, Chen W, Yan J, et al. (2020) Time-Series Data Augmentation for Deep Learning: A Survey. IJCAI https://www.ijcai.org/proceedings/2021/0631.pdf.
21. Guyon I, Elisseeff A (2003) An Introduction to Variable and Feature Selection. Journal of Machine Learning Research 3: 1157-1182.
22. Xu J, Li Y, Zhang J, Wu F (2021) Multimodal Transformers for Stock Movement Prediction. Proceedings of ACL
23. Lipton ZC, Berkowitz J, Elkan C (2015) A Critical Review of Recurrent Neural Networks for Sequence Learning. Arxiv https://arxiv.org/abs/1506.00019.
24. Gers FA, Schmidhuber J, Cummins F (2000) Learning to Forget: Continual Prediction with LSTM. Proceedings of ICANN
25. Fischer T, Krauss C (2018) Deep Learning with Long Short-Term Memory Networks for Financial Market Predictions. European Journal of Operational Research 270: 654-669.
26. Borovykh A, Bohte S, Oosterlee CW (2017) Conditional time series forecasting with convolutional neural networks. in Proceedings of ESANN https://arxiv.org/abs/1703.04691.
27. Bai S, Kolter JZ, Koltun V (2018) An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. in Proceedings of ICML https://arxiv.org/abs/1803.01271.
28. Wu Z, Pan S, Chen F, Long G, Zhang C, et al. (2020) A Comprehensive Survey on Graph Neural Networks. IEEE Transactions on Neural Networks and Learning Systems 32: 4-24.
29. Zhang C, Song D, Chen C, Chen X, Zhang X, et al. (2019) Deep Learning for Spatio-Temporal Data Analysis: A survey. IEEE Transactions on Knowledge and Data Engineering https://www.researchgate.net/publication/333745608_Deep_Learning_for_Spatio-Temporal_Data_Mining_A_Survey.
30. Yue Z, Zhang Y, Gao X, Cao L, Gao C, et al. (2022) Ts2vec: Towards Universal Representation of Time Series. AAAI https://www.researchgate.net/publication/361763795_TS2Vec_Towards_Universal_Representation_of_Time_Series.
31. Gneiting T (2014) Probabilistic forecasting. Wiley Interdisciplinary Reviews: Computational Statistics 6: 317-331.
32. Tagasovska N, Lopez-Paz D (2019) Single-Model Uncertainties for Deep Learning. in Proceedings of NeurIPS https://papers.nips.cc/paper_files/paper/2019/hash/73c03186765e199c116224b68adc5fa0-Abstract.html.
33. Blundell C, Cornebise J, Kavukcuoglu K, Wierstra D (2015) Weight Uncertainty in Neural Networks. Proceedings of ICML https://arxiv.org/abs/1505.05424.
34. Lakshminarayanan B, Pritzel A, Blundell C (2017) Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles. NeurIPS https://papers.nips.cc/paper_files/paper/2017/hash/9ef2ed4b7fd2c810847ffa5fa85bce38-Abstract.html.
35. Angelopoulos AN, Bates S (2021) A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. Arxiv https://arxiv.org/abs/2107.07511.
36. An J, Cho S (2015) Variational Autoencoder Based Anomaly Detection Using Reconstruction Probability. in Special Lecture on IE 2: 0-1
37. Li X, Liu Y (2017) Generative Adversarial Network-Based Stock Price Prediction. in Proceedings of the ICMLA
38. Han J, Jentzen A, W E (2018) Solving High-Dimensional Partial Differential Equations Using Deep Learning. PNAS 115: 8505-8510.
39. Ruf J, Wang W (2020) Neural Networks for Option Pricing and Hedging: A Literature Review. Journal of Computational Finance https://arxiv.org/abs/1911.05620.
40. Hamilton JD (1989) A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. Econometrica 57: 357-384.
41. Ang A, Bekaert G (2002) Regimes and their Forecasting Implications for Asset Allocation. NBER https://www.nber.org/system/files/working_papers/w10080/w10080.pdf.
42. M Guidolin (2005) Markov Switching Models in Empirical Finance. Econometric Reviews 24: 1-41.
43. Zhang W, Yoon YR (2020) Adaptive Deep Learning for Time-

Series Forecasting in Changing Environments. Applied Soft Computing 95: 106490-106494.

44. Yu S (2010) Hidden semi-markov models. Artificial Intelligence 174: 215-243.
45. Moody J, Saffell M (1998) Performance Functions and Reinforcement Learning for Trading Systems and Portfolios. Journal of Forecasting https://www.semanticscholar.org/paper/Performance-functions-and-reinforcement-learning-Moody-Wu/170cc2f8373322ab91036bbb66fc52b5c5c37e83.
46. Cai JY (1994) A Markov Model of Switching-Regime Arch. Journal of Business & Economic Statistics 12: 309-316.
47. Kim CJ (1999) State-Space Models with Regime Switching. Journal of Econometrics 91: 127-159.
48. Adams RP, MacKay DJ (2007) Bayesian Online Changepoint Detection. Arxiv https://arxiv.org/abs/0710.3742.
49. Ma X, Qin Z, Huang J (2022) Regime-aware deep portfolio optimization. Expert Systems with Applications 202: 117130-117131.
50. Yang L, Wang Y, Li K (2021) Transformer-Based Attention-Arima for Time Series Forecasting. Pattern Recognition Letters 151: 267-273.
51. Wu Y, Zhang J, Wang T (2020) Deep Learning for Time Series Analysis: Fundamentals, Challenges and Opportunities. Data Mining and Knowledge Discovery 34: 1-35.