U.PORTO

FEP FACULDADE DE ECONOMIA
UNIVERSIDADE DO PORTO

THE ROLE OF MARKET SENTIMENT IN PORTFOLIO CONSTRUCTION: EVIDENCE FROM THE EQUITY MARKET

Diana Sofia Martins Fernandes

Dissertation

Master in Finance

Supervised by

Prof. Ana Paula Serra

2026

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Professor Ana Paula Serra, for her guidance and constructive feedback throughout the development of this dissertation. I would also like to thank Pedro Reis for generously sharing his expertise on the Black-Litterman model, which was instrumental in shaping the methodological approach of this work.

I am also grateful to Professor Nora Naffa, whose work and insights inspired me to pursue this research topic and explore the intersection of market sentiment and portfolio optimization.

On a personal note, I would like to thank my dear friend Catarina Marinho for her unwavering support and encouragement throughout this journey. Her presence made all the difference. I am also deeply grateful to my family and friends for their constant support, with a special word of thanks to my mother and godmother, whose love and encouragement have been a source of strength throughout this journey.

Abstract

This study examines whether standardized Bloomberg sentiment indicators improve risk-adjusted portfolio performance when embedded as investor views within a Black-Litterman optimization framework. Applied to 455 of the S&P 500 constituents over January 2024 to December 2025, composite News and Twitter sentiment scores are used to construct weekly relative return views. Different levels of confidence in the views (τ) are tested. Results show that performance improvements are contingent on conservative τ calibration: at $\tau = 0.01$, Sharpe ratios approach the benchmark (1.034-1.104 vs. 1.123) with substantially lower turnover, while performance deteriorates monotonically as τ increases. Signal diagnostics reveal that the view vector exhibits mean-reverting behavior at the weekly frequency, locating the primary constraint not in the sentiment signal itself but, in the method, used to translate it into the investor views, as the historical return spread embeds a momentum assumption that the data does not support. Positive sentiment portfolios uniformly outperform negative ones over the full sample, though subperiod analysis reveals greater resilience of negative sentiment strategies during the market stress of 2025, providing partial support for behavioral asymmetry hypotheses.

Key words: Market Sentiment; Portfolio Optimization; Black-Litterman Model; Behavioral Finance; Bloomberg Sentiment Indicators.

Resumo

Este estudo analisa se indicadores de sentimento de mercado padronizados, da Bloomberg, melhoram o desempenho ajustado ao risco de carteiras quando incorporados como perspetivas dos investidores num modelo de otimização Black-Litterman. Aplicado a 455 constituintes do S&P 500 entre janeiro de 2024 e dezembro de 2025, o estudo utiliza indicadores compostos de sentimento provenientes de notícias e do Twitter para construir perspetivas relativas de rendibilidade com frequência semanal. Foram testados diferentes níveis de confiança nas perspetivas (τ).

Os resultados mostram que as melhorias de desempenho dependem de uma calibração conservadora de τ: para $\tau = 0.01$, os rácios de Sharpe aproximam-se do benchmark (1,034 -1,104 face a 1,123), acompanhados por uma redução substancial da rotatividade da carteira. Em contrapartida, o desempenho deteriora-se de forma monotónica à medida que τ aumenta. A análise dos sinais revela que o vetor de perspetivas apresenta um comportamento de reversão à média numa frequência semanal, sugerindo que a principal limitação não reside no sinal de sentimento em si, mas no método utilizado para o traduzir em perspetivas dos investidores, uma vez que o diferencial histórico de rendibilidades incorpora uma hipótese de momentum que não é suportada pelos dados.

As carteiras associadas a sentimento positivo superam consistentemente as carteiras de sentimento negativo ao longo de todo o período analisado. Contudo, a análise por subperíodos revela uma maior resiliência das estratégias baseadas em sentimento negativo durante o período de turbulência de mercado observado em 2025, fornecendo suporte parcial para hipóteses de assimetria comportamental.

Palavras-Chave: Sentimento de Mercado; Otimização de Carteiras; Modelo Black-Litterman; Finanças Comportamentais; Indicadores de Sentimento Bloomberg.

Index

1. Introduction 1
2. Literature Review 4
2.1 Market Sentiment 4
2.2 Evolution of Portfolio Theory 8
2.2.1 Mean-Variance Optimization 8
2.2.2 Alternative Approaches to Classical Mean-Variance Optimization 10
2.2.3 Bayesian Portfolio Optimization and the Black-Litterman Framework 11
2.3 Market Sentiment-based Investment Strategies 13
3. Methodology 16
3.1 Construction of Sentiment Signals 16
3.2 Black-Litterman Portfolio Construction Framework 17
3.2.1 Market Equilibrium (prior expected returns) 18
3.2.2 Investor views (sentiment-based signals) 19
3.2.3 Posterior expected return and portfolio construction 20
3.3 Portfolio Construction 21
3.4 Construction of Model Inputs 22
3.5 Performance Evaluation 24
4. Data 25
4.1 Data Collection and Sample Period 25
4.2 Investment Universe and Filtering Criteria 25
4.3 Data Processing and Variable Construction 26
4.4 Descriptive Statistics 28
5. Results 34
5.1 RQ1: Performance of Sentiment-Informed Portfolios 36
5.2 RQ2: Asymmetry Between Positive and Negative Sentiment 38
5.3 Subperiod Analysis 42
5.4 Signal Diagnostics: The Predictive Content of the View Vector 44
5.5 Practical Implications: Portfolio Turnover and Transaction Costs 47
6. Conclusions 48
Appendix 51
A. Study Comparison 51
B. Data & Descriptive Statistics 52

C. View Vector Construction 55
D. Model Inputs & Diagnostics 56
E. Full Performance Results 59
F. Results: Subperiod Analysis 63
G. Results: Robustness 68
H. Transaction Costs Sensitivity 70
References 71

Table Index

Table 1 - Summary Statistics of Main Variables 28
Table 2 - Descriptive Statistics of Portfolio Return Differentials Used to Construct the View Vector (Q) 29
Table 3 - Portfolio Composition of Sentiment-Classified Groups 31
Table 4 - Portfolio Composition Statistics 34
Table 5 - RQ1 BL Portfolio Performance Across Tau Values 36
Table 6 - RQ2 Portfolios Performance 39
Table 7 - Subperiod Analysis (RQ1) 42
Table A.1 - Methodological Comparison: This Study vs. Hung et al. (2024) and Colasanto et al. (2022) 51
Table B.1 - Weekly Cross-Sectional Return Distribution 52
Table B.2 - Summary Statistics of Sentiment Variables 52
Table B.3 - Most Persistently Classified Stocks as the 10% Cutoff 53
Table C.1 - Spread and View Vector Descriptive Statistics 55
Table D.1 - Covariance Matrix Diagnostics 56
Table D.2 - Portfolio Composition Statistics 57
Table D.3 - Average and Median Correlations Between Posterior Expected Returns (μBL) and Prior Equilibrium Returns (Π) 58
Table E.1 - RQ1 Full Portfolio Performance Across All Tau Values 59
Table E.2 - RQ2 Full Portfolio Performance Across All Tau Values 60
Table F.1 - Subperiod Analysis: RQ1 Portfolio Performance by Year, All τ Values 63
Table F.2 - Subperiod Analysis: RQ2 Portfolio Performance by Year, All τ Values 64
Table G.1 - Sharpe and Sortino ratios for Alternative View Construction 68
Table G.2 - View Vector Predictability Regressions 69
Table H.1 - Transaction Cost Sensitivity Analysis (Weekly Average Returns) 70

1. Introduction

Classical asset pricing models were developed under the belief that investors are fully rational and that markets are efficient. Under this traditional finance paradigm, asset prices are expected to reflect all available public and private information at all times, which is not compatible with systematic mispricing and predictable excess returns (Treynor, 1962; Sharpe, 1964; Lintner, 1965; Mossin, 1966).

Nevertheless, a growing number of empirical papers have challenged this view. Stock prices exhibit volatility that cannot be fully explained by changes in fundamentals (Shiller, 1981), and investors display systematic behavioral biases, including overconfidence, loss aversion, and representativeness, that cause prices to deviate from fundamental values in predictable ways (Kahneman and Tversky, 1979; De Bondt and Thaler, 1985; Barberis, Shleifer, and Vishny, 1998). In this context, Behavioral Finance emerges as a response to these limitations by incorporating insights from psychology into the decision-making process, giving rise to systematic patterns commonly referred to as market sentiment.

Market sentiment captures aggregated shifts in investors' beliefs that influence the demand and the valuation of financial assets beyond fundamentals and rational expectations (Baker & Wurgler, 2007). Empirical evidence suggests that sentiment has predictive effects on returns. For example, Baker and Wurgler (2006, 2007) were able to isolate the irrational component of sentiment and linked it to mispricing's. They found that periods of high sentiment are followed by low returns. During periods of high sentiment, stocks that are difficult to value become overpriced relative to fundamentals, and as sentiment reverts, prices adjust accordingly. Furthermore, sentiment effects appear asymmetric, meaning that negative sentiment may amplify losses more strongly than positive sentiment reduces them, consistent with loss aversion and negativity bias in investor behavior (Kahneman and Tversky, 1979; Baumeister et al., 2001).

If sentiment systematically affects asset prices, it follows that incorporating sentiment signals into portfolio optimization could improve investment outcomes. However, translating this theoretical insight into a practical and disciplined allocation framework is not straightforward. Direct use of sentiment signals in forecasting returns introduces the same estimation error problems that have been identified in classical mean-variance optimization (Markowitz, 1952; Michaud, 1989). Even well-specified return forecasts can produce extreme and

unstable portfolio weights when used as direct inputs to optimization (Black and Litterman, 1992). A more principled approach is offered by the Black-Litterman (BL) framework, which combines market equilibrium returns, implied by the CAPM through reverse optimization, with investor views, in a Bayesian updating structure (Black and Litterman, 1992; He and Litterman, 1999). By anchoring expected returns to market equilibrium and allowing views to tilt the portfolio only to the extent justified by the specified confidence, the framework provides a disciplined mechanism for incorporating sentiment signals while controlling estimation error and maintaining portfolio stability.

Despite the growing literature on sentiment-based investing, important gaps remain. Studies that document performance improvements from sentiment-informed portfolios typically rely on custom-built NLP pipelines or proprietary sentiment models that are difficult to replicate and compare across studies (Hung et al., 2024; Colasanto et al., 2022; Koratamaddi et al., 2021). As a result, it is unclear whether performance improvements reflect genuine informational content of sentiment or model-specific design choices.

This study addresses this gap by integrating standardized sentiment indicators provided by Bloomberg, Bloomberg News Sentiment and Bloomberg Twitter Sentiment, into a Black-Litterman portfolio optimization framework applied to 455 of the S&P 500 constituents over the period January 2024 to December 2025. To address RQ1, Bloomberg News and Twitter sentiment indicators are aggregated into a composite firm-level weekly sentiment score. Firms are classified into sentiment groups using percentile cutoffs of 10%, 20%, and 30%. For RQ2, two alternative classifications are considered. The first requires both Bloomberg News and Twitter indicators to simultaneously agree in sign, regardless of percentile rank. The second applies a sign constraint to the percentile-based classification, retaining only firms whose composite sentiment score is strictly positive or strictly negative. At each weekly rebalancing date, sentiment-based relative views are incorporated into the Black-Litterman framework alongside market equilibrium returns, producing posterior expected returns used for long-only portfolio optimization. A key methodological contribution is the treatment of the prior uncertainty parameter τ as a sensitivity parameter tested across five values, $\tau \in \{0.01, 0.05, 0.10, 0.20, 1.00\}$, rather than setting the parameter to a specific value or adopting the convention in which τ cancels in the posterior. A low tau means that an investor trusts the market equilibrium and therefore weighs more the market implicit expect returns relative to her views. The analysis focuses on two research questions:

RQ1. To what extent does incorporating Bloomberg sentiment indicators within a Black-Litterman portfolio optimization framework improve risk-adjusted portfolio performance relative to the S&P 500 benchmark?

RQ2. Do positive and negative sentiment signals have asymmetric effects on the performance of sentiment-informed portfolios?

This study contributes to the growing literature on sentiment-based portfolio construction by providing evidence on whether standardized, institutionally available sentiment indicators can replace the custom-built NLP models commonly used in prior research. By embedding these indicators within a theoretically grounded Black-Litterman framework, the study offers a transparent and replicable approach to sentiment-informed asset allocation that is directly applicable in practice. The analysis also contributes methodologically by combining portfolio performance evaluation with signal diagnostics, allowing a cleaner identification of whether limitations arise from the Q construction methodology or from the model implementation.

The results indicate that Bloomberg sentiment indicators can improve risk-adjusted portfolio performance within a Black-Litterman framework, but only when the prior uncertainty parameter τ is calibrated to small values that limit deviation from market equilibrium. At $\tau = 0.01$, the three cutoff specifications achieve Sharpe ratios close to the benchmark. Predictive regressions reveal that the investor view vector Q exhibits mean-reverting rather than momentum-like behavior at the weekly frequency, explaining why small tau values that limit the influence of the views systematically outperform larger ones. Positive sentiment portfolios uniformly outperform negative sentiment portfolios across all specifications on every risk-adjusted performance metric, a finding inconsistent with the behavioral finance prediction of asymmetric sentiment effects but consistent with the sustained bull market environment of the sample period. Subperiod analysis reveals, however, that this dominance narrows in 2025, providing partial support for the asymmetry hypothesis under market stress.

The remainder of this dissertation is structured as follows. Section 2 reviews the relevant literature on market sentiment and portfolio management theory. Section 3 describes the methodology. Section 4 presents the data and descriptive statistics. Section 5 reports the results and robustness analysis, and Section 6 concludes.

2. Literature Review

The literature review we present below sets the foundations for the framework and methodological approach we propose in the following chapters.

First, we review the empirical evidence on the impact of sentiment on equity pricing and analyze how market sentiment has been measured in previous studies.

