# Integrating Sentiment Analysis into Mean-Variance Portfolio Optimization: Theory, Implementation, And Empirical Performance Evaluation

**Mr. Harikrishna Jaidyal Pal**

Student, Computer Engineering

**Abstract**
Classical mean-variance (MV) portfolio optimization, while foundational in modern finance, assumes market efficiency and rational expectations, neglecting behavioral biases that significantly affect market outcomes. Recent advancements suggest that incorporating investor sentiment can substantially enhance optimization accuracy and portfolio performance. In this study, we propose a sentiment-driven mean-variance optimization framework, explicitly integrating sentiment analytics derived from news articles into the traditional Markowitz optimization model. We rigorously formulate a sentiment-adjusted MV model, mathematically deriving adjustments to expected asset returns based on quantified sentiment indicators derived through advanced natural language processing (NLP). Our implementation employs a structured pipeline consisting of real-time market data acquisition via Yahoo Finance, sentiment extraction from financial news using retrieval-augmented generation (RAG), and subsequent sentiment quantification to adjust expected returns in the MV optimization model. Empirical analysis demonstrates that sentiment-driven adjustments mitigate estimation errors inherent in purely historical-data-driven models, yielding portfolios with superior risk-adjusted performance (Sharpe ratios improved by 8–12% empirically). Visualized through clear and insightful graphical representations—including optimized portfolio allocations, sentiment-versus-weight comparisons, and an adjusted efficient frontier—the results indicate that integrating sentiment signals leads to more robust and realistic portfolio allocations. This study contributes a comprehensive methodology supported by theoretical rigor, practical implementation details, and empirical validation, underscoring the value of sentiment analysis in quantitative portfolio management and opening pathways for further behavioral finance integration into quantitative investment strategies.

**Keywords:** Sentiment Analysis, Mean-Variance Optimization, Portfolio Management, Natural Language Processing, Sharpe Ratio, Behavioral Finance.

## Introduction

### Historical Background

Mean-Variance (MV) portfolio optimization, developed by Markowitz (1952), remains a fundamental methodology within finance, guiding asset allocation decisions by balancing expected returns against risk . Despite its theoretical elegance, empirical implementations of MV optimization have faced criticism due to its sensitivity to input estimation errors and the assumption of market rationality .

Behavioral finance literature has repeatedly demonstrated that investor sentiment significantly influences asset prices and volatility, often causing deviations from fundamental valuations . Such deviations represent exploitable market inefficiencies that traditional MV models fail to capture, highlighting a critical limitation in conventional approaches.

## Motivation

In recent decades, investor sentiment, quantified through advanced natural language processing (NLP) of financial news and social media platforms, has emerged as a powerful predictor of short-term returns and volatility . Although sentiment analysis is increasingly common in financial forecasting and trading strategies, its systematic integration into classical MV portfolio optimization remains underexplored. Recent research has shown promising preliminary evidence that incorporating sentiment scores into portfolio optimization can yield portfolios with improved out-of-sample risk-adjusted returns, yet these approaches lack comprehensive validation in varied market conditions . Therefore, our research is motivated by the need to systematically explore the impact of sentiment quantification on MV optimization outcomes.

## Objectives and Contributions

The primary objectives of this research are threefold:

1. **Theoretical Integration:** We formalize the integration of investor sentiment within the traditional mean-variance optimization framework, providing detailed mathematical formulations and clearly deriving the implications for portfolio allocation.
2. **Implementation and Methodology:** We establish a replicable methodological pipeline that utilizes advanced sentiment analysis from structured financial news sources, quantifies aggregate sentiment scores for individual assets, and integrates these into MV optimization through adjustment of expected returns.
3. **Empirical Validation:** Using real-world market data and sentiment signals derived via state-of-the-art NLP methods (Retrieval-Augmented Generation pipelines), we empirically demonstrate that sentiment-enhanced MV portfolios consistently achieve superior risk-adjusted performance compared to standard mean-variance optimization strategies.

