Stock Return Prediction Based on Financial Text Sentiment Analysis and Dynamic Lead-Lag Relationships in FF5 Residual Networks

Juntao Zhou*

School of Business, Macau University of Science and Technology, Macau, China

*2250030500@student.must.edu.mo

Abstract. Stock return prediction remains challenging because public information, investor attention and cross-stock signals are incorporated into prices at different speeds. This study develops a two-layer empirical framework for a fixed universe of 150 U.S.-listed stocks. The prediction layer first estimates positive-forward-return probabilities with a PatchTST-based price model and then combines them with lagged sentiment features from investor comments, financial news and formal disclosures to measure the incremental contribution of text. Separately, the relationship layer removes common factor-driven co-movements with FF5 residuals and constructs dynamic Lead-Lag networks to identify Missing Link opportunities in delayed cross-stock reactions. The contribution of text varies across sources and horizons: investor-comment information is more useful at short and short-to-medium horizons, whereas comments combined with formal disclosures are more stable at medium and long horizons. The strongest representative Missing Link results occur in 60-70-day graph windows with 5-7-day follower periods, while the broader short-follower-period grid shows that positive evidence is concentrated in 40-70-day dynamic graph windows.

Keywords: stock return prediction, financial text sentiment, FF5 residuals, Lead-Lag network, Missing Link

1 Introduction

An accurate forecast of stock returns is still an important problem in empirical finance. Machine learning demonstrates its capability to detect nonlinear return trends in the out-of-sample prediction taking into account transaction costs [1]. Deep learning models can analyze complicated risk structures [2], and research on asset pricing emphasizes the importance of out-of-sample validation [3]. In addition to the price-related information, investor discussions, financial news and corporate announcements may also involve useful information such as attention, sentiment and events which are not explicitly shown in the stock prices. Sentiment obtained from texts has been observed to enhance the prediction of stock returns [4]. Since raw returns might fluctuate at both the overall market and style levels, the Fama-French five-factor model is adopted to reduce the

© The Author(s) 2026
V. Erokhin et al. (eds.), *Proceedings of the 6th International Conference on Internet Finance and Digital Economy (ICIFDE 2026)*, Advances in Economics, Business and Management Research 412,
https://doi.org/10.2991/978-94-6239-830-6_15

effect of these common factors [5]. Both price data and text data can be applied for predicting stock movements [6]. Graph-based prediction makes use of the connections between different stocks [7], and event-driven relational forecasting indicates that stock movements may propagate through some specific relationships [8]. Therefore, the purpose of this study is twofold. The main objective is to examine whether financial texts influence the directional prediction in the machine learning techniques. The secondary objective is to find the delayed inter-stock responses using FF5 residual-based dynamic Lead-Lag networks under a strict no-lookahead constraint. FF5 residuals are utilized to decrease the co-movements of general market and style factors before constructing the Lead-Lag network.

The contributions are fourfold. First, the two-layer framework separates return prediction from delayed cross-stock analysis. Second, multi-source text is evaluated as an increment to PatchTST-based prediction on identical samples. Third, FF5 residuals support dynamic Lead-Lag construction. Finally, Missing Link provides event-driven decision support under strict no-lookahead constraints.

The empirical workflow contains three linked but distinct steps. First, daily market sequences are used by the PatchTST-based model to generate price-only positive-return probabilities. Second, timestamped sentiment features from investor discussions, financial news and formal disclosures are aggregated to the stock-day level and fused with those probabilities. Third, daily stock returns and FF5 factor returns are used separately to construct residuals for Lead-Lag and Missing Link delayed-reaction analysis; machine-learning probabilities do not enter the relationship layer.

2 Data and Methodology

2.1 Sample and Data

This study examines a fixed pool of 150 actively traded U.S.-listed stocks, mainly from NYSE and NASDAQ, over 2013-2025. Candidate stocks are screened for continuous trading, sufficient price history, public-information and text coverage, and network comparability; the final pool spans different market-capitalization and share-price ranges and major sectors. Keeping the same stocks preserves comparable graph nodes across the prediction and relationship analyses and makes multi-source text collection and timestamp matching feasible. This design may introduce survivorship and liquidity-selection bias, so the findings are interpreted only within the fixed universe rather than generalized to the entire U.S. equity market. The dataset contains daily opening, high, low and closing prices, trading volume, related technical indicators, anonymized investor comments, anonymized financial news, formal SEC/company disclosures and Fama-French five-factor data. The platforms are anonymized because the analysis compares their informational functions rather than their identities: the text sample comprises three investor-comment platforms, three financial-news platforms and official disclosures. Seven horizons are examined: 5, 10, 20, 60, 120, 180 and 252 trading days. For stock i at date t and horizon h, the forward return is defined as:

$$R_{i,t,t+h} = \frac{P_{i,t+h} - P_{i,t}}{P_{i,t}} \quad (1)$$

where R denotes the forward return of stock i from date t to t+h, and P denotes the closing price used to calculate the forward return.

