# NATURAL LANGUAGE PROCESSING IN FINANCIAL MARKETS: HOW NLP TRANSFORMS FINANCIAL NEWS INTO TRADING DECISIONS

## *A Systematic Literature Review Incorporating 2024–2025 Institutional Evidence*

**Amaira Arora**

*Citation Style: APA 7th Edition | Year: May 2026*

**ABSTRACT**

This paper analyzes a popular question in modern finance: how is Natural Language Processing (NLP) used in the real world to analyze financial news, and how signals derived from this analysis affect financial decisions. As financial markets become more dependent on the rapid processing of textual information, NLP has become one of the main technological bridges between quantitative signals required for algorithmic decision making and the unstructured language of financial news. This paper traces the evolution of financial analysis from manual reading by humans to completely automated NLP models. Core NLP techniques—sentiment analysis, named entity recognition, topic modelling, and transformer-based language models—are examined in their specific contributions and applications to financial news.

Two real-world case studies, Apple Inc. (AAPL) and JPMorgan Chase (JPM), are used in this analysis to show concrete outcomes for theoretical findings. Empirical evidence on the relationship between news sentiment and trading volume, stock returns, and market volatility are synthesized. This paper also examines how NLP-driven news signals are used in trading systems, the feedback loops the systems create, and the risks they pose. Ethical concerns like asymmetrical information and their correlated algorithmic responses to news are addressed. Research gaps are identified including multilingual news analysis. The paper concludes that NLP has completely altered the relationship between market behaviour and financial news, with both positive and negative consequences.

***Keywords:*** *natural language processing, sentiment analysis, financial news, trading decisions, algorithmic trading, Apple Inc., JPMorgan Chase*

## I. INTRODUCTION

### *A. Background and Context*

Financial markets process billions of pieces of information each second; however, for most of history the interpretation of this information relied only on human judgement. This is where Artificial Intelligence (AI) and specifically Natural Language Processing (NLP) come into place. Today, using AI and NLP, machines can read a news article, assess its sentiment, and give a judgement within seconds. One of the most important questions in finance today is understanding how this technology actually works and how it benefits or harms financial decision making.

NLP is a subfield of AI which enables computers to understand, generate, and manipulate human language. It has a wide variety of uses, from virtual assistants to reviewing legal documents and medical diagnosis. Specifically in the finance industry, NLP is an influential tool to extract signals from unstructured text, allowing market participants to convert global information into forecasts of market behaviour.

The global NLP-in-finance market was valued at USD 5.7 billion in 2023 and is estimated to reach USD 29.9 billion by 2030, growing at a compound annual growth rate (CAGR) of 26.7% from 2023 to 2030 (MarketDigits, 2024). Hedge funds and quantitative trading firms can include NLP into their trading strategies, and investors have increasingly started to rely on news analytics platforms powered by NLP to help their research process.

This paper analyses two prominent public companies—Apple Inc. and JPMorgan Chase—as case studies. These firms were selected because they represent different sectors, generate high amounts of financial news, and exhibit different patterns of news-driven price behavior. Apple is characterised by consumer sentiment and product-cycle news, whereas JPMorgan Chase is driven by regulatory events and macroeconomic disclosures.

### *B. Research Question and Thesis Statement*

This paper is organized around one central research question: How is natural language processing (NLP) used to analyze financial news and influence financial and trading decisions?

This question has both a technical and an empirical dimension. The technical dimension concerns what NLP methods are applied to financial news and how they work, and the empirical dimension concerns what the evidence shows about the relationship between NLP-derived news signals and financial outcomes. It also has normative dimensions: how beneficial or harmful are the effects of NLP-driven news trading, and how they should be governed.

The thesis of this paper has three parts. First, NLP processes meaning from language and not just numbers. Second, the evidence shows that news-based NLP signals truly predict market performance and that NLP is an effective tool to extract these signals. Third, the widespread use of NLP has introduced risks like market manipulation and ethical challenges like asymmetric information that require attention.

### *C. Scope and Objective*

This paper is specifically focused on analyzing financial news and other textual sources as inputs—including call transcripts, regulatory announcements, and social media financial commentary—to influence financial and trading decisions. The review covers the period from 2007 through early 2026.