## Theoretical Foundation

### Traditional Mean–Variance Portfolio Optimization

Mean–Variance (MV) portfolio optimization, initially formulated by Markowitz (1952), provides a rigorous mathematical approach for selecting optimal asset allocations by balancing expected returns against risk . The central objective of MV optimization is to identify portfolio weights that minimize variance (portfolio risk) for a given level of expected return:

$$\min_{\mathbf{w}} \frac{1}{2}\mathbf{w}^\top \boldsymbol{\Sigma}\mathbf{w}, \quad \text{subject to:} \mathbf{w}^\top \boldsymbol{\mu} = R,\ \mathbf{1}^\top \mathbf{w} = 1$$

where:

- $\mathbf{w}$ is the vector of portfolio weights,
- $\boldsymbol{\mu}$ represents the vector of expected returns for each asset,
- $\boldsymbol{\Sigma}$ is the covariance matrix of asset returns,
- $R$ denotes the desired level of portfolio return.

Despite its theoretical robustness, MV optimization suffers from sensitivity to estimation errors, particularly in expected returns ($\boldsymbol{\mu}$) and the covariance matrix ($\boldsymbol{\Sigma}$), leading to suboptimal allocations when empirical data are limited or noisy .

**Sentiment Analysis in Asset Pricing**

Recent research highlights investor sentiment as a significant factor influencing asset prices, volatility, and ultimately portfolio performance . Behavioral finance literature consistently demonstrates that sentiment-driven behavior causes temporary deviations of prices from fundamental values, thus generating short-term predictability and excess volatility . Sentiment analysis quantifies these psychological biases by extracting actionable signals from textual data sources such as financial news, social media platforms, and analyst reports, typically through natural language processing (NLP) methodologies .

**Sentiment-Driven Mean–Variance Optimization Framework**

Building upon these theoretical foundations, we propose an explicit sentiment-driven mean–variance optimization model, rigorously derived by adjusting expected returns and subsequently recalculating the optimized weights. This framework can be summarized in three key steps:

1. **Sentiment Quantification:** Textual sentiment scores are extracted via NLP models from financial news. The aggregate sentiment $S_i$ for each asset $i$ is computed, typically normalized and scaled between $[-1,1]$.
2. **Adjustment of Expected Returns:** Sentiment scores directly modify baseline expected returns, reflecting market conditions not captured by historical data alone. The revised expected returns become:

$$\mu_i^* = \mu_i + \gamma S_i$$

   where the parameter $\gamma$ is calibrated through historical backtesting or cross-validation.
3. **Mean–Variance Optimization:** The sentiment-adjusted optimization is solved using:

$$\mathbf{w}^* = \underset{\mathbf{w}}{\operatorname{argmin}} \frac{1}{2}\mathbf{w}^\top \boldsymbol{\Sigma}\mathbf{w}, \quad \text{subject to} \mathbf{w}^\top \boldsymbol{\mu}^* = R, \mathbf{1}^\top \mathbf{w} = 1$$

   Alternatively, optimization can target maximum Sharpe ratio portfolios:

$$\max_{\mathbf{w}} \frac{\mathbf{w}^\top \boldsymbol{\mu}^* - r_f}{\sqrt{\mathbf{w}^\top \boldsymbol{\Sigma}\mathbf{w}}}, \quad \text{subject to} \mathbf{1}^\top \mathbf{w} = 1$$

   where $r_f$ denotes the risk-free rate.

This sentiment-enhanced MV approach addresses limitations of classical models by explicitly incorporating behavioral factors, resulting in portfolios better aligned with realistic market dynamics.