The chronological split is fixed before final evaluation: 2016-01-04 to 2021-12-31 is used for machine-learning fitting, 2022-01-03 to 2023-12-29 for validation and parameter selection, and 2024-01-02 to 2025-12-31 for final price-model evaluation. Text enhancement follows the price-only prediction. For 5d, 10d and 20d, labeled 2024 observations initialize the fusion pool, and the model is refitted on each 2025 prediction date using lagged text; on date t, the latest eligible training observation is t-h for horizon h, so its forward label is already complete. For 60d, 120d, 180d and 252d, stock-level 2024 text snapshots and 2024 probability-scaling statistics are fixed before combination with 2025 price probabilities. The last eligible 2025 prediction dates are 2025-12-31 for 5d/10d/20d, 2025-12-18 for 60d, 2025-09-24 for 120d, 2025-06-30 for 180d and 2025-03-17 for 252d. Prices observed through 2026-03-18 are used only ex post to complete these labels and never enter model inputs, fitting, graph construction or signal generation.

2.2 Machine Learning and Text Enhancement

The baseline prediction layer uses a PatchTST-based time-series classifier to calculate the probability of a positive future return for each stock-date-horizon observation. The task is a binary directional classification rather than exact price prediction. The model uses patch embedding, two encoder layers, four attention heads, a hidden dimension of 64, a feed-forward dimension of 128, dropout of 0.1, AdamW, weight decay of 0.0001, gradient clipping of 1.0, early-stopping patience of 2 and random seed 42. Prediction models are selected by validation AUC rather than test performance; resampling diagnostics use separately fixed seeds. StandardScaler is fitted on the training set, and missing market variables are filled with training-set medians. This design follows the PatchTST concept of patch-based long-term time-series modeling [9]. Stock market index forecasting over different periods can also be performed through horizon-specific modeling in non-stationary financial series [10]. Further research on sequence models indicates that temporal dependencies and incremental updates should be specifically considered in stock return prediction [11]. The raw signs are 1 if the future return is positive and 0 otherwise. Clean labels eliminate the observations with very small absolute future returns by using specific quantile cutoffs based on the horizon length prior to the final test. Soft clean, utilized for the 180 days ahead forecast, includes all observations but assigns less weight to those with small returns. The 252 days ahead specification employs raw sign labels and is preliminarily tested due to the limited availability of complete prediction dates. Hybrid sentiment prediction methods propose combining investor sentiment with deep learning for stock return prediction [12]. ESG sentiment studies also show that textual sentiment can be integrated with technical indicators [13]. PatchTST settings are fixed before the final test and vary by horizon: 30-180-day lookbacks, patch/stride pairs from 4/2 to 20/10, batch sizes from 96 to 384, learning rates of 0.001 or 0.0008 and 3-4 maximum

epochs. Text sources comprise three anonymous investor-comment platforms, three anonymous financial-news platforms and SEC/company disclosures. For prediction date t, only records available by the previous trading day are used. Each source is processed separately by DeBERTa-v3-finance, ModernFinBERT and StockTwits-RoBERTa. Document outputs are aggregated by stock and date into mean sentiment, sentiment divergence, seven-day decay, document count and sentiment intensity; zero-valued text features and document count identify no-text days without future imputation. These lagged features are combined with the price model's positive-return probability. Candidate sources, sentiment models, lookback windows and source combinations are compared under each horizon protocol. Short-horizon fusion is refitted daily on 2025 prediction dates from a 2024 initial pool. For 60d, 120d, 180d and 252d, 2024 stock-level mean-sentiment snapshots are held fixed and combined with 2025 probabilities standardized using 2024 statistics. Table 1 reports the highest observed AUC within these candidate comparisons; the text-specification comparison is therefore descriptive rather than a second untouched holdout. No future text or unfinished label is used to generate an individual prediction.