Second, we briefly review the literature on portfolio optimization and asset allocation. The goal is to highlight how portfolio theory evolved from the traditional mean-variance approach to more recent and flexible frameworks such as the Black-Litterman model that we adopted in this research. In addition, we present some recent sentiment-based investment strategies and discuss their limitations.

Finally, the two streams of the literature are brought together by examining the existing literature integrating market sentiment within portfolio optimization: we review the methodologies adopted in recent research, the sentiment indicators that are more commonly used, and the results of those strategies.

2.1 Market Sentiment

Market sentiment is distinct from fundamentals, expectations or risk preferences. First, market sentiment differs from information. While new information can rationally justify price adjustments, sentiment refers to price movements that could be linked to news but are not justified by fundamentals alone (Baker & Wurgler, 2006). Second, sentiment is not a synonym for risk or uncertainty. Price changes can be associated with changing risk-premia, reflecting a change in macroeconomic conditions, while sentiment-driven changes arise from biased expectations, not objective shifts in economic risk (Aggarwal, 2022).

Under this framework, market sentiment can be understood as a time-varying, non-fundamental component of asset demand that systematically affects prices, volatility, and expected returns (Baker & Wurgler, 2006, 2007).

The relevance of market sentiment comes from its documented impact on equity prices, and volatility, particularly for assets that are more difficult to value or arbitrage. Baker

and Wurgler (2006, 2007) developed an important framework in which they show that sentiment has cross-sectional effects, in particular, small-cap, young, highly volatile, and non-dividend-paying stocks are more sensitive to sentiment. Given that these securities are typically harder to arbitrage, mispricing's driven by sentiment can last for longer periods. Consistent with this view, periods of high sentiment tend to be followed by lower subsequent returns, while periods of low sentiment are followed by higher future returns, indicating that sentiment contributes to predictable variation in returns (Baker & Wurgler, 2006, 2007).

While Baker and Wurgler (2006, 2007) identify which types of stocks are most affected by sentiment, Yu and Yuan (2011) examine how sentiment influences the overall magnitude of return anomalies. They show that a large number of return anomalies are significantly stronger following periods of high investor sentiment. This suggests that sentiment changes over time and can strengthen mispricing, especially when market friction prevents prices from quickly adjusting (Stambaugh et al., 2012). In this sense, sentiment can both affect specific groups of stocks and intensify anomaly profits during high-sentiment periods.

Beyond return predictability, market sentiment plays an important role in shaping market risk and volatility by reflecting changes in investors' risk appetite. Kumar and Persaud (2002) argue that declines in risk appetite can trigger contagion-like dynamics, where investors simultaneously reduce exposure to risky assets, leading to broad price declines that are not necessarily tied to fundamentals. Similarly, Bandopadhyaya and Jones (2006) describe how shifts in risk attitudes become part of the price-setting process when investor beliefs become more aligned. Complementing this risk-channel interpretation, empirical evidence based on professional news sentiment measures shows that sentiment retains explanatory power even after controlling for standard risk factors, and that the effect is particularly pronounced in the left tail of the return distribution (very negative returns) (Allen et al., 2019).

Recent contributions further explore the mechanisms through which sentiment-driven demand can distort prices. Extrapolative expectation models, such as Barberis et al. (2018), demonstrate how investor overreaction and belief extrapolation can generate bubbles and intensified trading activity. These theoretical insights complement empirical evidence showing that sentiment-driven mispricing can persist when arbitrage is limited.

Finally, the literature also notices that sentiment effects can be asymmetric and state dependent. Evidence from periods of financial distress suggests that negative sentiment can

amplify losses more strongly than positive sentiment can lessen them (Stambaugh et al., 2012; Allen et al., 2019). Taken together, these findings support the view that market sentiment influences not only expected returns (and therefore prices) but also volatility and the correlation structure of stock returns, aspects that are directly relevant for portfolio optimization.

Given that market sentiment is not directly observable, the literature identified several indirect proxies to measure it. These measures vary considerably in data source, frequency, construction methodology, and economic interpretation. Sentiment indicators are usually categorized into the following categories: survey-based measures, market-based proxies, and text-based indicators. Aggarwal (2022) highlights that the lack of a unified definition of sentiment has contributed to heterogeneity in measurement approaches, which raises concerns regarding comparability and empirical robustness across studies.

Survey-based measures aim to capture investors' beliefs directly. Fisher and Statman (2000) use survey data to illustrate how cognitive biases, such as overconfidence and the "illusion of validity," translate into systematic optimism or pessimism. Brown and Cliff (2005) show that survey-based sentiment is significantly related to aggregate market valuation, with excessive optimism preceding lower subsequent returns as prices revert toward fundamentals. The primary advantage of survey measures comes from their conceptual clarity, as they directly reflect investor expectations and beliefs (Kresta, et al., 2024). However, survey-based measures are typically available at relatively low frequencies and provide aggregate indicators of investor optimism or pessimism (Sun et al., 2016). In addition, because survey data relies on self-reported expectations, they may be affected by reporting and sampling biases (Rosenman et al., 2011).

Market-based proxies infer sentiment indirectly from observable trading behavior and financial variables. Baker and Wurgler (2006, 2007) construct a widely used composite sentiment index by aggregating six indicators, including the closed-end fund discount, NYSE turnover, IPO activity, equity issuance share, and dividend premium. Other commonly employed proxies include the VIX "fear index" and the put-call ratio (Whaley, 2000; Dennis & Mayhew, 2002). Closely related are risk-appetite indices, such as the Equity Market Sentiment Index (EMSI) developed by Kumar and Persaud (2002) and Bandopadhyaya and Jones (2006), which interpret sentiment as a time-varying willingness to bear risk. While market-based proxies benefit from high frequency and broad market coverage, they often mix sentiment with liquidity conditions, macroeconomic risk, or volatility expectations, complicating

interpretation. As noted in Aggarwal (2022), such proxies may capture multiple latent factors rather than a pure measure of sentiment.

Text-based sentiment measures have become more popular with advances in Natural Language Processing (NLP) and the increased availability of textual data from news and social media platforms. Tetlock (2007) was one of the first researchers exploring this option, in his study, he quantified media pessimism in a Wall Street Journal column and showed that negative tones predict short-term price pressure that is followed by a reversal. Loughran and McDonald (2011) highlight important methodological concerns, pointing out that general-purpose dictionaries frequently misclassify financial terminology and propose the use of finance-specific word lists to improve classification accuracy. More recent machine-learning approaches, such as the supervised learning framework introduced by Ke et al. (2023), extract return-predictive information from news and are often able to outperform simpler dictionary-based methods.

Social media platforms provide an additional source of textual sentiment data. Bollen et al. (2011) show that mood states derived from Twitter feeds, particularly the "Calm" dimension, can forecast short-term movements in the Dow Jones Industrial Average. Similarly, Cookson and Niessner (2020) use StockTwits data to measure investor disagreement and document its relationship with return dynamics. While Text-based measures offer important advantages, including more frequent updates, information at the individual company level, and the ability to process large volumes of data efficiently, their construction depends heavily on methodological choices. As Aggarwal (2022) notes, differences in data sources, classification techniques, and model design contribute to variation in sentiment measures across studies. As a result, text-based sentiment indicators may lack full comparability when different preprocessing methods or algorithms are employed.

Overall, although sentiment has been extensively documented and measured through diverse methodologies, the heterogeneity of proxies, differences in data handling, and absence of standardized processes present significant challenges for comparability and practical implementation. As highlighted by Aggarwal (2022), weak correlations across sentiment indicators suggest that different measures may capture distinct behavioral components. For portfolio optimization purposes, this raises an important consideration whether the economic value of sentiment signals depends not only on their predictive power but also on their stability, transparency, and replicability. Addressing these measurement challenges is

therefore essential when integrating sentiment into systematic asset allocation frameworks.

2.2 Evolution of Portfolio Theory

Portfolio theory addresses problems of capital allocation across risky assets under uncertainty. Early portfolio optimization frameworks were grounded in the mean-variance theory of Markowitz, which was later integrated into equilibrium asset pricing models such as the CAPM. These were developed under the assumption of rational investors and efficient markets, which derived the optimal allocations from estimations of expected returns and risk. Within this framework, portfolio decisions relied on personal preferences and the estimation of parameters (Markowitz, 1952; Sharpe, 1964; Lintner, 1965; Mossin, 1966)[1].

Over time, however, important limitations related to the sensitivity of errors in estimations and to the questionable assumptions of the models were identified, leading to the development of new models and frameworks (Michaud, 1989; Britten-Jones, 1999). Within this evolutionary process, the Black-Litterman model emerged as a structured response to the instability of mean-variance optimization while preserving its core theoretical insights (Black & Litterman, 1992).

2.2.1 Mean-Variance Optimization

The modern portfolio theory originated with Markowitz's (1952) mean-variance framework. The model establishes that investors face a trade-off between expected return and risk, measured by portfolio variance. In this setting, investors are assumed to be rational and risk-averse, choosing portfolios that minimize variance for a certain level of expected return. In this setting, an important aspect is diversification, because stock returns are not perfectly correlated, combining them can reduce the overall variance of the portfolio without reducing the expected return. This optimization process identifies an efficient frontier of portfolios that offer minimum variances for different levels of expected return (Markowitz, 1952).

[1] Under some assumptions, the optimal risky portfolio is the same for all investors regardless of preferences. CAPM predicts further that the optimal risky portfolio is the market portfolio.

Subsequent contributions extended this framework, incorporating risk-free assets and market equilibrium. For instance, Tobin (1958) introduced the separation theorem, showing that all investors will hold the same risk-free and risky assets in different proportions. Later on, Treynor (1962), Sharpe (1964), Lintner (1965), and Mossin (1966) developed the CAPM, establishing a linear relationship between systematic risk and expected return. Within this model, because idiosyncratic risk can be eliminated through diversification, investors are only compensated for taking higher levels of systematic risk measured by the portfolio beta.

Although useful, the mean-variance framework relies on several strong assumptions regarding investor behavior and market structure. Investors are assumed to be rational and risk-averse, notions challenged by behavioral finance (Kahneman & Tversky, 1979). This representation is exact under either quadratic utility or normally distributed returns, conditions under which the mean and variance fully characterize investor preferences and the return distribution. However, empirical evidence suggests that asset returns exhibit skewness and excess kurtosis, and that higher-order moments may affect expected returns (Kraus & Litzenberger, 1976; Harvey & Siddique, 2000). Subsequent extensions, like CAPM, assume homogeneous expectations and frictionless markets, ideas that have received both theoretical and empirical criticisms (Roll, 1977; Fama & French, 1992). While the assumptions of the framework enable analytical tractability, they limit the descriptive realism of the model in practical settings.

Moreover, risk is defined as the variance of returns, meaning that both positive and negative deviations are considered in the same way (Markowitz, 1952). However, this definition may not fully reflect investor preferences. Behavioral evidence suggests that investors are more affected by negative outcomes than what they would be by positive outcomes on the same magnitude (Kahneman & Tversky, 1979). From this perspective, treating upside volatility as undesirable may be economically inconsistent with actual risk perceptions.

Other limitations include the high sensitivity of mean-variance optimization framework to estimation errors. Sampling error in estimated expected returns and covariances translates into large variation in estimated allocation weights (Britten-Jones, 1999). Michaud (1989) argues that mean-variance optimization can "error maximize," producing extreme and unintuitive allocations when inputs are noisy. This instability is not only a theoretical concern, DeMiguel, Garlappi, and Uppal (2009) show that, out-of-sample, optimized mean-variance

strategies often fail to outperform simple diversification strategies such as equal-weighting (1/N), highlighting the practical consequences of parameter uncertainty.

2.2.2 Alternative Approaches to Classical Mean-Variance Optimization

Given the weaknesses found in mean-variance optimization, subsequent research focused on developing alternatives that addressed its four main limitations: estimation error and parameter stability, the measurement of risk, incomplete return specification, and its static nature.

One of the main focus was to reduce the reliance on expected return estimates. For instance, the Global Minimum-Variance Portfolio (GMVP), which is a direct implication of the original mean-variance framework (Markowitz, 1952), gained renewed attention in later literature due to its practical advantages. Because the GMVP calculates asset weights based on the covariance matrix of returns, it produces more stable estimations given that it does not rely on expected return estimations (Michaud, 1989; Britten-Jones, 1999). Additionally, Clarke, de Silva, & Thorley (2011) document that minimum-volatility strategies can exhibit competitive performances in practice. Similar frameworks emerged, such as Risk Parity portfolios (see, for example, Maillard, Roncalli, and Teiletche, 2010). These translate into creating a minimum-variance portfolio subject to the constraint that each asset contributes equally to the portfolio risk. Models like the Most Diversified Portfolio (MDP) try to maximize the diversification ratio while balancing the weights of each asset (Choueifaty & Coignard, 2008). Yet empirical studies show that optimized portfolios often fail to outperform simple equal-weighted strategies (DeMiguel, Garlappi, & Uppal, 2009). This motivated the development of more robust optimization techniques that explicitly account for parameter uncertainty.

A second line of research focused on the definition of risk. While the classical framework treats downsides and upsides symmetrically (Markowitz, 1952), alternative measures were introduced to modify the objective function of optimization. Konno and Yamazaki (1991) propose a mean-absolute deviation model as a linear alternative to variance-based optimization. Downside-focused measures such as semi variance (Markowitz, 1952) restrict attention to negative deviations from a target return. In addition, coherent tail-risk measures, formalized by Artzner et al. (1999) and operationalized in portfolio selection by Rockafellar and Uryasev (2000), incorporate Conditional Value-at-Risk (CVaR) into the optimization

framework. Although these approaches modify the risk metric within the optimization problem, they preserve the mean-risk structure of portfolio selection and therefore continue to require expected return inputs.