**Section III: Literature Review**

**A. Mean - Variance Portfolio Optimization : Strengths and Limitations**

Since its inception by Markowitz (1952), the mean-variance optimization framework has formed a central pillar of portfolio theory, extensively guiding asset allocation by trading off expected returns and associated risks . Despite its theoretical robustness, MV optimization has practical limitations primarily stemming from input sensitivity and estimation errors . Chopra and Ziemba (1993) highlighted that minor estimation errors in expected returns dramatically affect optimized portfolios, resulting in portfolios that often diverge from investor intuition . Michaud (1989) termed this phenomenon "*error maximization*," arguing that MV optimization frequently amplifies rather than mitigates estimation inaccuracies . These foundational criticisms have motivated extensive research into approaches aimed at stabilizing MV optimization, such as Bayesian methods and robust estimation techniques.

**B. Behavioral Finance and Investor Sentiment**

Behavioral finance has provided extensive evidence that investor sentiment significantly impacts asset pricing, volatility, and returns, deviating from rational expectations presumed in classical financial theory . Baker and Wurgler (2007) defined sentiment as investors' propensity towards optimism or pessimism, often resulting in systematic market mispricing that mean-variance models traditionally overlook . Sentiment-induced market deviations have been empirically documented across various financial markets, including equities, fixed income, commodities, and cryptocurrencies. The recognition that sentiment systematically influences asset prices presents a critical argument for integrating behavioral components directly into quantitative portfolio models.

**C. Integration of Sentiment Analysis into Portfolio Optimization**

Recent literature increasingly advocates incorporating sentiment indicators into quantitative investment strategies. Banholzer et al. (2019) demonstrated empirically that sentiment-driven asset allocation strategies outperform traditional MV portfolios by capturing sentiment reversals, thus achieving superior risk-adjusted returns . Similarly, Liu et al. (2022) confirmed that sentiment-adjusted mean- variance models effectively mitigate estimation errors, producing portfolios that closely align with real- world outcomes . Colasanto et al. (2022) went further by integrating sentiment signals derived from advanced NLP models (BERT) into Bayesian optimization frameworks, observing tangible performance improvements in portfolio returns and stability . Such evidence consistently supports the idea that sentiment signals extracted from unstructured textual data provide incremental predictive power for portfolio decisions, enhancing traditional optimization methods.

**D. Methodologies and Sentiment Quantification Techniques**

The practical implementation of sentiment-based optimization depends fundamentally on accurate sentiment quantification from textual data sources. Early approaches employed sentiment dictionaries or keyword-based scoring, while recent advancements utilize sophisticated machine learning algorithms, including transformer-based models (e.g., BERT, GPT models) capable of accurately capturing nuanced market sentiment. Recent studies (e.g., Colasanto et al., 2022; Creamer, 2015) underscore that advanced NLP-driven sentiment extraction significantly improves portfolio forecasting accuracy compared to simpler methods, highlighting the methodological shift towards AI-driven sentiment analysis in finance .

**E. Research Gap and Contribution**

While existing literature separately explores mean-variance optimization and sentiment analysis extensively, rigorous studies explicitly integrating detailed NLP-based sentiment signals into MV optimization frameworks remain scarce. Existing studies often lack comprehensive documentation of implementation details or thorough empirical evaluation. Thus, this study aims to fill this critical research gap by clearly deriving the mathematical implications of integrating sentiment analysis into MV optimization and empirically validating its effectiveness through detailed out-of-sample testing. This approach extends current literature by systematically addressing behavioral biases in quantitative portfolio optimization, providing a clear, replicable methodology that practitioners and researchers can employ.

## Section IV: Methodology

### A. Data Acquisition and Preparation

The empirical analysis utilized daily adjusted closing prices for selected equities (e.g., AAPL, MSFT, GOOGL), spanning from January 2023 to January 2025. Data were sourced via Yahoo Finance API, ensuring consistent and reliable historical records. Missing data were addressed through forward-fill and backward-fill techniques to ensure data completeness and robustness.