2.3 FF5 Residual Network and Missing Link Definition

The relationship layer starts from raw daily stock returns and daily Fama-French five-factor returns obtained from the Kenneth French Data Library. The datasets are aligned by trading date before residual construction. To avoid the lookahead bias of full-sample factor regression, factor exposures are not estimated from the full 2013-2025 sample or from 2024-2025 test outcomes. Instead, each stock's daily excess return is regressed on MKT, SMB, HML, RMW and CMA only during the pre-test window from 2021-04-05 to 2023-12-29. The estimated intercept and factor loadings are then locked before the 2024-2025 test period. For each stock, the daily excess return is decomposed as:

During 2024-01-02 to 2025-12-31, the locked coefficients and same-day factor returns are used to calculate the factor-implied component, and the FF5 residual is the realized excess return minus that component. It therefore represents stock-specific abnormal movement after common market and style-factor effects are removed. These residuals, rather than raw returns or machine-learning probabilities, form the Lead-Lag network and Missing Link signals; no 2024-2025 stock returns are used to estimate the main FF5 betas. As a separate no-lookahead robustness check, expanding pre-date betas are re-estimated using only data available before each test date. Under this alternative residual definition, the 60-day graph-window result remains significantly positive, while the 70-day and 120-day results remain positive but are only marginally significant at the 10% level. The expanding-beta test therefore supports the medium-window pattern without replacing the fixed pre-test-beta main specification.

$$r_{i,t} = \alpha_i + \beta_{i,MKT}MKT_t + \beta_{i,SMB}SMB_t + \beta_{i,HML}HML_t + \beta_{i,RMW}RMW_t + \beta_{i,CMA}CMA_t + \varepsilon_{i,t} \quad (2)$$

where MKT, SMB, HML, RMW and CMA are the Fama-French factor returns, and epsilon is the stock-specific residual used for network construction. For each potential leader A and follower B, the rolling-window Lead-Lag correlation is:

$$\rho^{(l)}_{A\to B} = \mathrm{corr}(\varepsilon_{A,t}, \varepsilon_{B,t+l}), \quad l = 1,2,3,4,5 \tag{3}$$

where l denotes the candidate lag between leader A and follower B, and the correlation is calculated within the historical rolling graph window. The lag with the largest absolute correlation is selected as the candidate edge. The trading direction is determined by the sign of the best residual correlation and the sign of the leader residual shock:

$$\mathrm{Direction} = \mathrm{sign}(\mathrm{best_corr}) \times \mathrm{sign}(\mathrm{leader_z}) \tag{4}$$

where the direction combines the sign of the selected residual correlation with the sign of the leader's residual shock. Missing Link is a delayed-reaction test rather than proof of structural causality. At each graph-update date, the preceding historical window is used to estimate lagged FF5-residual relationships; the strongest 1,000 candidate edges are retained, the graph is refreshed every 20 trading days and is applied only to the following 20 trading dates. On each signal date, the strongest 50 available edges are checked under a no-sign-flip filter. A signal formed after the close requires an absolute 60-day standardized leader-residual z-score of at least 2.5, an absolute follower z-score no greater than 0.25, and a direction-aligned lagged formal-disclosure score above zero. That score combines sentiment acceleration, intensity and document count with weights 0.55, 0.30 and 0.15. Follower returns are measured from the next trading day over 1d, 2d, 3d, 5d or 7d. Graph construction, text confirmation and signal triggering never use future residuals, returns, text or edges.

Here, z denotes the rolling standardized FF5 residual. Missing Link is a sparse event-window filter for delayed follower responses rather than a classifier or causal model. It is therefore evaluated with cost-adjusted returns, active-day win rates, mark-to-market performance, benchmark and factor diagnostics, bootstrap and placebo tests, and no-lookahead audits instead of AUC.

3 Empirical Results

3.1 Text-enhanced Prediction Results

Table 1 reports 2025 ROC-AUC ablation results for the baseline machine-learning and text-enhanced models on exactly the same stock-date observations within each horizon. Uplift therefore measures the text increment rather than a sample-composition change. Text improves ROC-AUC at every horizon; the largest uplift occurs at 120 days, while the 252-day result remains exploratory.