A third strand of literature challenges the single-factor structure of the CAPM. Empirical evidence suggests that the market beta alone is not fully capable of explaining cross-sectional variation in expected returns (Fama & French, 1992). Multifactor models, such as the Fama-French three-factor model (Fama & French, 1993) and the Arbitrage Price Theory (Ross, 1976), introduce additional determinants of expected returns. Although these models enhance the economic interpretation of risk premia by providing an alternative structure for estimating expected returns (Fama & French, 1993; Ross, 1976), they do not address the sensitivity of portfolio weights to estimation errors.

Finally, dynamic extensions emphasize that portfolio choice is inherently intertemporal. Merton (1969) develops an intertemporal asset pricing framework in which investors account for changes in investment opportunities and hedge against shifts in state variables. While dynamic models are more realistic, they also increase model complexity and parameter requirements (Campbell & Viceira, 2002), often increasing estimation challenges.

Taken together, these alternative models address important weaknesses of the classical framework. Nonetheless, they either reduce reliance on return forecasts, substitute them with alternative parametric models, or increase model complexity without creating an integrated solution that allows the incorporation of additional information into portfolio allocation while controlling estimation errors. This tension between incorporating forward-looking information and controlling parameter instability motivates the adoption of Bayesian portfolio optimization methods. In particular, the Black-Litterman framework combines equilibrium returns with subjective views in a statistically disciplined manner, thereby reducing sensitivity to extreme inputs (Black & Litterman, 1992).

2.2.3 Bayesian Portfolio Optimization and the Black-Litterman Framework

Bayesian portfolio optimization offers a response to the expected return estimation problem by treating expected returns as random variables and combining prior information with new signals in a probabilistic framework (Klein & Bawa, 1976).

The Black-Litterman model (Black & Litterman, 1992) represents the most influential implementation of Bayesian reasoning in asset allocation. Instead of estimating expected returns based on historical averages, the model starts from the equilibrium returns implied by market capitalization weights through reverse optimization under the CAPM structure. These implied returns serve as a prior distribution reflecting the market consensus. Investor views are then incorporated through a linear mapping matrix and combined with the prior using Bayesian updating to generate posterior expected returns (Black & Litterman, 1992). The posterior returns replace the original expected return vector.

The key innovation of the Black-Litterman framework comes from its shrinkage mechanism. Expected returns derived from investor views are not used directly, instead, they are blended with equilibrium returns according to the specified confidence in each view. This anchoring in market equilibrium reduces the extreme and unintuitive portfolio weights often produced by classical mean-variance optimization (Black & Litterman, 1992; Idzorek, 2005). In contrast to global minimum-variance or minimum-volatility approaches, which improve stability by avoiding expected returns, the Black-Litterman model retains forward-looking information while controlling estimation error. Similarly, unlike multifactor asset pricing models that provide alternative expected return estimates, Black-Litterman modifies the optimization process itself by embedding return uncertainty within a Bayesian updating structure (Satchell and Scowcroft, 2000).

Empirical research further supports the practical relevance of the framework. Bessler, Opfer, and Wolff (2017) show that Black-Litterman portfolios can outperform traditional mean-variance and naive diversification strategies, particularly when return forecasts are unstable. Giacometti et al. (2007) extend the model to non-normal return distributions and demonstrate its robustness under alternative distributional assumptions.

Despite its advantages, the Black-Litterman model has its limitations. Its implementation requires choosing two key inputs: τ, which reflects the uncertainty about the market equilibrium returns (and therefore the CAPM model itself), and Ω, which captures the confidence in investor views, both of which can influence posterior results, and the literature has not come up with a unique specification for these (Satchell & Scowcroft, 2000; Meucci, 2008). Furthermore, because the model ultimately relies on a mean-variance optimization structure, the framework remains sensitive to the estimation of the covariance matrix.

The flexibility of the Black-Litterman framework allows the incorporation of systematic signals beyond traditional financial variables. In this dissertation, we use sentiment indicators to form relative return views and embed these within the Bayesian updating structure. This enables forward-looking information to influence portfolio weights while remaining anchored to market equilibrium.

2.3 Market Sentiment-based Investment Strategies

This section integrates the previous discussions on market sentiment and portfolio optimization theory in order to examine how sentiment signals can be incorporated into asset allocation decisions.

Recent research increasingly investigates the role of sentiment signals in asset allocation. Across a variety of methodological frameworks, the empirical evidence suggests that incorporating sentiment in asset allocation strategies can improve portfolio performance when compared to price-based approaches. Studies using machine learning architectures including long short-term memory networks, reinforcement learning, and transformer-based models consistently find that sentiment-enhanced strategies generate higher returns and improved risk-adjusted performance (Malandri et al., 2018; Koratamaddi et al., 2021; Mantshimuli, 2025; Konstantinidis et al., 2023; Feng et al., 2024). Similar improvements have been documented within more structured portfolio optimization settings, including Black-Litterman frameworks, where sentiment signals are used to generate portfolio allocation weights (Hung et al., 2024; Colasanto et al., 2022). The integration of sentiment has also been explored within alternative risk-focused models, yielding improved performance (Muthivhi & van Zyl, 2021; Taheripour et al., 2025). Taken together, these findings suggest that sentiment contains economically meaningful information that can enhance portfolio allocation decisions.

Within the Black-Litterman literature specifically, two studies are most directly relevant to this dissertation. Hung et al. (2024) construct sentiment-informed portfolio views by fine-tuning a BERT classifier on news articles scraped across seven S&P 500 stocks. The binary directional output of the classifier, predicting whether a stock will rise or fall the following day, is passed as a feature into three recurrent neural network architectures (RNN, LSTM, and GRU) alongside historical price and volume data, and the resulting predicted

returns from the view vector Q. Over a 224-day test period in 2019, the GRU-based Black-Litterman portfolio achieves an annualized return of 46.6% and Sharpe and Sortino ratios of 13.0% and 17.9%, outperforming the S&P 500, an equal-weight portfolio, and the Markowitz mean-variance strategy. However, this performance comes at a significant practical cost: portfolio turnover reaches 68% for the GRU specification and 84% for the RNN specification, levels that would generate substantial transaction costs and that the authors acknowledge as a key limitation. Critically, because the sentiment signal functions as an auxiliary input to a price prediction pipeline rather than directly as the view, the independent contribution of sentiment to performance is difficult to isolate from the predictive power of the deep learning architecture itself. Colasanto et al. (2022) adopt a more direct approach, applying FinBERT, a domain-specific BERT variant pre-trained on financial text, to Financial Times articles to extract a continuous polarity score for each stock. This score is then used within a Monte Carlo simulation framework, formally grounded in the Wiener measure, to generate a predicted future price that determines the absolute view Q in the Black-Litterman model. Over a single rebalancing window in early 2021, the sentiment-informed portfolio achieves a Sharpe ratio of 1.14 compared to 1.07 for the no-views portfolio. While this approach establishes a more direct link between the raw sentiment signal and the portfolio view, several limitations constrain its generalizability. The evaluation covers only a single five-day rebalancing period, and confidence levels in the uncertainty matrix Ω are set manually by the researchers, reintroducing the subjective judgment that the Black-Litterman framework was originally designed to reduce.

Despite the promising empirical results, the literature exhibits substantial heterogeneity in how sentiment signals are constructed and implemented, ranging from custom-built NLP pipelines and task-specific classification models (Kucukaslan & Tas, 2024; Konstantinidis et al., 2023; Koratamaddi et al., 2021; Mantshimuli, 2025) to proprietary scoring systems and domain-specific pre-trained models such as FinBERT (Feng et al., 2024; Taheripour et al., 2025; Colasanto et al., 2022). Because sentiment scores are highly sensitive to design and preprocessing choices, such as tokenization, lemmatization, and negation handling, minor differences across studies can lead to substantially different signals (Todd et al., 2024), making it difficult to determine whether documented performance improvements reflect the genuine informational content of sentiment or model-specific tuning (Jiang, 2020). Furthermore, the reliance on complex pipelines often produces unintuitive results that are hard to

interpret or replicate, limiting the practical applicability of these approaches for institutional investors (Rudin, 2019; Consoli et al., 2022).

Building on the literature discussed, this study examines how standardized sentiment indicators can be incorporated into equity portfolio optimization and whether doing so improves portfolio performance relative to a passive investment benchmark. To this end, Bloomberg sentiment indicators are integrated into a Black-Litterman optimization framework, allowing sentiment signals to be incorporated as investor views within a disciplined equilibrium-based allocation model. Unlike Hung et al. (2024) and Colasanto et al. (2022), which rely on custom-scraped data and NLP pipelines applied to narrow stock universes, this study uses institutionally available, consistently constructed sentiment measures applied to S&P 500 constituents over 2024-2025, enabling a more transparent and practically reproducible evaluation. Table A.1 (see Appendix) summarizes the key methodological differences between the two studies. In addition, the study also investigates whether the informational content of sentiment signals differs depending on their direction. Behavioral finance research suggests that financial markets may react asymmetrically to information, as investors often respond more strongly to negative news than to positive news due to mechanisms such as loss aversion and negativity bias (Kahneman & Tversky, 1979; Chopra & Ziemba, 1993). As a result, sentiment signals may not contribute symmetrically to portfolio construction, and negative sentiment may contain stronger information regarding future price movements than positive sentiment, a question that neither Hung et al. (2024) nor Colasanto et al. (2022) address directly.

Our analysis addresses the following research questions:

RQ1. To what extent does incorporating Bloomberg sentiment indicators within a Black-Litterman portfolio optimization framework improve risk-adjusted portfolio performance relative to the S&P 500 benchmark?

RQ2. Do positive and negative sentiment signals have asymmetric effects on the performance of sentiment-informed portfolios?

3. Methodology

In this chapter, we present the empirical methodology used to address the research questions outlined in the previous chapter. In particular, the analysis investigates whether incorporating market sentiment into portfolio optimization improves risk-adjusted portfolio performance and whether positive and negative sentiment signals exhibit asymmetric effects. To do so, we obtained the sentiment indicators from Bloomberg. These are the basis of our views in the Black-Litterman Bayesian framework. This approach provides a structured way to integrate sentiment information into expected returns while maintaining consistency with the underlying risk structure conveyed by the market portfolio. The following sections describe how sentiment signals are constructed and how they are incorporated into the Black-Litterman model. Finally, we present the measures we use to evaluate portfolio performance.

3.1 Construction of Sentiment Signals

Sentiment signals are constructed using firm-level weekly composite sentiment scores obtained by combining Bloomberg News and Twitter sentiment indicators into a single measure. These scores are used to capture market sentiment at each rebalancing date.

For the baseline specification, firms are ranked cross-sectionally according to their composite sentiment scores. Based on this ranking, firms are classified into sentiment-based groups using percentile cutoffs of 10%, 20%, and 30% of the distribution. These groupings are used to construct relative investor views within the Black-Litterman framework, reflecting the expectation that firms with higher sentiment will outperform firms with lower sentiment (Tetlock, 2007). This approach is consistent with the idea that sentiment signals are more informative for ranking assets cross-sectionally than for predicting absolute return levels (Baker and Wurgler, 2006).

To address the second research question, two alternative sentiment classification approaches are considered in order to isolate the effects of positive and negative sentiment signals.

In the first approach, firms are included in positive or negative sentiment portfolios only when both Bloomberg News and Twitter sentiment indicators are aligned in sign. Specifically, a firm is classified as positive (negative) only if both indicators are simultaneously

positive (negative). This approach focuses on stronger and more consistent sentiment signals, at the cost of a reduced cross-sectional sample.

In the second approach, the percentile-based classification used in the baseline specification is retained, but an additional constraint is imposed on the sign of sentiment. Firms are selected from the upper (lower) percentile cutoffs only if their composite sentiment score is strictly positive (negative). This ensures consistency between the strength and direction of sentiment while maintaining a broader set of assets.

These alternative constructions allow the analysis to evaluate whether the impact of sentiment on portfolio performance depends on the strength, consistency, and direction of the signal.

3.2 Black-Litterman Portfolio Construction Framework[2]

A growing number of studies in portfolio optimization adopt the Black-Litterman framework as a structured approach for incorporating external information into asset allocation decisions while maintaining consistency with market equilibrium (Black & Litterman, 1992; He & Litterman, 1999; Satchell & Scowcroft, 2000; Idzorek, 2005).

In this study, the Black-Litterman framework is used to incorporate sentiment-based signals into the portfolio construction process. The model is particularly suitable in this context because market sentiment represents noisy and rapidly changing information, which can influence expected returns but may also introduce instability if used directly in return forecasting (Tetlock, 2007).

The model consists of three main components: market equilibrium returns (prior expected returns), investor views derived from sentiment signals, and the posterior expected returns used for portfolio construction.

[2] Black & Litterman, 1992; Satchell & Scowcroft, 2000; Idzorek, 2005; Kolm et al., 2021

3.2.1 Market Equilibrium (prior expected returns)

The starting point of the Black-Litterman framework is the expected excess returns implied by the market prices in equilibrium. Rather than directly forecasting expected returns, the model infers the returns that are consistent with the observed market portfolio (The optimal risky portfolio under CAPM) and its level of risk. These implied returns provide a neutral benchmark for expected returns, which can subsequently be adjusted when incorporating additional information such as sentiment signals.

The implied equilibrium excess returns (Π) depend on three variables: the market capitalization weights (w_m), the covariance matrix of asset returns (Σ), and the risk-aversion parameter (δ).

Market-implied equilibrium excess returns are calculated as:

$$\Pi = \delta \Sigma w_m \tag{1}$$

The covariance matrix is estimated using the historical weekly excess returns of the 455 constituents over a rolling window of 52 weeks, allowing the risk structure of the portfolio to evolve over time while remaining consistent with the rebalancing frequency. Given its high dimensionality of the estimation problem (the number of assets is greater than the number of observations), an analytical nonlinear shrinkage method is adopted to ensure a numerically stable and inventible matrix, following the approach of Ledoit and Wolf (2020).

The risk aversion parameter (δ) is computed using an implied market-based approach as:

$$\delta = \frac{E(R_m - Rf)}{\sigma_m^2} \tag{2}$$

where R_m represents the return of the S&P 500 index and R_f the risk-free rate. The parameter is estimated from market data and treated as constant over the sample period, following He and Litterman (1999) and Grinold & Kahn (1999), who suggest that the implied risk aversion coefficient can be approximated as the ratio of the expected excess return of the market portfolio to its variance.