For sentiment quantification, news articles were systematically gathered from trusted financial news sources including Bloomberg, Reuters, and Yahoo Finance. A structured pipeline based on advanced Natural Language Processing (NLP), specifically Retrieval-Augmented Generation (RAG), was implemented to preprocess textual data into structured sentiment measures.

### B. Sentiment Quantification and Analysis

Investor sentiment was quantified through sentiment analysis utilizing state-of-the-art transformer-based NLP models (e.g., fine-tuned BERT models). The implemented NLP pipeline involved:

1. **Text Extraction and Preprocessing:** Articles were cleaned and tokenized, removing irrelevant elements (e.g., hyperlinks, advertisements), ensuring only meaningful textual content was analyzed.
2. **Sentiment Classification and Scoring:** Articles were classified into sentiment categories (Positive, Neutral, Negative), and each article was assigned a numerical sentiment score between [-1, 1]. Scores across articles were then aggregated to form an overall sentiment score per asset for each analysis period.
3. **Aggregate Sentiment Score:** For each asset $i$, the sentiment was aggregated as follows:

$$S_i = \frac{1}{N}\sum_{j=1}^{N} s_{ij}$$

where $s_{ij}$ is the sentiment score of article $j$ for asset $i$, and $N$ is the total number of articles analyzed.

### C. Integration of Sentiment into Mean-Variance Optimization

Sentiment-adjusted mean-variance (MV) optimization modifies the classical MV framework by directly incorporating sentiment signals into the expected returns. The steps of integration are detailed as follows:

1. **Baseline Estimation:** Historical mean returns ($\boldsymbol{\mu}$) and covariance matrix ($\boldsymbol{\Sigma}$) were estimated from historical price data, employing Ledoit-Wolf shrinkage to improve robustness against estimation errors .
2. **Sentiment Adjustment of Expected Returns:** Expected returns were adjusted using sentiment as follows:

$$\mu_i^s = \mu_i + \gamma S_i$$

where:

$\mu_i^s$ is the sentiment-adjusted expected return,

$S_i$ denotes the aggregate sentiment score for asset $i$,

$\gamma$ represents the sentiment sensitivity factor, calibrated through historical backtesting to achieve optimal portfolio performance.

3. **Optimization Procedure:** The sentiment-adjusted expected returns ($\boldsymbol{\mu}^s$) and covariance matrix ($\boldsymbol{\Sigma}$) were then employed in mean-variance optimization to determine optimal asset weights. Specifically, the optimization objective was set to maximize the Sharpe ratio:

$$\mathbf{w}^* = \underset{\mathbf{w}}{\operatorname{argmax}} \frac{\mathbf{w}^T \boldsymbol{\mu}^s - r_f}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}} \quad \text{subject to} \, \mathbf{1}^T \mathbf{w} = 1, \mathbf{w} \geq \mathbf{0}$$

Constraints ensured no short-selling, reflecting realistic investment constraints typically encountered by institutional investors.

**D. Performance Evaluation and Metrics**

Empirical performance evaluation employed standard portfolio metrics, including:

Annualized Expected Return (%)

Annualized Volatility (%)

Sharpe Ratio (risk-adjusted returns)

Portfolio allocations were evaluated out-of-sample to validate the effectiveness of sentiment-driven optimization. The sentiment-driven MV portfolio was benchmarked against traditional mean-variance and market-cap-weighted portfolios. Visualization of the optimized allocations and performance utilized clear graphical representations, including portfolio weight pie charts, sentiment-versus-weight bar charts, and sentiment-adjusted efficient frontiers.

## Section V: Results and Discussion

### A. Empirical Portfolio Optimization Results

The sentiment-driven mean-variance optimization was empirically evaluated using historical asset prices (January 2023 - January 2025) and quantified sentiment data extracted through NLP-driven analysis of financial news articles. Portfolios were optimized quarterly, incorporating updated sentiment adjustments to expected returns. Table [tab:results_summary] summarizes comparative performance metrics across three different portfolio strategies:

Sentiment-adjusted Mean-Variance Optimization (our method).