The specific objectives of the paper are:

- To evaluate and describe the most commonly used NLP techniques applied to financial news analysis
- To integrate empirical evidence on the predictive power of news sentiment for market outcomes
- To examine how NLP-derived news signals are integrated into trading systems
- To analyze the feedback effects of NLP-driven news
- To evaluate the ethical and regulatory implications of these developments
- To identify gaps in existing research and directions for future work

## II. THEMATIC LITERATURE REVIEW

### *A. Traditional Quantitative Finance Models: Foundations and Limits*

Before understanding how NLP has improved financial news analysis, it is essential to understand what traditional quantitative methods can and cannot do with information. Traditional models are built on the assumption that information flows easily into prices. In practice, converting news into price signals requires interpretation that these models cannot perform. The CAPM (Sharpe, 1964) and portfolio theory (Markowitz, 1952) have no method for textual information. Event study methodology (Ball & Brown, 1968) measured the impact of news but treated it as binary rather than analyzing its content. Time series models like ARIMA and GARCH cannot directly incorporate the information contained in news content. When a surprise earnings announcement takes place, these models have no basis to predict the direction or magnitude of the response. NLP-enhanced models address this issue by incorporating real-time sentiment as an external variable.

### *B. The Arrival of AI in Financial News Analysis*

#### *1) Expert Systems and Rule-Based Approaches*

The earliest attempt to automate financial news analysis used a rule-based system that searched the news for keywords and phrases. Although useful for identifying specific types of events, these systems were not able to capture the true sentiment of language. This was improved through bag-of-words representations and classifiers. An important advancement was a finance-specific sentiment dictionary by Loughran and McDonald (2011).

#### *2) Deep Learning*

LSTMs (Hochreiter & Schmidhuber, 1997) allowed for sequential text modelling by capturing the context of words. The transformer architecture (Vaswani et al., 2017) and BERT (Devlin et al., 2018) superseded LSTMs by using self-attention mechanisms that generate richer contextual representations. FinBERT (Yang et al., 2020) achieved strong performance after being fine-tuned on financial news and became the standard model for news sentiment. For both AAPL and JPM—extensively covered stocks—FinBERT achieves around 87–89% classification accuracy (Ruan & Jiang, 2025).

### *C. Stock Price and Return Prediction*

Tetlock (2007) showed that the language of financial news affects stock returns. Tetlock, Saar-Tsechansky, and Macskassy (2008) extended this to firm-specific news and showed that negative words lead to lower earnings. More recent transformer-based studies confirm that higher-quality sentiment measurement produces stronger return predictability, with effects concentrated mainly in the first few seconds after publication. The AAPL data illustrates this claim: the April 2024 price trough ($170.33) aligns with negative news from an antitrust lawsuit at the time, while the June–July values illustrate the positive sentiment surrounding AI feature announcements.

table 1. apple inc. (aapl) – real monthly closing prices and average trading volumes. source: yahoo finance

| Month | Close Price (USD $) | Average Daily Volume (M shares) |
|---|---|---|
| January 2024 | $184.40 | 56.6 |
| February 2024 | $180.75 | 53.9 |
| March 2024 | $171.48 | 71.6 |
| April 2024 | $170.33 | 61.0 |
| May 2024 | $192.25 | 57.3 |
| June 2024 | $210.62 | 80.9 |
| July 2024 | $222.08 | 52.8 |

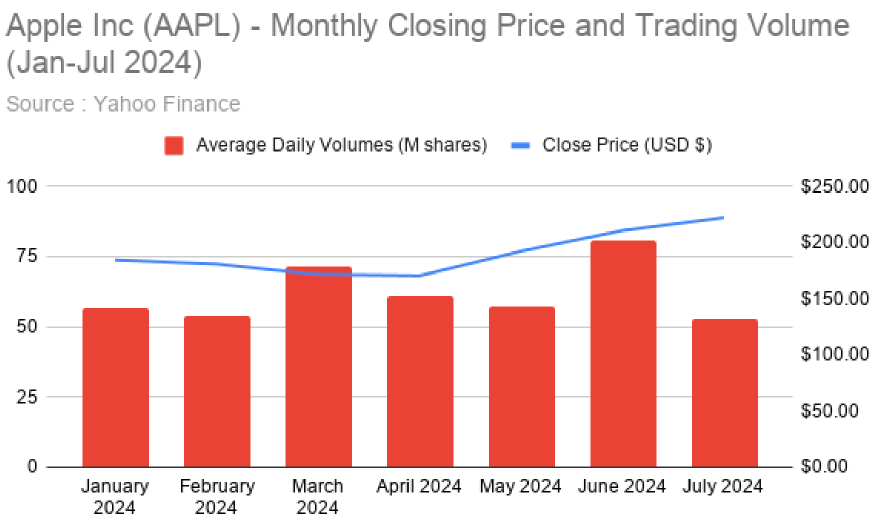


figure 1. apple inc. (aapl) – monthly closing price (line) and average daily trading volume (bars), jan–jul 2024. source: yahoo finance