3.2.2 Investor views (sentiment-based signals)

Investor views represent the component of the Black-Litterman framework that allows portfolio allocations to deviate from market equilibrium based on additional information. In this study, the views are derived from firm-level sentiment indicators and are formulated in relative terms, expressing beliefs about the performance of groups of stocks rather than individual assets.

However, specifying relative views introduces additional methodological complexity. As noted by Barua and Sharma (2023), the Black-Litterman literature has predominantly focused on absolute return views, with relatively limited guidance on the implementation of relative views. As a result, the construction of sentiment-based views in this study requires additional modeling choices regarding the definition of groups, the specification of the pick matrix, and the calibration of view uncertainty.

In the model, views are defined through three elements. First, μ denotes the vector of the true but unobservable expected returns of all assets. Second, P is the pick matrix, which identifies the assets involved in each view. Finally, Q represents the expected return differential associated with each view.

Investor views are encoded through the relationship:

$$P\mu = Q \tag{3}$$

At each rebalancing date, stocks are ranked according to their sentiment scores and classified into high-sentiment and low-sentiment groups. A relative view is then formulated, stating that firms in the high-sentiment group are expected to outperform those in the low-sentiment group.

The pick matrix (P) is constructed such that stocks in the high-sentiment group receive positive weights and stocks in the low-sentiment group receive negative weights. These weights are proportional to the market capitalization of each stock within their respective groups and are normalized so that the weights in the high-sentiment group sum to +1 (100%) and those in the low-sentiment group sum to -1, following Idzorek (2005). When addressing the second research question, views are also specified as relative views. For the positive sentiment portfolio, the pick matrix assigns positive weights to firms in the positive sentiment

group and negative weights to all remaining firms, normalized so that the long leg sums to +1 and the short leg sums to -1. For the negative sentiment portfolio, the pick matrix assigns positive weights to all remaining firms and negative weights to firms in the negative sentiment group, with the same normalization. In both cases, weights within each leg are proportional to market capitalization. This formulation expresses the view that positive-sentiment firms are expected to outperform the rest of the market, and that the market is expected to outperform negative-sentiment firms, respectively.

The view vector Q is constructed using historical return differentials between the long and short legs of each portfolio, providing a transparent and data-driven approach consistent with common empirical practices (Kara et al., 2019). For RQ1, this differential is the realized weekly return of the high-sentiment group minus the return of the low-sentiment group. For RQ2, the differential depends on the direction of the view, for positive sentiment portfolios, Q is constructed from the realized return of the positive-sentiment classified group minus the return of the remaining market universe, while for negative sentiment portfolios, Q is constructed from the realized return of the remaining market universe minus the return of the negative-sentiment classified group. In both cases, a positive Q value indicates that the long leg outperformed the short leg in recent periods, which is used as the basis for the forward-looking view. Full details of the construction procedure, including the exponential weighting scheme, are provided in Section 3.4.

3.2.3 Posterior expected return and portfolio construction

The posterior expected returns are obtained by combining market-implied equilibrium returns with sentiment-based investor views. These represent the updated expected returns used for portfolio construction after incorporating both sources of information.

The posterior expected returns depend on the following components: the market-implied equilibrium expected returns ($\Pi \in \mathbb{R}^{N\times 1}$); the covariance matrix ($\Sigma \in \mathbb{R}^{N\times N}$); the pick matrix ($P \in \mathbb{R}^{K\times N}$); the view vector ($Q \in \mathbb{R}^{K\times 1}$), where *K* denotes the number of views; the uncertainty matrix ($\Omega \in \mathbb{R}^{K\times K}$), a scalar parameter ($\tau$), which scales the uncertainty associated with the market-implied equilibrium returns.

Posterior expected returns are obtained through Bayesian updating and are given by:

$$\mu^{BL} = [(\tau\Sigma)^{-1} + P'\Omega^{-1}P]^{-1}[(\tau\Sigma)^{-1}\Pi + P'\Omega^{-1}Q] \quad (4)$$

The uncertainty matrix Ω is specified as:

$$\Omega = P\Sigma P' \quad (5)$$

Where P is the pick matrix identifying the assets involved in each view and Σ is the covariance matrix of excess returns estimated from historical data. This anchors the uncertainty of each view to the historical variance of the long-short portfolio associated with that view, providing a transparent and data-driven specification.

This specification differs from the proportional approach of He and Litterman (1999), in which $\Omega = \tau P\Sigma P'$ causes τ to enter both the prior uncertainty term ($\tau\Sigma$) and the view uncertainty matrix simultaneously. Under that formulation, varying τ rescales both terms in tandem, making it impossible to isolate the effect of prior confidence on the posterior returns. By specifying $\Omega = P\Sigma P'$ independently of τ, the scalar operates exclusively on the prior distribution, so that its role is to control the weight placed on the equilibrium prior relative to the views. Higher values of τ place greater uncertainty on the prior and pull $E[R]$ toward the views, while lower values tighten the prior and pull $E[R]$ toward Π.

Given the absence of consensus on τ in the literature, with recommended values ranging from 0.01 to 0.05 calibrated against tracking error (Idzorek, 2005), to approximately $1/T$ based on interpreting $\tau\Sigma$ as the standard error of the estimate of Π (Blamont and Firoozye, 2003), to 1 (Satchell and Scowcroft, 2000), τ is treated as a free parameter and results are reported across its plausible range.

3.3 Portfolio Construction

The posterior expected returns obtained are used as inputs for portfolio construction. Portfolio weights are obtained through a long-only optimization that maximizes the portfolio Sharpe ratio using the posterior expected returns and covariance matrix. This constraint is imposed to reflect realistic investment constraints faced by many institutional investors and to ensure comparability with the benchmark (Molk & Partnoy, 2019).

Formally, the optimization problem is defined as:

$$\max_{w} \frac{w'\mu^{BL}}{\sqrt{w'\Sigma w}} \tag{6}$$

Subject to:

- $w_i \geq 0$ for all i (long-only constraint)
- $\sum w_i = 1$ (full investment constraint)

where w denotes the vector of portfolio weights, μ^{BL} represents the vector of posterior expected excess returns obtained from the BL framework, and Σ is the covariance matrix.

Three portfolio specifications are considered. The first incorporates the full sentiment signal and represents the baseline sentiment-informed allocation. Its performance is evaluated relative to the S&P 500 index to address RQ1. To address RQ2, two additional portfolio specifications are constructed using only positive sentiment signals and only negative sentiment signals, respectively. This separation allows the analysis to assess whether negative sentiment generates stronger portfolio effects than positive sentiment, consistent with the behavioral finance literature on asymmetric investor responses to information.

3.4 Construction of Model Inputs

The inputs required for the Black-Litterman model are constructed at each weekly rebalancing date, using only information available prior to that date to avoid look-ahead bias.

Market-capitalization weights are computed by normalizing firm-level market capitalizations so that the weights sum to one across all assets. These weights represent the market portfolio and are used to derive the implied equilibrium excess returns via equation (1).

Sentiment-based signals are constructed using the cross-sectional ranking of firms according to their weekly composite sentiment scores. At each rebalancing date, firms are sorted and classified into high-sentiment and low sentiment groups based on percentile cutoff thresholds of 10%, 20%, and 30%. These alternative cutoffs allow the analysis to assess the robustness of the results to different definitions of sentiment strength.

The pick matrix ($P \in \mathbb{R}^{1\times 455}$) is constructed as described in section 3.2.2. To address RQ1, stocks are selected based on their groups (high or low sentiment) that were created

based on the combined sentiment score. Stocks in the high sentiment group have a combined weight, in the vector, of +1 and stocks in the low sentiment group have a combined weight of -1, resulting in a combined weighting of 0 for P, translating a relative view setting. To address RQ2, stocks are selected based on directional sentiment or sentiment agreement variables. The same logic applies to both RQ2 classification approaches, with the only difference being how firms are assigned to the positive or negative sentiment group, as described in Section 3.1. For the positive sentiment portfolio, the pick matrix assigns positive weights to firms in the positive sentiment group and negative weights to all remaining firms, normalized so that the sum of positive weights equals +1 and the sum of negative weights equals -1, expressing a relative view that positive-sentiment firms outperform the rest of the market. For the negative sentiment portfolio, the pick matrix assigns positive weights to all remaining firms and negative weights to firms in the negative sentiment group, with the same normalization, expressing a relative view that the market outperforms negative-sentiment firms. In both cases, weights within each leg are proportional to market capitalization.

The view vector ($Q \in \mathbb{R}^{1\times 1}$) is time-varying. It is defined as the expected return differential between high and low sentiment portfolios, estimated using a 52-week rolling window of realized spreads with exponential decay weighting. Let s_{t-k} denote the realized return spread between high- and low-sentiment portfolios k weeks prior to rebalancing date t. The expected spread used to construct the view vector is defined as:

$$\hat{q}_t = \sum_{k=1}^{52} w_k \; s_{t-k} \tag{7}$$

where the weights w_k follow an exponential decay structure:

$$w_k = \frac{\lambda^{k-1}}{\sum_{j=1}^{52} \lambda^{j-1}}, 0 < \lambda < 1 \tag{8}$$

This weighting scheme assigns greater importance to more recent observations while still incorporating information from the full historical window. Prior research suggests that the impact of sentiment on stock prices is short-lived and tends to revert relatively quickly to fundamentals (Tetlock, 2007), implying that recent information should play a more prominent role in the estimation of expected return differentials. At the same time, relying exclusively on a short window may introduce substantial noise. The use of an exponentially weighted average over a longer horizon provides a balance between responsiveness and stability in the estimation of the view vector. The decay parameter λ is set to 0.95 throughout.

For the second research question, the same exponential weighting procedure is applied using the realized return differential of the corresponding relative portfolio, the return of the long leg minus the return of the short leg, ensuring consistency between the pick matrix and the view interpretation.

The uncertainty associated with each view is computed as the variance of its corresponding long-short portfolio ($p_k' \Sigma\, p_k$), as described in equation (5). This specification ensures that views with higher return variability are assigned lower confidence, as the uncertainty of each view scales directly with the variance of its corresponding long-short portfolio. As discussed in Section 3.2.3, Ω is specified independently of τ, tau operates on the weight of the prior implicit returns and is associated with the confidence of the market equilibrium model. A lower tau means more confidence in the expected returns implied in prior market capitalizations. To assess the sensitivity of the posterior to this parameter, results are reported for $\tau \in \{0.01, 0.05, 0.10, 0.20, 1.00\}$, spanning the full range of values proposed in the literature from near-zero (Idzorek, 2005) to unity (Satchell and Scowcroft, 2000).

3.5 Performance Evaluation

Portfolio performance is evaluated using seven standard metrics. Average return and volatility capture mean profitability and total risk. The Sharpe and Sortino ratios, which allow the comparison between portfolios with different levels of risk (Sharpe, 1966; Sortino and Price, 1994). Maximum drawdown captures the largest peak-to-trough loss over the sample period. The information ratio measures annualized active return per unit of tracking error relative to the benchmark (Grinold and Kahn, 1999). Turnover is computed as the sum of absolute weight changes relative to drift-adjusted end-of-period weights, divided by two to reflect one-way turnover (Brandt, 2010). Given that τ controls the relative weight placed on prior equilibrium versus sentiment-based views, all metrics are reported across five tau specifications to assess parameter sensitivity. To address RQ1, the baseline portfolio is evaluated relative to the S&P 500. To address RQ2, portfolios constructed using only positive sentiment signals are compared with those using negative sentiment signals, assessing whether negative sentiment signals are more influential more, consistent with the behavioral finance findings on asymmetric investor reactions to adverse information (Da et al., 2011).

4. Data

This chapter describes the data used in the empirical analysis and the steps taken to construct the variables required for portfolio construction. It outlines the data collection process, the definition of the investment universe, the cleaning procedures applied to the raw data, and the construction of the final dataset.

4.1 Data Collection and Sample Period

The empirical analysis focuses on US stocks, specifically firms included in the S&P 500 index. The US stock market is selected due to its size, liquidity, and extensive coverage by professional news outlets and social media platforms, which is particularly relevant for sentiment-based analysis.

While the S&P 500 serves as the benchmark, the investment universe is restricted to a subset of its constituents. The sample includes firms that remained constituents of the index throughout the entire sample period and for which sufficient financial and sentiment data are available. All data is obtained from Bloomberg. The dataset includes Bloomberg News and Twitter sentiment indicators, daily stock close adjusted prices, market capitalization data, S&P 500 index returns, and one-month U.S. Treasury bill yields. The primary analysis covers the period from January 2024 to December 2025. To ensure that all model inputs are available from the first rebalancing date, additional data from January 2023 is included as a pre-sample warm-up period. This pre-sample window supports the construction of the 52-week rolling covariance matrix and the exponentially weighted view vector Q, both of which require a full year of prior observations before the first portfolio is formed.

4.2 Investment Universe and Filtering Criteria

To ensure data consistency across variables, a filtering procedure is applied based on data availability. Firms are excluded if more than 5% of weekly observations are missing in the variables required for portfolio construction.

For sentiment data, Bloomberg News and Twitter indicators are first aggregated to

weekly frequency and combined into a composite sentiment measure. The missing data filter is applied to this combined weekly sentiment series, and firms with more than 5% missing observations are excluded. The final investment universe is defined as the intersection of firms for which weekly stock returns, market capitalization, and combined sentiment data are simultaneously available. After applying these filters, the final sample consists of 455 firms ($N = 455$).

4.3 Data Processing and Variable Construction

The raw datasets obtained from Bloomberg are first cleaned to ensure consistency and data quality. This includes removing duplicate observations, verifying that values lie within economically plausible ranges, and confirming that sentiment scores remain within their theoretical bounds.