The text effect varies by source and horizon. Candidate sources, sentiment models, lookback windows and source combinations are compared under the applicable protocol. The highest-observed 5d, 10d and 20d specifications are comment-based and use ModernFinBERT with 60d, 20d and 60d text windows, respectively. For 60d, 120d and 180d, the highest-observed specifications combine comments with formal disclosures, use StockTwits-RoBERTa and fix 2024 mean-sentiment snapshots before 2025 evaluation. Rank IC, Precision@Top10% and Top-Bottom Spread lead to conclusions consistent with ROC-AUC, indicating improved ranking as well as directional classi-

fication. The 252d result remains exploratory because complete one-year labels are limited. Investor-attention studies suggest that investor-generated discussions can reflect short-term market attention [14]. Multi-source sentiment studies show that stock-return effects can vary across information sources [15].

Table 1. Ablation results evaluated by ROC-AUC under different information inputs

Horizon	Protocol	N	Baseline ML AUC	Text-enhanced ML AUC	Uplift	Evidence Role
5d	clean + daily refit	15,019	0.5235	0.5381	+0.0146	Supporting Evidence
10d	clean + daily refit	18,724	0.5274	0.5359	+0.0084	Supporting Evidence
20d	clean + daily refit	18,812	0.5339	0.5741	+0.0402	Supporting Evidence
60d	clean + 2024 snapshot	17,334	0.5578	0.5898	+0.0321	Supporting Evidence
120d	clean + 2024 snapshot	17,674	0.5715	0.6260	+0.0545	Supporting Evidence
180d	soft clean + 2024 snapshot	18,300	0.5144	0.5552	+0.0408	Supporting Evidence
252d	raw sign + 2024 snapshot	7,500	0.5606	0.6003	+0.0397	Exploratory

3.2 Missing Link Delayed Reaction Results

Table 2. Missing Link results under walk-forward rolling graph construction

Graph window	Follower period	Trades	Active days	Net active-day mean	Win rate	p-value	Evidence Role
40d	7d	48	42	1.218%	66.67%	0.0236	Short-medium support
50d	3d	52	42	0.770%	57.14%	0.0455	Short-medium support
60d	5d	63	50	1.360%	68.00%	0.0048	Main evidence
60d	7d	63	50	1.426%	72.00%	0.0087	Main evidence
70d	5d	61	50	1.306%	58.00%	0.0278	Robust main evidence
70d	7d	61	50	1.409%	64.00%	0.0223	Robust main evidence
120d	5d	38	31	2.084%	67.74%	0.0202	Low-frequency support
180d	1d	39	35	0.129%	57.14%	0.3439	Boundary check
252d	1d	30	25	0.297%	52.00%	0.2185	Boundary check

Table 2 reports representative dynamic walk-forward specifications from the complete 65-specification short-follower grid, including its strongest, supporting and boundary windows.

In Table 2, Trades is the number of activated Missing Link signals, Active days counts dates with at least one signal, Net active-day mean is the 25-bps cost-adjusted mean return on active days, and Win rate is the share of positive active days. The p-value is the one-sided bootstrap probability of a nonpositive active-day mean; smaller values indicate stronger evidence of a positive mean but do not establish a permanent or causal relationship. Graph forecasting research supports the use of inter-stock network information in stock return prediction [16]. Dynamic routing graph attention supports adaptation to changing stock relationships [17]. Dynamic hypergraph studies support higher-order relationships [18]. Dynamic graph-learning research further supports time-varying relationships [19].

To avoid showing only the best parameter pair, the diagnostic grid combines 13 graph windows (20, 30, 40, 50, 60, 70, 80, 90, 100, 110, 120, 180 and 252 days) with 1d, 2d, 3d, 5d and 7d follower periods, giving 65 dynamic walk-forward specifications. Of these, 50 have positive active-day means, 11 have unadjusted p-values below 5% and 21 below 10%. Evidence is concentrated in the 40-70-day family: all 20 specifications are positive, with 9 below 5% and 11 below 10%. The 60-70-day windows are the strongest representatives in Table 2; 120 days provides low-frequency support, whereas 20-30 and 180-252 days do not support the main pattern. Missing Link is therefore interpreted as a short-to-medium dynamic Lead-Lag pattern rather than an isolated best-pair result.