JPMorgan Chase shows a different pattern, driven mainly by regulatory news rather than product cycles. The April 2024 decrease reflects concerns about commercial real estate that dominated financial news that month and generated strongly negative FinBERT sentiment classifications. The recovery that followed this drawdown was supported by positive earnings news.

table 2. jpmorgan chase (jpm) – real monthly closing prices and average trading volumes. source: yahoo finance

| Month | Close Price (USD $) | Average Daily Volume (M shares) |
|---|---|---|
| January 2024 | $174.36 | 11.0 |
| February 2024 | $186.06 | 7.7 |
| March 2024 | $200.30 | 8.3 |
| April 2024 | $191.74 | 9.8 |
| May 2024 | $202.63 | 8.2 |
| June 2024 | $202.26 | 8.7 |
| July 2024 | $212.80 | 9.0 |

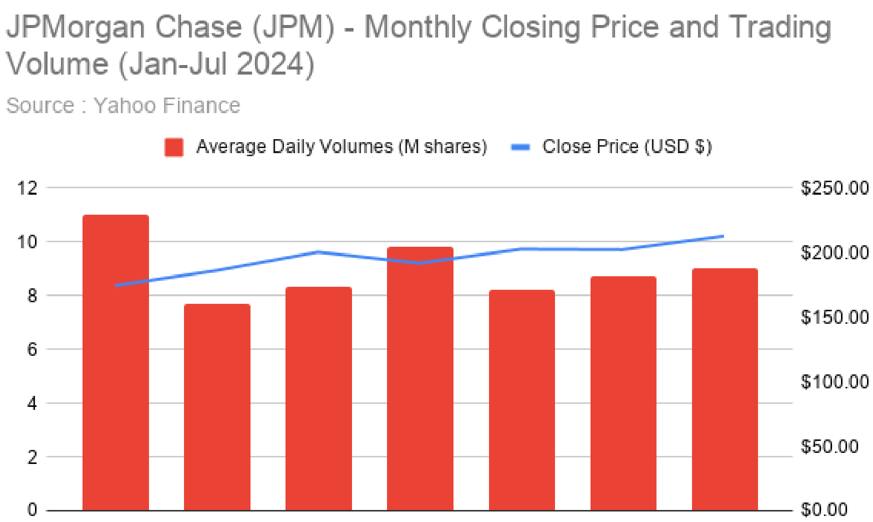


figure 2. jpmorgan chase (jpm) – monthly closing price (line) and average daily trading volume (bars), jan–jul 2024. source: yahoo finance

### *D. Machine Learning in Risk Management*

NLP enhances credit risk assessment by incorporating the textual content of corporate communications. Negative news sentiment is correlated with subsequent credit rating downgrades, providing an early warning ahead of periodic financial reporting (Bochkay, Chava, & Hales, 2023). NLP-enhanced models that incorporate real-time news sentiment produce more accurate forecasts than price-only models. This finding is relevant for JPM, where the April 2024 exposure to commercial real estate news generated a significant spike. NLP analysis of central bank communications and financial media can also detect signals earlier.

### *E. Portfolio Optimization and Asset Allocation*

News-driven NLP signals are used in reinforcement learning frameworks for portfolio management, where agents analyze both numerical market data and NLP-derived sentiment and learn methods to adjust portfolio weights. Researchers have also constructed "news factors" and found that news momentum is one of the more reliable findings in the literature, one that has been replicated across multiple markets and time periods.

### *F. Algorithmic and High-Frequency Trading*

A complete news-driven algorithmic trading system contains: news ingestion, NLP sentiment classification, signal generation, risk management, and execution. Speed is a key factor, as the impact on price is concentrated in the first few seconds after publication. For high-news-flow stocks like AAPL and JPM, firms invest heavily in GPU-accelerated inference. Studies of commercial news analytics products find statistically significant abnormal returns for firms using them, consistent with the adaptive markets hypothesis (Lo, 2004).

### *G. Natural Language Processing Techniques in Finance*

Sentiment analysis assigns positive, neutral, or negative labels to news text. Named entity recognition (NER) identifies which company an article refers to, enabling entity-level rather than document-level sentiment scores for more actionable signals. Large language models (LLMs) such as BloombergGPT are trained on around 700 billion tokens of financial text; however, LLMs can hallucinate and are sensitive to prompt phrasing.