Following this initial cleaning, all variables are converted to a weekly frequency to ensure consistency across datasets and alignment with the portfolio rebalancing strategy. For price-based variables, including stock prices and the S&P 500 index, the last available observation within each week is retained. Weekly stock returns are then computed as the weekly percentage price change including dividends. Excess returns are obtained by subtracting the corresponding weekly risk-free rate from stock returns. The same procedure is applied to the S&P 500 index to compute benchmark excess returns.

The risk-free rate is proxied by the yield on the one-month U.S. Treasury bill. The last available annualized yield observed in each week is retained and converted into a weekly rate by dividing by 52. Market capitalization data is also converted to a weekly frequency by retaining the last available observation within each and are subsequently used to compute time-varying market-capitalization weights by normalizing firm-level market capitalizations at each rebalancing date.

Bloomberg News Sentiment and Bloomberg Twitter Sentiment are first aggregated to a weekly frequency by computing the average of daily observations within each week. A composite weekly sentiment score is then constructed by averaging the two indicators when both are available for a given firm/week and retaining the observed value when only one is available. This approach maximizes data availability while preserving cross-firm consistency.

In addition to the composite score, directional sentiment variables are constructed. The positive sentiment variable retains the composite score when it is strictly positive and assigns zero otherwise. In the same way, the negative sentiment variable retains the composite score when it is strictly negative and assigns zero otherwise. Agreement indicators are also constructed: a positive agreement dummy takes the value of one when both Bloomberg News and Twitter sentiment indicators are simultaneously positive for a given firm-week, and zero otherwise, a negative agreement dummy takes the value of one when both are simultaneously negative. These variables allow the identification of firms for which sentiment signals are both strong and directionally consistent across sources.

After constructing the weekly sentiment series, the missing-data filter is applied to the combined sentiment measure. Firms with more than 5% missing weekly observations in the composite sentiment series are excluded. This filtering step is applied after combining the sentiment indicators to avoid unnecessarily discarding firms for which at least one sentiment measure is available.

The covariance matrix of asset returns is estimated using weekly excess returns over a rolling window of 52 weeks. However, as mentioned by Reis et al. (2026), because the number of assets ($N = 455$) substantially exceeds the number of observations in the estimation window ($T = 52$), the sample covariance matrix suffers from poor conditioning and is prone to estimation errors. To address this, the Analytical Nonlinear Shrinkage method of Ledoit and Wolf (2020) is adopted, which optimally adjusts the eigenvalues of the sample matrix using a kernel-based spectral density estimator and a Hilbert transform to correct for the eigenvalue bias that arises when the concentration ratio (N/T) is high.

Following estimation, symmetry is enforced by averaging the matrix with its transpose. A spectral decomposition is then performed, and any negative eigenvalues are set to zero to ensure the matrix is positive semidefinite (Reis et al., 2026). This minimum-Frobenius-distance correction yields a numerically stable and invertible covariance matrix suitable for portfolio optimization.

The risk-aversion parameter (δ) is estimated as the ratio of the average excess return to the variance of the market portfolio (equation (2)), proxied by the S&P 500 index, and is assumed constant over the sample period following Grinold and Kahn (1999).

4.4 Descriptive Statistics

This section presents descriptive statistics for the main variables and provides an overview of the key properties of the dataset.

The final sample covers the period from January 2024 to December 2025 and consists of 104 weekly observations. The investment universe includes 455 US stocks drawn from the S&P 500 index, restricted to firms with continuous data availability across returns, sentiment measures, and market capitalization. All variables are constructed at weekly frequency. The resulting panel is effectively balanced, with missing observations accounting for less than 0.2% of the sentiment data and no missing values in returns or market capitalization.

Table 1 - Summary Statistics of Main Variables

Variable	Observations	Mean	Std. Dev.	Min	Median	Max
Excess returns	47320	0.0019	0.0410	-0.4562	0.0026	0.5574
Composite sentiment	47245	0.0018	0.0815	-0.8197	0.0014	0.5955
Market capitalization	47320	114,667	345,131	4,899	38,301	4,920,507
Risk-free rate	104	0.0009	0.0001	0.0007	0.0008	0.0010

Note: Observations are at the firm-week level over the full sample period (January 2024 - December 2025, 104 weekly rebalancing dates). Returns and sentiment variables are expressed as weekly decimals. Market capitalization is reported in USD thousands. The lower observation count for composite sentiment (47,245 vs. 47,320) reflects a small number of firm-weeks with missing Bloomberg data. The risk-free rate has 104 observations, one per week, as it is a time-series variable rather than a cross-sectional one.

Weekly excess returns display an average of approximately 0.19% (10,4% compounded annualized) with a standard deviation of around 4.1% (29,6% annualized), indicating substantial dispersion in realized returns across firms and time. The average within-week cross-sectional dispersion across stocks is approximately 3.6%, with values ranging between 1.5% and 5.9% across weeks, confirming sufficient cross-sectional variation to support ranking-based strategies (see Table B.1 in Appendix).

The composite sentiment measure has a mean close to zero and a standard deviation of approximately 0.081, consistent with its bounded nature and its construction as an average of underlying indicators. The distribution spans both positive and negative values. When decomposed, positive sentiment is strictly non-negative and negative sentiment is strictly non-positive by construction. The average magnitude of positive sentiment is slightly larger than that of negative sentiment in absolute terms, suggesting a mild asymmetry in the distribution of sentiment signals (see table B.2 in Appendix).

Market capitalization exhibits substantial dispersion, reflecting the heterogeneity of firm sizes within the S&P 500 universe. The risk-free rate remains stable over the sample period with low variability, as expected for short-term Treasury yields.

Table 2 - Descriptive Statistics of Portfolio Return Differentials Used to Construct the View Vector (Q)

Portfolios	Mean	Std. Dev.
RQ1 Ranking-based spreads		
10% cutoff	0.0566	0.0277
20% cutoff	0.0439	0.0210
30% cutoff	0.0368	0.0172
RQ2 Sign-based spreads		
Positive sentiment (10%)	0.0350	0.0191
Positive sentiment (20%)	0.0291	0.0152
Positive sentiment (30%)	0.0268	0.0145
Negative sentiment (10%)	0.0275	0.0193
Negative sentiment (20%)	0.0247	0.0151
Negative sentiment (30%)	0.0243	0.0135
RQ2 Agreement spreads		
Positive agreement	0.0270	0.0164
Negative agreement	0.0190	0.0159

Note: Realized weekly return differentials used as inputs to construct the view vector Q. The view vector is constructed by applying an exponentially weighted average to the presented past realized return differentials. For RQ1, the spread reflects the return differential between the high- and low-sentiment groups. For RQ2, the spread reflects the return differential between the long leg and the short leg of the corresponding relative portfolio, the sentiment-classified group minus the remaining universe for positive views, and the remaining universe minus the sentiment-classified group for negative views.

Table 2 reports descriptive statistics for the return differentials used to construct the view vector Q. The return differentials reported here are substantially larger in magnitude than the average individual stock return reported in Table 1. This is because Table 1 reports the equal weighted average return across all 455 stocks, while Table 2 reports the return differential between stocks at the extremes of the sentiment distribution. These extreme groups exhibit meaningful different realized returns, producing spreads of larger magnitude.

For the sentiment-ranked spreads portfolios, mean return differentials between high and low sentiment firms are positive across all cutoff thresholds and decline monotonically as the cutoff widens, consistent with the expectation that more extreme sentiment classifications yield stronger signals. Volatility also declines with the cutoff, reflecting the more diversified composition of the wider groups.

For the directional sentiment spread portfolios we observe positive mean return differentials regardless of sentiment direction, as both positive and negative views are expressed as relative views in which the long leg is expected to outperform the short leg. The magnitudes of the negative sentiment spread returns are slightly smaller than those of the corresponding positive sentiment across all cutoffs, providing preliminary descriptive evidence that the negative sentiment signal may generate smaller return differentials than the positive signal. Agreement-based spreads display similar directional patterns, though with slightly lower dispersion, reflecting the reduced sample size arising from the requirement that both sentiment sources agree in sign.

Table 3 - Portfolio Composition of Sentiment-Classified Groups

Portfolio	Long leg (positive view) avg MCap (b$)	Short leg (negative view) avg MCap (b$)	Avg N long	Avg N short
RQ1: Ranking-Based Constituents				
10% cutoff	110,7	103,5	45	45
20% cutoff	106,3	101,2	91	91
30% cutoff	104,7	108,4	136	136
RQ2: Sign-Based Constituents				
Positive sentiment (10%)	110,7	115,1	45	410
Positive sentiment (20%)	105,6	117,0	90	365
Positive sentiment (30%)	103,9	119,4	134	321
Negative sentiment (10%)	115,9	103,5	410	45
Negative sentiment (20%)	118,0	101,2	364	91
Negative sentiment (30%)	117,1	109,8	321	134
RQ2: Agreement Constituents				
Positive agreement	149,6	104,7	115	340
Negative agreement	99,3	193,9	377	78

Note: Average market capitalization (USD billions) and average number of stocks per leg, computed over the full sample period. For RQ1, long and short legs correspond to high- and low-sentiment groups. For RQ2, in Positive Sentiment, the long leg is the sentiment-classified group and the short leg the remaining universe of stocks. In Negative Sentiment, the short leg is the sentiment-classified group and the long leg remaining universe of stocks.

Table 3 reports the average market capitalization of the long and short legs of each portfolio, together with the average number of stocks in each leg, computed over the full sample period. Market capitalizations are reported in USD billions.

For the ranked sentiment spread portfolios, the long and short legs exhibit broadly comparable average market capitalizations across all cutoff thresholds, ranging from approximately $100bn to $111bn. This indicates that the high and low sentiment classifications do not systematically select for firms of different sizes, and that the observed return differentials

are unlikely to reflect a size premium. The number of stocks in each leg is symmetric by construction, increasing from approximately 45 stocks at the 10% cutoff to 136 at the 30% cutoff.

For the sign-based portfolios, the short leg, which comprises the full complement of non-classified stocks, is substantially larger than the long leg by construction, with approximately 410 stocks in the short leg against 45 in the long leg at the 10% cutoff. Market capitalizations across the two legs remain broadly comparable, suggesting no systematic size tilt in the relative view portfolios. The agreement-based portfolios display a more notable pattern: the positive agreement portfolio has a substantially larger long leg market cap ($149.6bn vs $104.7bn), while the negative agreement portfolio shows the reverse, with the short leg averaging $193.9bn against a long leg of $99.3bn. This suggests that when both sentiment sources simultaneously agree on negative sentiment, the flagged firms tend to be larger-cap companies, which may reflect the greater media coverage and analyst attention that large firms attract.

To further characterize the sentiment-classified groups, the most persistently classified stocks at the 10% cutoff are examined. Table B.3 (in Appendix) reports the twenty stocks most frequently classified as high or low sentiment across the sample period. High-sentiment stocks are predominantly concentrated in the technology and semiconductor sectors, including Broadcom, Lam Research, Arista Networks, Applied Materials, and Micron Technology, alongside strong performers in consumer discretionary and financials. This is consistent with the exceptional performance of AI-related and technology equities during the 2024-2025 sample period, which generated persistently positive media and social sentiment. Low-sentiment stocks are notably concentrated in healthcare, including UnitedHealth, Elevance Health, Stryker, IQVIA, and Humana, alongside utilities, consumer staples, and retail names such as Dominion Energy, Kraft Heinz, Dollar General, and Starbucks. The healthcare concentration in the low-sentiment group likely reflects regulatory uncertainty and earnings pressures facing managed care companies during this period. Interestingly, NextEra Energy appears in both groups across different weeks, reflecting the time-varying nature of sentiment classifications. These findings suggest that sentiment classification in this sample is meaningfully related to sector-level news flow and earnings dynamics rather than being randomly distributed across the cross-section.

Beyond the realized sentiment-based portfolios, the empirical framework relies on the construction of the view variable Q, obtained as an exponentially weighted average of past realized return differentials. This transformation introduces a smoothing effect that reduces the volatility of the signal relative to the underlying spreads. While realized spreads display standard deviations ranging from 1.4% to 2.8% depending on the cutoff (see table C.1 in Appendix), the exponential weighting scheme produces a considerably more stable series, with volatility reduced by roughly 80-85%. This is achieved by dampening short-term fluctuations and assigning greater weight to recent observations, ensuring that the views incorporated into the Black-Litterman framework remain responsive while avoiding excessive sensitivity to weekly noise.

Regarding the covariance matrix, diagnostic analysis of the unregularized sample covariance matrix reveals severe ill-conditioning, with near-singularity and unstable inversion. In contrast, the shrinkage-based covariance matrices exhibit substantially improved numerical properties. Across the sample period, the average condition number is approximately 310, with values remaining below 420 in all cases, indicating a well-conditioned and stable matrix. The minimum eigenvalue remains strictly positive, averaging 0.0006, confirming that the matrices are positive definite and suitable for inversion. The scale of the estimated variances is also consistent with the observed return volatility, with an average variance of 0.0016, corresponding to a weekly standard deviation of approximately 4.0% (see table D.1 in Appendix).

These results indicate that the shrinkage estimator effectively stabilizes the covariance matrix while preserving economically meaningful variation in the data. This is essential for the implementation of the Black-Litterman framework, where reliable covariance estimates are required to obtain robust portfolio allocations.

Taken together, the descriptive statistics provide strong support for the suitability of the dataset. Missing observations are negligible, accounting for less than 0.2% of the composite sentiment score, and the panel is effectively complete. The presence of substantial cross-sectional dispersion in returns, combined with persistent sentiment-based return differentials, confirms that the data is both statistically informative and economically meaningful. These properties provide a reliable foundation for the construction of model inputs and for the implementation of the Black-Litterman framework in the subsequent analysis.