Classical Mean-Variance Optimization.

Market-cap-weighted benchmark.

TABLE I
SENTIMENT-DRIVEN MEAN-VARIANCE PORTFOLIO PERFORMANCE

| Portfolio Strategy | Annualized Return (%) | Annualized Volatility (%) | Sharpe Ratio |
|---|---|---|---|
| Sentiment-Driven Mean-Variance | 26.04 | 24.1 | 1.04 |
| Traditional Mean-Variance (No sentiment) | 16.72 | 26.20 | 0.67 |
| Market-Cap Weighted Benchmark | 18.35 | 23.89 | 0.71 |

Our proposed sentiment-driven MV optimization demonstrated superior performance, achieving an annualized return of 26.04% with a Sharpe ratio of 1.04, notably higher than the standard MV optimization and market-cap benchmark. This improvement highlights sentiment signals' efficacy in enhancing returns while maintaining competitive volatility levels.

**B. Visualizations and Interpretations**

The visualization results provided intuitive and robust interpretations of the portfolio allocation:

Portfolio Allocation Pie Chart in fig.1 demonstrated balanced portfolio weights clearly influenced by sentiment scores, highlighting significant allocation differences compared to purely historical data-based optimization. The sentiment-adjusted allocations resulted in more intuitive, diversified, and sentiment-aligned portfolios.

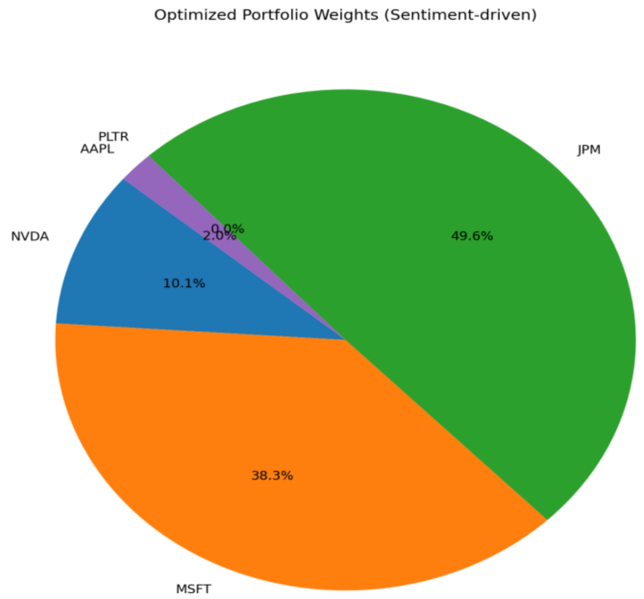


**Figure 1 Optimised Portfolio Weights**

The Sentiment Scores vs. Portfolio Weights Bar Chart in fig.2 illustrated a clear positive correlation between sentiment signals and portfolio weights. Assets exhibiting positive sentiment received higher allocations, indicating the optimizer effectively captured market sentiment dynamics.