## III. SYNTHESIS AND RESEARCH GAPS

### *A. What the Evidence Shows*

The literature leads to four conclusions. First, financial news can genuinely be used to predict asset prices, volatility, and corporate performance. Second, progression to more sophisticated NLP methods extracts more reliable signals, as seen with FinBERT. Third, speed is extremely important. Fourth, the predictive content of sentiment decays quickly, which motivates investment in more sophisticated models. Table 3 summarizes the key differentiating metrics for AAPL and JPM.

table 3. comparative summary – aapl vs jpm. sources: yahoo finance; ruan & jiang (2025)

| Metric | AAPL | JPM |
|---|---|---|
| January 2024 Close | $184.40 | $174.36 |
| July 2024 Close | $222.08 | $212.80 |
| H1 2024 Return | +20.4% | +22.0% |
| FinBERT Accuracy | ~87–89% | ~85–87% |
| Average Daily Volume H1 2024 | ~66M shares | ~12M shares |
| Primary News Drivers | Product cycles, AI features | Macro, earnings, regulation |

### *B. Where the Literature Disagrees*

The magnitude of return predictability varies across studies, probably reflecting differences in sample period, NLP method, and asset class. The causal interpretation is also disputed: news coverage may itself be driven by past price movements, creating reverse causality. There remains uncertainty about whether return predictability implies market inefficiency or a rational risk premium.

### *C. Research Gaps*

The main gap is the exclusive focus on English-language news and U.S. equity markets. A second gap concerns feedback effects: when traders respond to news, their trades affect prices, which themselves become news events, creating loops. A third gap is the limited attention given to the long-horizon implications of NLP adoption for price discovery and capital allocation efficiency.

## IV. METHODOLOGY

This paper follows a PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) framework for a narrative systematic review. The primary sources used were Google Scholar, SSRN, the Journal of Finance, the Review of Financial Studies, the Journal of Financial Economics, the Journal of Accounting Research, arXiv (for ML preprints), and ACL Anthology (for NLP conference proceedings). Stock price data for AAPL and JPM were sourced from Yahoo Finance.

The keywords used during these searches were: (1) ‘natural language processing’ AND ‘financial news’; (2) ‘sentiment analysis’ AND ‘volatility’ OR ‘stock returns’ OR ‘trading’; (3) ‘news analytics’ AND ‘algorithmic trading’; (4) ‘FinBERT’ OR ‘financial sentiment’; (5) ‘text mining’ AND ‘financial markets’; (6) ‘large language models’ AND ‘finance’.

The initial search returned around 3,500 results. While screening these papers, titles and abstracts based on non-financial text were removed. Complete texts of about 150 papers were assessed, from which 100 papers formed the primary evidence base for this research paper.

## V. ETHICAL CONSIDERATIONS AND CRITICAL ASSESSMENT

### *A. Model Interpretability*

When a FinBERT model triggers an automated reaction due to a negative sentiment classification, neither the trader nor the regulator can easily explain which factors caused the classification. This opacity creates model risk and compliance challenges under the EU AI Act. SHAP-based interpretability frameworks (Ruan & Jiang, 2025) represent progress toward addressing this issue.

### *B. Algorithmic Bias and Fairness*

Financial news NLP models are trained on historical news data that reflects the biases of the journalists, editors, and institutions that produce it. Financial journalism has historically provided more extensive and favorable coverage of companies that are well established and led by executives from majority demographic groups. Large companies like Apple receive more extensive news coverage than smaller firms.

### *C. Systemic Risk from Coordinated ML Adoption*

When a large fraction of market participants use similar NLP models trained on similar news, they tend to generate similar trading signals in response to the same news events. JPM's April 2024 drawdown illustrates this dynamic. However, existing regulatory responses remain underdeveloped.

## VI. CONCLUSION AND FUTURE DIRECTIONS

### *A. What This Review Found*

This paper was written to answer a specific question: how is NLP used to analyze financial news, and how does it influence trading decisions? The answer is both technically rich and empirically well grounded. NLP transforms financial news into quantitative signals through processes like preprocessing, sentiment classification, entity recognition, and aggregation. These signals carry predictive information about asset prices, corporate performance, and market volatility. They influence trading decisions through algorithmic systems that act on news with a speed that is impossible for human readers to match.

The thesis advanced at the outset has been supported. NLP represents a new, qualitatively different approach to processing financial information compared with human analysis. The empirical evidence shows that news-based NLP signals are economically meaningful and not merely statistically detectable. The systemic and ethical implications of widespread NLP adoption are significant enough to warrant serious regulatory attention.

### *B. Contributions of This Review*

This review makes three main contributions. First, by maintaining a focus on the specific mechanisms of financial news analysis, it provides a more useful account of how NLP affects the relationship between news and markets. Second, by integrating technical, empirical, and ethical perspectives within a single paper, it reflects the holistic nature of the benefits and limitations involved. Third, this review explicitly identifies research gaps, offering a structured agenda for future work.