5. Results

This section presents the empirical results of the analysis and discusses their implications. Section 5.1 evaluates the risk-adjusted performance of the RQ1 sentiment-informed portfolios across all tau values. Section 5.2 examines the asymmetry between positive and negative sentiment signals, addressing RQ2. Section 5.3 presents a subperiod analysis assessing the regime-dependence of the results. Section 5.4 examines the predictive content of the view vector Q and the robustness of the findings to alternative view constructions, providing a diagnostic assessment of the Q construction and its underlying assumptions Section 5.5 addresses the practical implications of portfolio turnover and transaction costs. All results are based on a 52-week exponentially weighted window with $\lambda = 0.95$, and are reported across five values of the prior uncertainty parameter $\tau \in \{0.01, 0.05, 0.10, 0.20, 1.00\}$.

Table 4 - Portfolio Composition Statistics

Panel A - RQ1 tests

Portfolio	Avg N stocks	HHI	Wtd avg beta	Wtd avg mcap ($bn)
		RQ1 portfolios		
10% cutoff	352	0.177	0.86	787
20% cutoff	335	0.167	0.92	1,058
30% cutoff	355	0.151	0.87	1,193

Panel B – RQ2 tests

Portfolio	Avg N stocks	HHI	Wtd avg beta	Wtd avg mcap ($bn)
RQ2 - Agreement				
Positive sentiment	308	0.178	0.99	1,646
Negative sentiment	347	0.053	0.63	871
RQ2 - Sign (10%)				
Positive sentiment	353	0.198	0.78	791
Negative sentiment	310	0.078	0.85	1,296
RQ2 - Sign (20%)				
Positive sentiment	334	0.164	0.85	1,036
Negative sentiment	316	0.087	0.83	1,465
RQ2 - Sign (30%)				
Positive sentiment	346	0.139	0.80	1,149
Negative sentiment	331	0.093	0.82	1,499

Note: Average portfolio characteristics computed over the full sample period. Statistics are based on the original view specification ($\lambda = 0.95$, 52-week window). Avg N stocks is the average number of assets receiving non-zero portfolio weight. HHI is the Herfindahl-Hirschman Index of portfolio concentration, computed as Σw_i^2 at each date and averaged over time. Weighted average beta is the portfolio-weight-average of individual stock market betas estimated from a full-sample OLS regression of weekly excess returns on S&P 500 excess returns. Weighted average market capitalization is reported in USD billions.

The average number of stocks receiving non-zero portfolio weight ranges from 308 to 355 across all specifications, confirming that the BL framework distributes weights broadly rather than concentrating positions in a small number of assets. HHI values are consistently low, ranging from 0.053 to 0.198, with the agreement negative portfolio being by far the most diversified, reflecting its well-populated long leg of approximately 377 stocks relative to a short leg of only 78. Weighted average market betas are uniformly below one across all specifications, ranging from 0.63 to 0.99, indicating that the BL optimization systematically tilts toward lower-beta large-cap stocks rather than introducing additional systematic risk relative

to the market. This pattern is most pronounced in the agreement negative portfolio, which exhibits the lowest weighted beta of 0.63 and correspondingly the lowest portfolio volatility in the analysis. Weighted average market capitalizations range from USD 787bn to USD 1.65 trillion, consistent with the large-cap composition of the S&P 500 universe and reflecting the tendency of the BL shrinkage mechanism to assign greater weight to the largest constituents. These characteristics are broadly stable across tau values and cutoff specifications, confirming that the choice of τ primarily affects the degree of deviation from equilibrium rather than the fundamental nature of the active portfolios.

5.1 RQ1: Performance of Sentiment-Informed Portfolios

Table 5 reports the performance of three baseline sentiment-informed portfolios across all tau values for the different cutoffs alongside the S&P 500 benchmark over the full sample period from January 2024 to December 2025.

Table 5 - RQ1 BL Portfolio Performance Across Tau Values

τ	Avg Ret	Std Dev	Sharpe	Sortino	Max DD	Turnover	IR
Benchmark	**0.32%**	**2.07%**	**1.123**	**1.500**	**-17.47%**	**-**	**-**
			10% cutoff				
0.01	0.30%	2.09%	1.034	1.427	-19.04%	32.5%	-0.391
1.00	0.33%	2.72%	0.875	1.690	-16.68%	80.4%	0.024
			20% cutoff				
0.01	0.32%	2.09%	1.104	1.539	-19.09%	30.8%	-0.039
1.00	0.23%	2.77%	0.587	0.825	-20.51%	72.9%	-0.290
			30% cutoff				
0.01	0.30%	2.07%	1.039	1.435	-18.22%	31.1%	-0.444
1.00	0.28%	2.57%	0.773	1.144	-17.54%	70.0%	-0.154

Note: Average return and standard deviation are weekly figures. Sharpe and Sortino ratios are annualized by $\sqrt{52}$, assuming independent weekly returns. Maximum drawdown is the largest peak-to-trough cumulative loss.

Turnover is average weekly one-way portfolio turnover. IR is the annualized information ratio. BL Sentiment Views portfolios are long-only Sharpe-ratio maximizing portfolios with the long leg comprising the high-sentiment group and the short leg the low-sentiment group, weights proportional to market capitalization.

The results reveal a clear and systematic relationship between tau and portfolio performance. At $\tau = 0.01$, the posterior expected returns remain very close to the equilibrium prior, the average correlation between μ_{BL} and Π is 0.997 across all three cutoffs (see table D.3 in Appendix), resulting in conservative deviations from the market portfolio and substantially lower turnover, ranging from 30.8% to 32.5% compared to 70.0% to 80.4% at $\tau = 1.0$. Despite this conservative positioning, at $\tau = 0.01$, the three specification achieve values that approach the benchmark's 1.123, (1.034, 1.104, and 1.039 for the 10%, 20%, and 30% cutoffs respectively). As tau increases, the posterior diverges further from equilibrium, turnover rises substantially, and risk-adjusted performance deteriorates across all cutoff specifications. At $\tau = 1.0$, the portfolios produce Sharpe ratios of 0.875, 0.587, and 0.773, all below the benchmark.

This monotonic pattern is the central finding of RQ1 and is directly connected to the quality of the sentiment signal. Since the view vector Q exhibits mean-reverting rather than momentum-like behavior at the weekly frequency, as confirmed by the signal diagnostics in Section 5.4, large deviations from equilibrium are systematically harmful. The posterior is being pushed away from equilibrium by views that do not accurately predict future returns, introducing tracking error without appropriate return improvements. Smaller tau values limit this deviation, preserving the portfolio's proximity to the efficient market allocation while still incorporating a modest directional tilt from the sentiment signal.

Importantly, the underperformance at $\tau = 1.0$ is not driven by lower average returns. The three portfolios earn weekly average excess returns of 0.33%, 0.23%, and 0.28%, broadly comparable to the benchmark's 0.32%. Rather, it reflects higher realized volatility, ranging from 2.57% to 2.77%, indicating that large sentiment-based tilts introduce tracking error without appropriate return improvements. Despite this overall picture, the 10% cutoff portfolio at $\tau = 1.0$ provides a partial exception on downside risk metrics. Its Sortino ratio of 1.690 exceeds the benchmark's (1.500), suggesting that a portion of the excess volatility corresponds to upside rather than downside deviations. Its maximum drawdown of -16.68% is also marginally better than the benchmark's -17.47%, providing some evidence that the most

concentrated sentiment specification offers a degree of downside protection during adverse market episodes.

Taken together, the evidence from RQ1 provides nuanced support for the hypothesis that Bloomberg sentiment indicators improve risk-adjusted portfolio performance within a Black-Litterman framework. At small tau values, all three cutoff specifications achieve Sharpe ratios approaching the benchmark's 1.123, while at $\tau = 1.0$ all fall below it, a monotonic deterioration consistent with the mean-reverting nature of the sentiment signal documented in Section 5.4. This finding contrasts with Colasanto et al. (2022) and Hung et al. (2024), who document outperformance from sentiment-informed BL portfolios, however, this difference is more likely attributable to their use of forward-looking Q construction methods, deep-learning price predictions and Monte Carlo simulation respectively, than to any superiority of their underlying sentiment measures. In terms of risk-adjusted performance, the best-performing specification in this study achieves a Sharpe ratio of 1.104, marginally lower than Colasanto et al.'s (2022) reported value of 1.14, though their evaluation covers a single 5-day window. Hung et al.'s (2024) reported Sharpe of 13.0, which is not directly comparable given their 7-stock universe and single-year evaluation period (see Table A.1 in Appendix for a full methodological comparison). This finding is consistent with Bessler, Opfer, and Wolff (2017), who show that Black-Litterman portfolios tend to outperform passive benchmarks only when the investor views incorporated into the model have genuine predictive content, a condition that appears to hold only partially in the current setting. More broadly, the difficulty of outperforming a passive benchmark in large-cap U.S. equity markets is well documented (DeMiguel et al., 2009), and the 2024-2025 sample period, characterized by a sustained bull market, represents a particularly challenging environment for cross-sectional sentiment-based strategies (Stambaugh et al., 2012).

5.2 RQ2: Asymmetry Between Positive and Negative Sentiment

Table 6 reports the performance of portfolios constructed using positive and negative sentiment signals separately, for both the agreement-based and sign-based classification approaches across different values of tau.

Table 6 - RQ2 Portfolios Performance

τ	Avg Ret	Std Dev	Sharpe	Sortino	Max DD	Turnover	IR
			Agreement positive				
0.01	0.33%	2.07%	1.146	1.602	-18.35%	23.0%	0.159
1.00	0.19%	2.46%	0.571	0.951	-26.22%	61.1%	-0.486
			Agreement negative				
0.01	0.33%	2.07%	1.157	1.578	-18.06%	14.9%	0.328
1.00	0.11%	1.78%	0.442	0.669	-17.16%	62.8%	-1.078
			Sign positive 10%				
0.01	0.32%	2.01%	1.152	1.575	-16.76%	29.3%	-0.036
1.00	0.39%	2.57%	1.103	1.879	-15.02%	73.6%	0.216
			Sign positive 20%				
0.01	0.35%	2.03%	1.250	1.730	-16.92%	26.1%	0.662
1.00	0.23%	2.59%	0.644	0.873	-18.78%	65.7%	-0.290
			Sign positive 30%				
0.01	0.34%	2.04%	1.205	1.650	-17.26%	25.7%	0.418
1.00	0.21%	2.31%	0.668	1.069	-16.51%	62.4%	-0.397
			Sign negative 10%				
0.01	0.32%	2.13%	1.096	1.518	-18.97%	16.1%	0.018
1.00	0.16%	2.26%	0.510	0.748	-16.36%	61.5%	-0.698
			Sign negative 20%				
0.01	0.32%	2.14%	1.073	1.484	-19.09%	19.6%	-0.082
1.00	0.15%	2.12%	0.523	0.800	-18.83%	58.7%	-0.759
			Sign negative 30%				
0.01	0.29%	2.11%	0.982	1.375	-19.28%	23.3%	-0.706
1.00	0.09%	2.09%	0.313	0.475	-24.43%	61.4%	-1.010

Note: Average return and standard deviation are weekly figures. Sharpe and Sortino ratios annualized by √52,

assuming independent weekly returns. Maximum drawdown is the largest peak-to-trough cumulative loss. Turnover is average weekly one-way portfolio turnover. IR is the annualized information ratio relative to the S&P 500 benchmark. For brevity, only $\tau = 0.01$ and $\tau = 1.00$ are reported; full results across all five τ values are in Table E.2. Agreement-based portfolios classify firms as positive (negative) only when both Bloomberg News and Twitter sentiment indicators are simultaneously positive (negative). Sign-based portfolios apply a percentile cutoff with an additional constraint requiring the composite sentiment score to be strictly positive (negative). For positive sentiment portfolios, the long leg comprises the sentiment-classified group and the short leg the remaining universe. For negative sentiment portfolios, the long leg comprises the remaining universe and the short leg the sentiment-classified group. In both cases, weights within each leg are proportional to market capitalization.

The results reveal two consistent patterns. First, as documented for RQ1 in Section 5.1, performance improves monotonically as tau decreases across all specifications. At τ=0.01, five out of eight specifications exceed the benchmark's 1.123, with the sign positive 20% cutoff achieving the highest Sharpe ratio of 1.250, while the three sign-based negative sentiment portfolios fall below the benchmark, ranging from 0.982 to 1.096, At $\tau = 1.0$, only the sign positive 10% cutoff approaches the benchmark, with a Sharpe ratio of 1.103. This confirms that the tau sensitivity documented in Section 5.1 is a general feature of the framework rather than a characteristic specific to the baseline specification.

Second, and more substantively for RQ2, positive sentiment portfolios consistently outperform their negative sentiment counterparts on every performance metric across all tau values and cutoff specifications. For the sign-based approach at $\tau = 1.0$, the positive 10% cutoff achieves a Sharpe ratio of 1.103 compared to 0.510 for the corresponding negative portfolio, and a Sortino ratio of 1.879 against 0.748. As the cutoff widens to include a larger proportion of the cross-section, performance deteriorates for both sentiment directions, while the gap between positive and negative portfolios remains consistent.

The Sortino ratios reinforce this pattern. The sign positive 10% portfolio at $\tau = 1.0$ achieves a Sortino ratio of 1.879, exceeding the benchmark's 1.500 and representing the strongest downside risk-adjusted performance observed across all scenarios. This finding suggests that, at the most concentrated specification, the positive sentiment signal provides meaningful protection against left-tail losses even as it fails to generate superior total risk-adjusted performance at higher tau values. By contrast, negative sentiment portfolios exhibit

Sortino ratios uniformly below the benchmark at $\tau = 1.0$, and the 30% cutoff negative portfolio reaches the lowest value in the analysis at 0.475.

The information ratios also confirm the pattern. The sign positive 10% portfolio at $\tau = 1.0$ achieves a mildly positive information ratio of 0.216, indicating that the active bet generates some return per unit of tracking error, while all negative sentiment information ratios are negative at $\tau = 1.0$, reflecting a consistent inability to generate active returns sufficient to justify deviations from the benchmark.