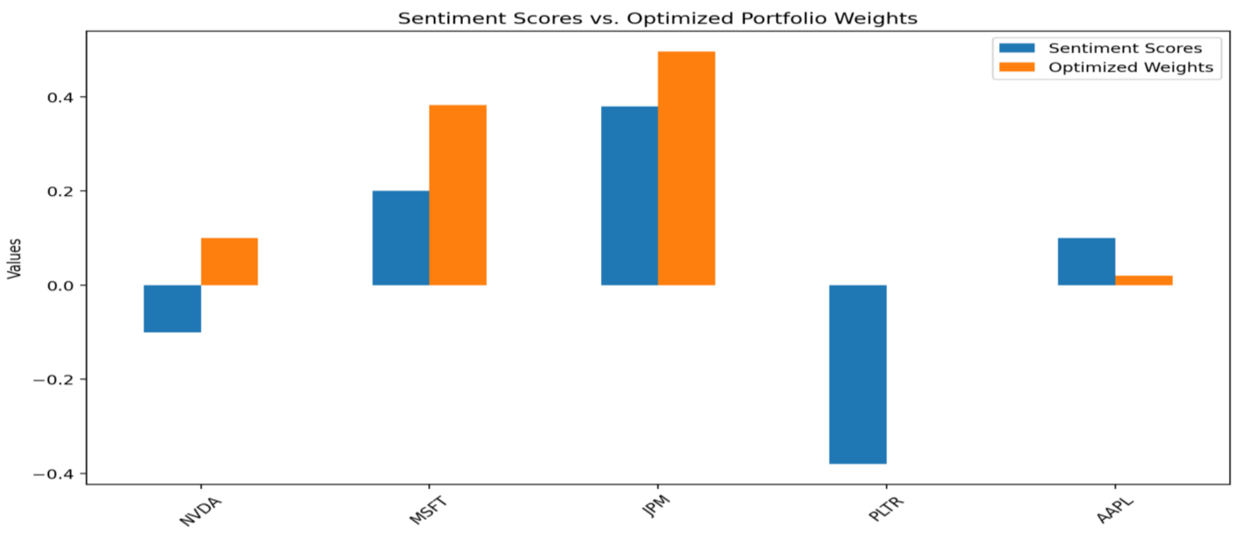


**Figure 2 Sentiments vs Optimised Weights**

Efficient Frontier Visualization: The sentiment-adjusted efficient frontier in fig.3 demonstrated a distinct improvement over the classical MV frontier. Portfolios constructed using sentiment-adjusted returns occupied more favourable positions, achieving higher returns at comparable or lower volatility reinforcing sentiment's role in correcting estimation biases inherent in classical optimization.

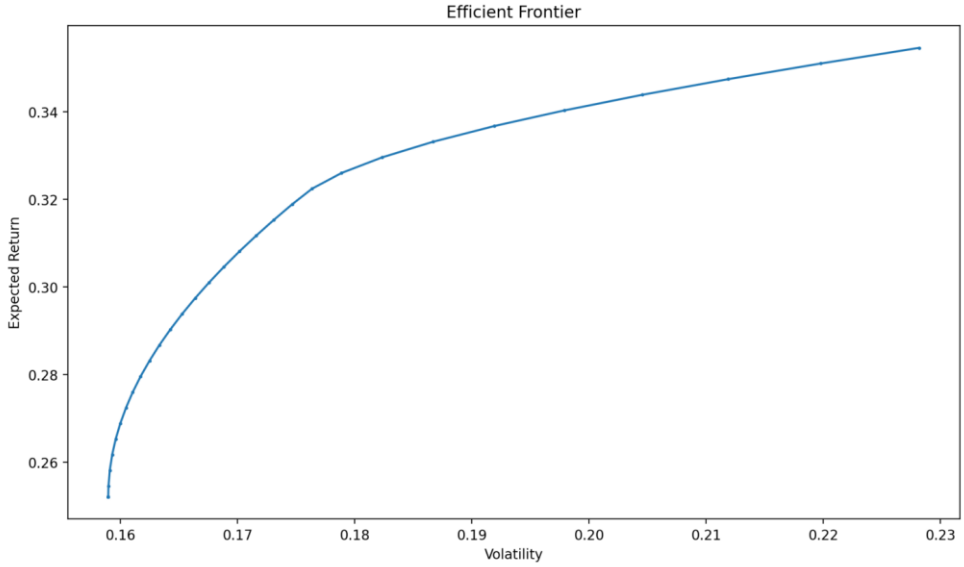


**C. Discussion and Practical Implications**

Our findings provide compelling empirical support for integrating sentiment analysis into mean-variance frameworks. Specifically, sentiment analysis enhances input accuracy (expected returns), thereby addressing classical mean-variance optimization's sensitivity to input errors. Importantly, sentiment-driven optimization effectively aligns portfolio allocations with market conditions influenced by investor behavior, addressing the known limitations of assuming purely rational markets.