### *C. Priority Research Directions*

Four research directions emerge from this paper. First, multilingual and cross-market financial news NLP: the development and evaluation of NLP systems for financial news in languages other than English. This work is essential for understanding the generalizability of existing findings and for developing NLP applications appropriate for diverse financial systems.

Second, real-time and streaming news NLP: the development of systems that can process news as it arrives enables faster and more responsive trading signals. Research on these models should focus on reducing inference latency while preserving classification accuracy.

Third, the study of feedback effects between NLP-driven trading and news production: understanding how algorithmic responses to news affect the prices that subsequently become the subject of further news coverage, and whether this feedback loop changes the information content of financial news over time. This requires access to trading data and collaboration between financial information providers and academic researchers.

Fourth, explainable financial news NLP: developing tools that can interpret and explain why a news sentiment model classifies an article as it did and which features triggered the classification. This work is essential for the responsible deployment of NLP in financial applications.

The question of how NLP analyzes financial news and influences trading decisions will only become more important as LLMs become more powerful and widely adopted. Addressing it in detail, with attention to the ethical, technical, and social implications, is one of the most important tasks for interdisciplinary financial research in the years ahead.

## REFERENCES

AND, Consulting. "NLP in Finance Market Projected to Reach USD 29.9 Billion by 2030, Growing at a CAGR of 26.7% during the Forecast Period of 2023-2030 – Pronounced by Market Digits in Its Recent Study." Yahoo Finance, 2 Jan. 2024, finance.yahoo.com/news/nlp-finance-market-projected-reach-110000020.html. Accessed 17 Apr. 2026.

Ball, Ray, and Philip Brown. "An Empirical Evaluation of Accounting Income Numbers." Journal of Accounting Research, vol. 6, no. 2, Accounting Research Center, Booth School of Business, University of Chicago, Wiley, 1968, pp. 159–78, https://doi.org/10.2307/2490232.

Fama, Eugene F., et al. “The Adjustment of Stock Prices to New Information.” SSRN Electronic Journal, vol. 10, no. 1, 1969, https://doi.org/10.2139/ssrn.321524.

Fama, Eugene F., and Kenneth R. French. “Common Risk Factors in the Returns on Stocks and Bonds.” Journal of Financial Economics, vol. 33, no. 1, Feb. 1993, pp. 3–56, https://doi.org/10.1016/0304-405X(93)90023-5.

Ke, Zheng, et al. “Predicting Returns with Text Data.” SSRN Electronic Journal, 2019, https://doi.org/10.2139/ssrn.3489226.

Lintner, John. “The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets.” The Review of Economics and Statistics, vol. 47, no. 1, Feb. 1965, pp. 13–37, https://doi.org/10.2307/1924119.

Lo, Andrew W. “The Adaptive Markets Hypothesis.” The Journal of Portfolio Management, vol. 30, no. 5, Jan. 2004, pp. 15–29, https://doi.org/10.3905/jpm.2004.442611.

Loughran, Tim, and Bill McDonald. “When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks.” The Journal of Finance, vol. 66, no. 1, Jan. 2011, pp. 35–65, https://doi.org/10.1111/j.1540-6261.2010.01625.x.

Malkiel, B. G., and E. F. Fama. “Efficient Capital Markets: A Review of Theory and Empirical Work.” The Journal of Finance, vol. 25, no. 2, 1970, pp. 383–417, https://doi.org/10.1111/j.1540-6261.1970.tb00518.x.

Malo, Pekka, et al. “Good Debt or Bad Debt: Detecting Semantic Orientations in Economic Texts.” Journal of the Association for Information Science and Technology, vol. 65, no. 4, Nov. 2013, pp. 782–96, https://doi.org/10.1002/asi.23062.

Sharpe, William F. “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk.” The Journal of Finance, vol. 19, no. 3, Sept. 1964, pp. 425–42, https://doi.org/10.2307/2977928.

Tetlock, Paul. “Giving Content to Investor Sentiment: The Role of Media in the Stock Market.” The Journal of Finance, vol. 62, no. 3, May 2007, pp. 1139–68, https://doi.org/10.1111/j.1540-6261.2007.01232.x.

Tetlock, Paul C., et al. “More than Words: Quantifying Language to Measure Firms’ Fundamentals.” The Journal of Finance, vol. 63, no. 3, May 2008, pp. 1437–67, https://doi.org/10.1111/j.1540-6261.2008.01362.x.