The evidence does not support the hypothesis that negative sentiment carries a stronger or more exploitable signal than positive sentiment. Across all cutoff specifications and classification approaches, positive sentiment portfolios generate higher average returns and better risk-adjusted performance, with Sharpe ratios uniformly exceeding those of the corresponding negative sentiment portfolios. This finding is inconsistent with the behavioral finance prediction that investors react more strongly to adverse information due to loss aversion and negativity bias (Kahneman and Tversky, 1979; Da et al., 2011), which would imply larger price dislocations among negative-sentiment stocks and greater portfolio benefits from incorporating negative signals. It also contrasts with evidence that sentiment effects are asymmetric across market conditions, with negative sentiment having a disproportionate influence on returns relative to positive sentiment (Da et al., 2011).

However, the studies identify sentiment asymmetry primarily during periods of market stress and elevated volatility, conditions that were largely absent from the 2024-2025 sample period, which was characterized by a sustained equity market rally and strong upward momentum in large-cap U.S. stocks. Baker and Wurgler (2006) show that sentiment effects are regime-dependent and most pronounced for stocks that are difficult to value or arbitrage, a category that large-cap S&P 500 constituents do not naturally belong to. In such an environment, positive sentiment is more likely to be aligned with the prevailing directional trend, while negative sentiment, which requires stocks to underperform the market, is systematically frustrated by the broad upward drift. This interpretation is consistent with the broader evidence that sentiment effects on returns are most pronounced during periods of market stress, elevated uncertainty, and limited arbitrage capacity (Baker and Wurgler, 2006; Shleifer and Vishny, 1997).

While these findings are at odds with theoretical predictions from the behavioral finance

literature, they are consistent with the specific market dynamics of the 2024-2025 sample period. The following section examines whether the full-sample results mask meaningful variation across different market environments, and whether the regime-dependence of sentiment effects can explain the patterns documented above.

5.3 Subperiod Analysis

To assess whether the full-sample results mask meaningful variation across different market environments, the sample period is divided into two annual subperiods: January to December 2024 and January to December 2025. The first subperiod was characterized by a stable and sustained bull market in U.S. large-cap equities, while the second included the sharp market correction of April 2025 driven by fears of a global trade war, followed by a rapid recovery. Table 7 reports Average Returns, Sharpe ratio, Sortino ratio, Information Ratio, and Average Turnover for the RQ1 portfolios at $\tau = 0.01$ and $\tau = 1.0$ across both subperiods. The full subperiod analysis including RQ2 specifications is reported in Appendix F.

Table 7 - Subperiod Analysis (RQ1)

Portfolio		2024					2025				
	τ	Avg Ret	*Sharpe*	Sortino	IR	Turnover	Avg Ret	Sharpe	Sortino	IR	Turnover
Benchmark		0.40%	1.691	2.558	-	-	0.25%	0.742	0.965	-	-
RQ1 10%	0.01	0.39%	1.621	2.253	-0.144	36.7%	0.21%	0.630	0.895	-0.655	28.4%
RQ1 10%	1.00	0.73%	1.697	3.458	0.915	78.3%	-0.07%	-0.210	-0.365	-1.111	82.5%
RQ1 20%	0.01	0.39%	1.680	2.441	-0.077	32.5%	0.25%	0.729	1.038	-0.001	29.2%
RQ1 20%	1.00	0.54%	1.253	1.567	0.391	68.7%	-0.09%	-0.284	-0.443	-1.142	77.0%
RQ1 30%	0.01	0.35%	1.491	2.156	-0.788	34.0%	0.24%	0.732	1.013	-0.079	28.3%
RQ1 30%	1.00	0.56%	1.407	1.984	0.476	65.5%	-0.01%	-0.020	-0.030	-0.883	74.5%

Note: Sharpe and Sortino ratios are annualized by √52, assuming independent weekly returns. IR is the annualized information ratio relative to the S&P 500 benchmark. Turnover is the average one-way weekly portfolio turnover. The first subperiod covers January to December 2024, characterized by a sustained bull market in U.S. large-cap equities. The second subperiod covers January to December 2025, which included the sharp

market correction of April 2025 driven by trade war concerns, followed by a rapid recovery. Full subperiod results across all τ values and all RQ2 specifications are reported in Tables F.1 and F.2.

The results reveal a striking divergence between the two years. In 2024, the benchmark achieves a Sharpe ratio of 1.691, while the RQ1 10% cutoff reaches 1.621 at $\tau = 0.01$ and 1.697 at $\tau = 1.0$, with the 20% and 30% cutoffs similarly strong, achieving Sharpe ratios of 1.680 and 1.491 at $\tau = 0.01$ and 1.253 and 1.407 at $\tau = 1.0$. Information ratios are positive for most $\tau = 1.0$ specifications, confirming that the sentiment-based active bets generated value relative to the benchmark during this period.

In 2025, the picture reverses entirely. The benchmark Sharpe ratio falls to 0.742, while all RQ1 portfolios at $\tau = 1.0$ produce negative Sharpe ratios, -0.210, -0.284, and -0.020 for the 10%, 20%, and 30% cutoffs respectively. At $\tau = 0.01$, portfolios hold up considerably better, with all three cutoffs maintaining positive Sharpe and Sortino ratios, 0.630, 0.729, and 0.732 respectively, remaining positive when the corresponding $\tau = 1.0$ specifications are all negative. This confirms that conservative tau calibration provides meaningful downside protection during volatile market environments, in addition to its pre-cost performance benefits documented in Section 5.1.

The return differential is equally stark: while $\tau = 0.01$ portfolios retain positive average weekly returns of 0.21%, 0.25%, and 0.24% across the three cutoffs in 2025, $\tau = 1.0$ portfolios generate negative average returns of -0.07%, -0.09%, and -0.01% respectively, confirming that the 2025 deterioration at high tau values reflects genuine capital losses rather than merely higher volatility around a positive mean.

Turnover is higher in 2025 across all specifications, the 10% cutoff rises from 78.3% to 82.5%, implying that transaction costs were most damaging precisely when returns were deteriorating. At $\tau = 0.01$, turnover remains stable and low across both subperiods (28%-37%), confirming that small τ values insulate the strategy from the sentiment classification instability that drives high rebalancing costs at $\tau = 1.0$.

The RQ2 subperiod results, reported in full in Table F.2 (in Appendix), exhibit the same general pattern but also shed additional light on the asymmetry between positive and negative sentiment documented in Section 5.2. In 2024, the performance gap between positive and negative sentiment portfolios is large, consistent with positive sentiment being aligned with the prevailing directional trend. In 2025, this gap narrows substantially, with sign positive

10% collapses while sign negative holds up. While negative sentiment does not dominate positive sentiment unconditionally in the full sample, it demonstrates greater relative resilience during the stress period of 2025, consistent with the prediction that negative sentiment effects are more pronounced during periods of market stress (Yu & Yuan, 2011). The full-sample finding of positive sentiment dominance documented in Section 5.2 therefore appears to be partly a consequence of the specific market regime rather than a general property of sentiment-based strategies.

Taken together, these findings reframe the full-sample results as regime-specific rather than unconditional. The sentiment strategy generated genuine risk-adjusted value in 2024 but proved fragile to the 2025 market shock, consistent with evidence that sentiment-based strategies are highly sensitive to the prevailing market regime (Baker and Wurgler, 2006). The signal diagnostics presented in Section 5.4 provide further insight into the underlying mechanism driving this pattern.

5.4 Signal Diagnostics: The Predictive Content of the View Vector

The performance results documented in Sections 5.1 and 5.2 raise an important question: does the underperformance of sentiment-informed portfolios reflect a failure of the portfolio construction mechanism, does it indicate that the view vector Q itself lacks predictive content, or both? These two explanations have different implications. If the views are informative but the portfolio construction is too conservative, for instance, due to excessive shrinkage toward equilibrium, high turnover, or the long-only constraint, then modifications to the implementation could potentially recover the signal. However, if the views themselves are not predictive of subsequent returns, then no adjustment to the portfolio construction would be expected to improve performance systematically.

To distinguish between these two possibilities, this section examines two complementary diagnostics. First, the average and median correlations between the posterior expected returns μ_{BL} and the prior equilibrium returns Π are reported, to assess the extent to which the Black-Litterman framework actually incorporates the sentiment views into the posterior. Second, a predictive regression is estimated to directly test whether the view vector Q contains statistically meaningful information about subsequent realized return differentials.

Table D.3 (in appendix) shows that the average correlations between μ_{BL} and Π provide reassuring evidence that the Black-Litterman framework is genuinely incorporating the sentiment views into the posterior expected returns. Average correlations range from 0.202 to 0.510 at $\tau = 1.0$, indicating that the posterior expected returns diverge substantially from the prior equilibrium in many periods. The median correlations are higher, ranging from 0.454 to 0.826, suggesting that the distribution of correlations is left-skewed, with a small number of weeks in which the posterior diverges sharply from the prior. Across all tau values, correlation at $\tau = 0.01$ are uniformly above 0.996, confirming that portfolios barely deviate from the market at small tau, while at $\tau = 1.0$ the results reflect substantial divergence driven by the sentiment views. This pattern implies that the Black-Litterman mechanism is functioning as intended, sentiment views do move the posterior, and that the underperformance documented in the previous sections is unlikely to be attributable to over-shrinkage toward equilibrium.

To assess whether the signal itself is informative, a predictive regression directly tests whether the view vector Q contains statistically meaningful information about subsequent realized return differentials:

$$r_{t+1} = \alpha + \beta \cdot Q_t + \varepsilon_{t+1} \tag{8}$$

where r_{t+1} denotes the realized return differential of the relevant portfolio in week $t+1$ and Q_t is the view vector at week t. A positive and statistically significant β would indicate that the view vector contains exploitable predictive content. Across almost all specifications, the estimated coefficient on Q is negative, indicating that higher values of the view vector are associated with lower realized return differentials in the following week, consistent with short-term mean reversion rather than the momentum-like persistence assumed by the exponential weighting scheme.

Table G.2 (in appendix) shows this pattern appears consistently across both RQ1 and RQ2 specifications, with $\beta(Q)$ estimates of -0.738, -0.395 and -0.850 for the RQ1 10%, 20%, and 30% cutoffs respectively, suggesting it is a general feature of the view construction approach. Most coefficients are statistically insignificant. The exceptions are the sign positive 30% specification ($p = 0.022$), which is significant at the 5% level, and the RQ1 30% cutoff ($p = 0.079$) and agreement positive ($p = 0.087$), both marginally significant at the 10%

level. The negative β estimates remain consistent across all three, confirming the mean-reverting pattern.

To assess whether the problem is specific to the exponential weighting scheme, two alternative view construction approaches are considered. The first replaces the 52-week exponentially weighted window with a shorter 4-week window and a lower decay parameter of $\lambda = 0.50$, producing a more reactive view vector that places substantially greater weight on the most recent observations. The second replaces the level of the sentiment-based spread with its week-on-week change, capturing shifts in the direction of sentiment rather than its accumulated level. The results, reported in Table G.1 (in Appendix), are qualitatively identical across all three view construction approaches, no specification consistently outperforms the benchmark on the Sharpe ratio, and the ranking of scenarios is preserved throughout. The most notable finding is that the short-window specification produces modestly higher Sharpe ratios for negative sentiment portfolios, for instance, the agreement negative portfolio improves from 0.442 to 0.617, consistent with negative sentiment effects being short-lived and better captured by a more reactive view. However, these improvements are modest and do not alter the fundamental conclusion.

Furthermore, using the prediction errors from the predictive regressions reported in Table G.2 (in Appendix) to construct a data-driven view uncertainty matrix, yields dynamic Omega values averaging between 0.0002 and 0.0008 across specifications, substantially smaller than the static $\Omega = P\Sigma P'$ values. A smaller Omega implies higher confidence in the views, meaning such a specification would assign greater weight to views that the data shows are mean-reverting and uninformative, reinforcing rather than resolving the fundamental problem.

Taken together, these findings allow the two hypotheses raised at beginning of this section to be assessed. The correlation diagnostics confirm that the first hypothesis, a portfolio construction failure can be largely ruled out: the BL framework is genuinely incorporating the views into the posterior, and conservative tau calibration already addresses the shrinkage concern. The predictive regressions point instead toward the second hypothesis: the view vector Q lacks statistically meaningful predictive content at the weekly frequency when constructed from historical return spreads. However, attributing this finding to the sentiment signal itself would be premature. Bloomberg's indicators are constructed on a large

professional corpus, provide a continuous sentiment scale rather than a binary positive/negative classification, and are at least as informative as the custom NLP measures employed in comparable studies that document stronger performance. The more precise diagnosis is that the limitation lies in the method used to construct Q from the signal. Using the historical return spread as a direct proxy for expected return differentials embeds a backward-looking, momentum-like assumption, an assumption that the predictive regressions show is inconsistent with the data. The problem is therefore not the model, not the signal, but the Q construction methodology: the way the sentiment information is translated into a forward-looking view. The Bloomberg sentiment signal, whether exponentially weighted over a long or short window or expressed as a level or a change, does not yield a predictive Q when used in this way. Unlike the present study, which derives Q directly from historical sentiment spreads, an alternative approach would be to estimate Q from predictive regressions mapping sentiment signals to forward return differentials, or to simulate future price paths via Monte Carlo methods as in Colasanto et al. (2022), both of which would replace the backward-looking spread with a genuine forecast.

5.5 Practical Implications: Portfolio Turnover and Transaction Costs

Portfolio turnover is high across all specifications, ranging from 61% to 81% on a weekly basis at $\tau = 1.0$ and dropping to 29%-37% at $\tau = 0.01$, with meaningful implications for after-cost performance. A full transaction cost analysis is reported in Appendix H, examining returns net of 5, 10, and 20 basis points per one-way transaction. The results confirm that at realistic cost levels of 10 basis points, after-cost returns fall well below the benchmark across most specifications at $\tau = 1.0$, while $\tau = 0.01$ portfolios prove considerably more resilient due to their substantially lower turnover, which reduces the transaction cost burden without eliminating it entirely. This is consistent with the broader literature documenting that high-frequency rebalancing imposes transaction cost burdens that are difficult to overcome with modest return signals (Novy-Marx & Velikov, 2016; Frazzini et al., 2014).