3.3 Robustness, Portfolio Validation, and No-lookahead Checks

Table 2 shows the active-day returns after a 25 basis point transaction cost. The reverse-direction placebos yield negative results and the median values of the random directions are near zero, whereas the no-lookahead audit indicates that the graph estimation, formal-text confirmation and signal activation rule out the use of future information. Since the strategy possesses sparse and event-driven features, the trading intensity is assessed by the number of trades and active days instead of the daily turnover rate. The FF5 residual calculation and FF5-alpha diagnostics are employed to examine the exposure to common factors, and the ETF-benchmark diagnostics are compared with the market. Considering that various models have been analyzed, these row-level p-values should be considered in conjunction with the block bootstrap, placebo and multiple testing diagnostics rather than being treated as single confirming tests. In the fixed-year assessment, the graph window, follower period and related signal rules are determined based on the 2024 data and are then kept unchanged for use in the 2025 holdout sample. The holdout evidence for the year 2025 is weaker; only 33 out of 65 specifications have positive active-day means, none achieve the 5% significance level, one has an unadjusted p-value less than 10%, and two are close to the 10% level. This result is consistent with the time-varying Lead-Lag relationships and may explain the stronger performance of the dynamic walk-forward updating compared with a fixed annual specification.

4 Discussion

Machine-learning prediction probabilities are informative but should not be interpreted as stand-alone trading signals, especially over short horizons. Prior research shows that predictive variables may contain return information, risk exposure and factor structures at the same time [20]. Latent-factor methods support extracting return-predictive signals from many original predictors [21]. Because factor performance can vary over time and depend on factor definition, the short-, medium- and long-horizon results are interpreted separately [22].

Investor-comment platforms mainly attract short-term interest, formal disclosures support confirmation over a medium and long period, and news texts explain the background of the events. Some recent studies indicate that large language models might enhance the understanding of financial texts, though this research does not apply LLMs for prediction directly [23]. Recent evidence evaluates LLM-based stock-return signals using news headlines beyond the model's knowledge cutoff [24]. The Missing Link layer is based on events rather than a general classifier. Although it provides positive but less strong evidence for the year 2025, it also suggests that the detected relationships change with time; therefore, this layer should be used as a decision-support tool instead of being considered a permanent structural or causal model.

5 Conclusion

This study proposes a two-layer framework for stock-return direction prediction and delayed cross-stock response detection. The first layer combines PatchTST price prediction with multi-source financial text and improves ROC-AUC at every horizon, with the largest uplift at 120 days; the 252-day result is exploratory. Comments contribute more at short and short-to-medium horizons, whereas comments combined with formal disclosures are more stable at medium and long horizons. The second layer builds dynamic Lead-Lag networks from FF5 residuals and identifies Missing Link opportunities. Representative results occur in 60-70-day graph windows with 5-7-day follower periods, while the complete grid concentrates broader positive evidence in 40-70-day windows. Cost, placebo, bootstrap, multiplicity and no-lookahead checks support a cautious decision-support interpretation; weaker 2025-only significance indicates time variation rather than permanent causality. Future work may examine larger or periodically reconstituted stock universes, adaptive graph updating and foundation-model-based text representations.

References

1. Leippold M, Wang Q, Zhou W. Machine learning in the Chinese stock market[J]. Journal of Financial Economics, 2022, 145(2): 64-82. DOI: 10.1016/j.jfineco.2021.08.017.
2. Chen L, Pelger M, Zhu J. Deep learning in asset pricing[J]. Management Science, 2024, 70(2): 714-750. DOI: 10.1287/mnsc.2023.4695.