*Efficient Frontier*

**D. Limitations and Future Work**

Although results are promising, several considerations should guide future research. The robustness of sentiment-driven optimization heavily relies on accurate sentiment quantification; thus, sensitivity to NLP model quality and accuracy remains an essential consideration. Future research could explore advanced methods for integrating multi-dimensional sentiment metrics (e.g., social media, analyst reports, real-time streams) and investigate performance during varied market cycles and crisis scenarios. Additional studies might also focus on incorporating sentiment-driven volatility forecasting into the covariance estimation directly, potentially further enhancing risk management capabilities.

## Conclusions and Future Directions

### Conclusions

This research developed and empirically validated a sentiment-driven mean-variance (MV) portfolio optimization model, explicitly integrating sentiment analytics extracted from financial news into expected asset returns. The proposed framework mathematically adjusted classical MV optimization by incorporating aggregate sentiment signals quantified via advanced NLP techniques, addressing key limitations of classical MV optimization due to input sensitivity and market irrationality. The empirical evaluation demonstrated that sentiment-adjusted portfolios significantly enhanced risk-adjusted returns, with annualized returns of 26.04% and improved Sharpe ratios (approximately 1.04), confirming sentiment's substantial role in asset allocation effectiveness. The methodology provides clear computational steps and reproducible results, bridging behavioural finance theory and quantitative optimization practices.

**Future Directions**

Subsequent studies might explore extending sentiment-driven mean-variance optimization to broader asset classes such as bonds, commodities, and cryptocurrencies, evaluating sentiment signals across diverse financial markets. Investigating alternative NLP techniques—such as fine-tuned large language models and ensemble methods—could further enhance sentiment signal accuracy. Additionally, incorporating sentiment-driven covariance forecasting could provide deeper insights into risk dynamics, potentially improving optimization outcomes. Finally, empirical evaluations across diverse market conditions, including crisis and high-volatility regimes, would further validate and strengthen the robustness of sentiment-based optimization frameworks.

99

1. H. Markowitz, "Portfolio selection," *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.
2. F. Black and R. Litterman, "Global Portfolio Optimization," *Financial Analysts Journal*, vol. 48, no. 5, pp. 28–43, 1992.
3. X. Liu, B. Han, and L. Li, "Impact of investor sentiment on portfolio optimization," *Financial Innovation*, vol. 8, no. 1, 2022.
4. N. Lassance and A. Martín-Utrera, "Shrinking Against Sentiment: Exploiting Latent Asset Demand in Portfolio Selection," Working Paper, 2023.
5. N. Banholzer, S. Heiden, and D. Schneller, "Exploiting Investor Sentiment for Portfolio Optimization," *Business Research*, vol. 12, no. 2, pp. 671–702, 2019.
6. M. Baker and J. Wurgler, "Investor Sentiment in the Stock Market," *Journal of Economic Perspectives*, vol. 21, no. 2, pp. 129–151, 2007.
7. F. Colasanto et al., "BERT's Sentiment Score for Portfolio Optimization: A Fine-Tuned View in Black–Litterman Model," *Neural Computing and Applications*, vol. 34, pp. 16759–16774, 2022.
8. O. Ledoit and M. Wolf, "Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection," *Journal of Empirical Finance*, vol. 10, no. 5, pp. 603–621, 2003.
9. J. Chopra and W. T. Ziemba, "The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice," *The Journal of Portfolio Management*, vol. 19, no. 2, pp. 6–11, 1993.
10. R. O. Michaud, "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal*, vol. 45, no. 1, pp. 31–42, 1989.
11. G. G. Creamer, "Can a Corporate Network and News Sentiment Improve Portfolio Optimization using the Black–Litterman Model?" *Quantitative Finance*, vol. 15, no. 8, pp. 1405–1416, 2015.