3. Gu S, Kelly B, Xiu D. Empirical asset pricing via machine learning[J]. The Review of Financial Studies, 2020, 33(5): 2223-2273. DOI: 10.1093/rfs/hhaa009.
4. Li Y, Bu H, Li J, Wu J. The role of text-extracted investor sentiment in Chinese stock price prediction with the enhancement of deep learning[J]. International Journal of Forecasting, 2020, 36(4): 1541-1562. DOI: 10.1016/j.ijforecast.2020.05.001.
5. Fama E F, French K R. A five-factor asset pricing model[J]. Journal of Financial Economics, 2015, 116(1): 1-22. DOI: 10.1016/j.jfineco.2014.10.010.
6. Xu Y, Cohen S B. Stock movement prediction from tweets and historical prices[C]//Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics. Melbourne: Association for Computational Linguistics, 2018: 1970-1979. DOI: 10.18653/v1/P18-1183.
7. Jafari A, Haratizadeh S. GCNET: graph-based prediction of stock price movement using graph convolutional network[J]. Engineering Applications of Artificial Intelligence, 2022, 116: 105452. DOI: 10.1016/j.engappai.2022.105452.
8. Xu W, Liu W, Xu C, Bian J, Yin J, Liu T-Y. REST: relational event-driven stock trend forecasting[C]//Proceedings of the Web Conference 2021. New York: Association for Computing Machinery, 2021: 1-10. DOI: 10.1145/3442381.3450032.
9. Nie Y, Nguyen N H, Sinthong P, Kalagnanam J. A time series is worth 64 words: long-term forecasting with transformers[C]//International Conference on Learning Representations. 2023. https://openreview.net/forum?id=Jbdc0vTOcol.
10. Ji Y, Luo Y, Lu A, Xia D, Yang L, Liew A W-C. Galformer: a transformer with generative decoding and a hybrid loss function for multi-step stock market index prediction[J]. Scientific Reports, 2024, 14(1): 23762. DOI: 10.1038/s41598-024-72045-3.
11. Chen D Y, Mao L. Research on stock price prediction integrating incremental learning and Transformer model[J]. Journal of Frontiers of Computer Science and Technology, 2024, 18(7): 1889-1899. DOI: 10.3778/j.issn.1673-9418.2401037.
12. Jing N, Wu Z, Wang H. A hybrid model integrating deep learning with investor sentiment analysis for stock price prediction[J]. Expert Systems with Applications, 2021, 178: 115019. DOI: 10.1016/j.eswa.2021.115019.
13. Lee H, Kim J H, Jung H S. Deep-learning-based stock market prediction incorporating ESG sentiment and technical indicators[J]. Scientific Reports, 2024, 14(1): 10262. DOI: 10.1038/s41598-024-61106-2.
14. Wan Y P, Zhao Y, Zhao L Y. Stock returns and investor sentiment: machine learning for textual analysis of online social media[J]. Financial Market Research, 2024(5): 65-77. Available: https://www.nafmii.org.cn/yj/jrscyj/qk/2024/202405/202406/P020250324548086473193.pdf.
15. Gao Y, Zhao C, Wang Y. Investor sentiment and stock returns: new evidence from Chinese carbon-neutral stock markets based on multi-source data[J]. International Review of Economics & Finance, 2024, 92: 438-450. DOI: 10.1016/j.iref.2024.02.049.
16. Cheng D, Yang F, Xiang S, Liu J. Financial time series forecasting with multi-modality graph neural network[J]. Pattern Recognition, 2022, 121: 108218. DOI: 10.1016/j.patcog.2021.108218.
17. Lei Z, Zhang C, Xu Y, Li X. DR-GAT: dynamic routing graph attention network for stock recommendation[J]. Information Sciences, 2024, 654: 119833. DOI: 10.1016/j.ins.2023.119833.
18. Liao S, Xie L, Du Y, Chen S, Wan H, Xu H. Stock trend prediction based on dynamic hypergraph spatio-temporal network[J]. Applied Soft Computing, 2024, 154: 111329. DOI: 10.1016/j.asoc.2024.111329.

19. Zhang L, Dai Z, Zeng Q, Lin N, Xia H, Zhao K, Zhang M, Liu S. Dynamic graph learning for integrating temporal relationships in stock prediction[J]. Expert Systems with Applications, 2026, 317: 131857. DOI: 10.1016/j.eswa.2026.131857.
20. Kelly B T, Pruitt S, Su Y. Characteristics are covariances: a unified model of risk and return[J]. Journal of Financial Economics, 2019, 134(3): 501-524. DOI: 10.1016/j.jfineco.2019.05.001.
21. Dixon M F, Polson N G, Goicoechea K. Deep partial least squares for empirical asset pricing[EB/OL]. SSRN, 2022. DOI: 10.2139/ssrn.4137647.
22. Arnott R D, Harvey C R, Kalesnik V, Linnainmaa J. Reports of value's death may be greatly exaggerated[J]. Financial Analysts Journal, 2021, 77(1): 44-67. DOI: 10.1080/0015198X.2020.1842704.
23. Guo T, Hauptmann E. Fine-tuning large language models for stock return prediction using newsflow[C]//Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track. Miami: Association for Computational Linguistics, 2024: 1028-1045. DOI: 10.18653/v1/2024.emnlp-industry.77.
24. Lopez-Lira A, Tang Y. Can ChatGPT forecast stock price movements? Return predictability and large language models[EB/OL]. SSRN, 2023, revised 2026. DOI: 10.2139/ssrn.4412788.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