6. Conclusions

This study examines whether standardized Bloomberg sentiment indicators improve risk-adjusted portfolio performance within a Black-Litterman framework, and whether positive and negative sentiment signals carry asymmetric informational content.

With respect to RQ1, the evidence provides nuanced support for the hypothesis that Bloomberg sentiment indicators improve risk-adjusted portfolio performance within a Black-Litterman framework. The answer depends critically on the calibration of τ. At $\tau = 0.01$, the three cutoff specifications achieve Sharpe ratios of 1.034, 1.104, and 1.039 for the 10%, 20%, and 30% cutoffs respectively, approaching the benchmark's 1.123, with average weekly returns of 0.30%, 0.32%, and 0.30% broadly comparable to the benchmark's 0.32%, indicating that the underperformance is driven by higher volatility rather than lower returns. Sharpe ratios fall well below the benchmark across all cutoffs. This monotonic deterioration reflects the mean-reverting nature of the sentiment signal documented in Section 5.4, which penalizes large deviations from equilibrium weights. The subperiod analysis further reveals that performance was meaningfully stronger in 2024, with $\tau = 1.0$ specifications marginally exceeding the benchmark Sharpe of 1.691 in that subperiod before proving fragile to the 2025 market shock, while $\tau = 0.01$ portfolios approached but did not consistently exceed the benchmark in either subperiod. Transaction cost analysis confirms a dual benefit of small tau: better pre-cost performance combined with substantially lower turnover, reducing but not eliminating the transaction cost burden at realistic cost levels of 10 to 20 basis points. Relative to the most closely related studies, this study reports more modest but more credible performance improvements, evaluated over a broader universe, a longer horizon, and using standardized rather than custom-built subjective or black-box sentiment measures.

With respect to RQ2, the full-sample results do not support an asymmetric advantage of negative sentiment signals over positive sentiment signals. Positive sentiment portfolios uniformly outperform their negative counterparts across all cutoff specifications and tau values on every performance metric considered. Yet the subperiod analysis reveals an important nuance: the performance gap between positive and negative sentiment portfolios narrows substantially in 2025, with negative sentiment portfolios demonstrating greater relative resilience during the market stress period. The Sharpe ratio for the agreement negative portfolio improves from a Sharpe of 0.338 in 2024 to 0.551 in 2025, providing partial support for the

behavioral finance asymmetry hypothesis (Baumeister et al., 2001). The full-sample dominance of positive over negative sentiment therefore appears to be partly a consequence of the sample period, a bull market regime, rather than a general property of sentiment-based strategies.

The signal diagnostics presented in Section 5.4 provide an important unifying insight. The BL framework is functioning as intended, sentiment views do move the posterior away from equilibrium, as confirmed by posterior-prior correlations ranging from 0.202 to 0.510 at $\tau = 1.0$. However, predictive regressions reveal that the view vector Q exhibits mean-reverting rather than momentum-like behavior at the weekly frequency, explaining why small tau values that limit the influence of the sentiment views systematically outperform larger ones. Alternative view constructions, including a short-window specification and a delta-sentiment approach, do not materially change this conclusion. Computing the variance of rolling prediction errors to construct a dynamic view uncertainty matrix yields values substantially smaller than the static $\Omega = P\Sigma P'$ specification, assigning even greater confidence to views that the data proves uninformative. Taken together, these robustness tests confirm that the primary constraint lies in the construction of Q rather than in the sentiment signal or the model. Bloomberg's standardized indicators are professionally constructed on a large corpus and are unlikely to be inferior to the custom NLP measures used in comparable studies; the failure lies in using the historical return spread as a direct proxy for expected future differentials, an assumption that the mean-reversion evidence squarely contradicts.

This study contributes to the literature in several ways. First, it provides evidence on the practical value of standardized, institutionally available sentiment indicators in a systematic portfolio construction context, complementing prior work that relies on custom-built NLP models such as Colasanto et al. (2022) and Hung et al. (2024). Second, it shows that the choice of τ is a first-order implementation decision with substantial performance implications, a dimension absent from comparable studies. Third, it provides a diagnostic decomposition of portfolio underperformance that distinguishes between the Q construction methodology and the model implementation components, a distinction that is rarely made explicitly in the sentiment-investing literature. Fourth, it documents that the asymmetry between positive and negative sentiment effects is regime-dependent, a finding with direct implications for practitioners. Finally, we construct a practical transparent framework that could be easily replicated by other researchers.

Several limitations constrain the generalizability of the findings, each pointing toward directions for future research. The sample period from January 2024 to December 2025 applied to 455 of the S&P 500 constituents was selected to capture the most recent available market environment, deliberately excluding the 2020-2023 period, which was heavily influenced by the exceptional dynamics of the COVID-19 pandemic and the war in Ukraine, including unprecedented monetary policy interventions and extreme volatility episodes that may not be representative of normal market conditions. The two-year window yields 104 weekly observations, consistent with the rolling estimation window used for the covariance matrix. While a longer sample spanning multiple market cycles would be preferable, extending the window backward would reintroduce the exceptional market dynamics this study deliberately sought to avoid, and forward extension remains constrained by data availability at the time of writing. The findings should therefore be interpreted as evidence on the performance of sentiment-based BL strategies in a specific bull market environment rather than as a general assessment across all market conditions. Future research should examine performance over longer samples, including periods of elevated volatility and market stress where sentiment asymmetry effects are theoretically more pronounced.

The construction of the view vector Q represents a further limitation and the most promising direction for future research. The exponentially weighted historical spread does not exhibit statistically significant predictive power at the weekly frequency, exhibiting mean reversion rather than the momentum assumed by the view construction. A natural next step would be to replace the backward-looking historical spread with a genuine forward-looking forecast, for instance by estimating Q from predictive regressions mapping sentiment signals directly to future return differentials, or by exploring higher rebalancing frequencies more closely aligned with the transitory nature of sentiment signals. The latter approach would need to be weighed against the substantially higher transaction costs that daily rebalancing would entail for a universe of this size, as the results of this study already demonstrate that turnover is a binding constraint even at weekly frequency.

Appendix

A. Peer-Studies Comparison

Table A.1 - Methodological Comparison: This Study vs. Hung et al. (2024) and Colasanto et al. (2022)

	This study	Hung et al. (2024)	Colasanto et al. (2022)
Sentiment source	Bloomberg News & Twitter (standardized)	BERT classifier on Thomson Reuters & CNBC news	FinBERT on Financial Times articles
Stock universe	455 S&P 500 constituents	7 S&P 500 stock	10 NASDAQ-100 stocks
Sample period	Jan 2024 – Dec 2025 (2 years, 104 weeks)	Jan 2015–Nov 2019 (training); Jan–Nov 2019 (testing, 224 days)	Single rebalancing window, Jan–Feb 2021 (5 days)
Portfolio framework	Black-Litterman	Black-Litterman	Black-Litterman
View construction	Historical return spread, exponentially weighted	GRU/LSTM/RNN predicted stock prices	Monte Carlo simulation using Wiener measure + FinBERT score
View type	Relative	Absolute	Absolute
τ sensitivity	Tested across 5 values	Fixed at 0.025	Not reported (fixed scalar s)
Ω specification	Data-driven (PΣP')	Diag ($P\tau\Sigma P'$)	Proportional to view variance
Rebalancing Frequency	Weekly	Daily	Single period
Best Sharpe ratio	1.104 (20% cutoff, τ=0.01)	13.0 (GRU-BL)	1.14 (FinBERT-BL)
Benchmark Sharpe	1.123 (S&P 500, same period)	10.9 (S&P 500, same period)[3]	1.07 (no-views BL portfolio)[4]
Turnover	30.8–80.4%(τ-dependent)	68% (GRU), 84% (RNN)	Not reported

Note: This table summarizes key methodological differences between this study and the two most directly related prior works. All figures are taken directly from the respective papers.

[3] Hung et al. (2024) report Sharpe and Sortino ratios with a percentage sign in their table header (Sh.R(%) and Sr.R(%)), without specifying their annualization methodology. These figures are therefore not directly comparable to the annualized weekly Sharpe ratios reported in this study and Colasanto et al. (2022). No conversion is attempted.

[4] Colasanto et al. (2022) do not use the S&P 500 as a benchmark. Their benchmark Sharpe of 1.07 refers to a no-views Black-Litterman portfolio, that is, the same BL framework applied without any sentiment-based views, evaluated over the same single 5-day rebalancing window. This differs from the passive index benchmark used in this study and makes direct performance comparison unreliable.

B. Data & Descriptive Statistics

Table B.1 - Weekly Cross-Sectional Return Distribution

	Mean	Std Dev	Min	Max
Cross-sectional mean return	0.21%	1.93%	-8.70%	6.00%
Cross-sectional std return	3.47%	0.83%	1.23%	5.86%
Cross-sectional min return	-16.43%	7.53%	-45.62%	-4.44%
Cross-sectional max return	17.53%	7.38%	5.15%	55.74%

Note: Each row summarizes the time-series distribution of a weekly cross-sectional statistic computed across the 455 sample stocks, over 104 weekly rebalancing dates (January 2024 - December 2025). Mean and Std Dev refer to the time-series mean and standard deviation of each cross-sectional statistic across weeks.

Table B.2 - Summary Statistics of Sentiment Variables

Variable	Mean	Std Dev	Min	Max
Composite sentiment	0.0018	0.0815	-0.8197	0.5955
Positive sentiment	0.0272	0.0490	0.0000	0.5955
Negative sentiment	-0.0254	0.0533	-0.8197	0.0000
Positive agreement	0.4612	0.4987	0.0000	1.0000
Negative agreement	0.2891	0.4533	0.0000	1.0000

Note: Statistics are computed at the firm-week level over the full sample period. Positive (negative) sentiment retains the composite score only when strictly positive (negative) and assigns zero otherwise. Agreement dummies equal one when both Bloomberg News and Twitter indicators are simultaneously positive (negative).

Table B.3 - Most Persistently Classified Stocks as the 10% Cutoff

Panel A

High Sentiment - persistent stocks			Low Sentiment - persistent stocks		
Ticker	**Company**	**Sector**	**Ticker**	**Company**	**Sector**
AVGO UW	Broadcom	Technology	D UN	Dominion Energy	Utilities
ORLY UW	O'Reilly Automotive	Consumer Discretionary	DG UN	Dollar General	Consumer Discretionary
COST UW	Costco	Consumer Staples	KHC UW	Kraft Heinz	Consumer Staples
LRCX UW	Lam Research	Technology (Semiconductors)	F UN	Ford	Consumer Discretionary
EOG UN	EOG Resources	Energy	UNH UN	UnitedHealth	Healthcare
CB UN	Chubb	Financials	WFC UN	Wells Fargo	Financials
ANET UN	Arista Networks	Technology	ELV UN	Elevance Health	Healthcare
BRK/B UN	Berkshire Hathaway	Financials	SBUX UW	Starbucks	Consumer Discretionary
DHR UN	Danaher	Healthcare	SYK UN	Stryker	Healthcare
AMAT UW	Applied Materials	Technology (Semiconductors)	DVN UN	Devon Energy	Energy

Panel B

High Sentiment - persistent stocks			Low Sentiment - persistent stocks		
Ticker	**Company**	**Sector**	**Ticker**	**Company**	**Sector**
HWM UN	Howmet Aerospace	Industrials	DOC UN	Healthpeak Properties	Real Estate
A UN	Agilent Technologies	Healthcare	IQV UN	IQVIA Holdings	Healthcare
MU UW	Micron Technology	Technology (Semiconductors)	MDLZ UW	Mondelez	Consumer Staples
PGR UN	Progressive	Financials	CNC UN	Centene	Healthcare
NEE UN	NextEra Energy	Utilities	KIM UN	Kimco Realty	Real Estate
AZO UN	AutoZone	Consumer Discretionary	LULU UW	Lululemon	Consumer Discretionary
RTX UN	RTX Corporation	Industrials	HUM UN	Humana	Healthcare
BKNG UW	Booking Holdings	Consumer Discretionary	NEE UN	NextEra Energy	Utilities
FSLR UW	First Solar	Energy (Clean)	TGT UN	Target	Consumer Discretionary
WDC UW	Western Digital	Technology	LOW UN	Lowe's	Consumer Discretionary

Note: The table reports the twenty stocks most frequently classified in the high-sentiment (positive) and low-sentiment (negative) groups at the 10% cutoff over the full sample period (January 2024 – December 2025, 104 weekly rebalancing dates). Frequency is the number of weeks a stock appears in the respective sentiment group out of a maximum of 104. Only stocks with non-missing sentiment data for at least 50% of weeks are included. Sector classifications are based on the Global Industry Classification Standard (GICS).

C. View Vector Construction

Table C.1 - Spread and View Vector Descriptive Statistics

	Spread Std Dev	Q Std Dev	Volatility Reduction	Spread Avg	Q Avg
RQ1 – BL Sentiment Views portfolios					
10% cutoff	2.77%	0.48%	82.6%	5.47%	0.48%
20% cutoff	2.10%	0.41%	80.4%	4.28%	0.41%
30% cutoff	1.72%	0.29%	83.2%	3.58%	0.29%
RQ2 - Agreement portfolios					
Positive agreement	1.64%	0.21%	87.3%	2.62%	0.21%
Negative agreement	1.59%	0.35%	78.2%	1.78%	0.35%
RQ2 - Sign-based positive					
10% cutoff	1.91%	0.29%	84.9%	3.37%	0.29%
20% cutoff	1.55%	0.27%	82.7%	2.80%	0.27%
30% cutoff	1.45%	0.22%	85.1%	2.58%	0.22%
RQ2 - Sign-based negative					
10% cutoff	1.93%	0.31%	84.1%	2.68%	0.31%
20% cutoff	1.51%	0.30%	80.2%	2.45%	0.30%
30% cutoff	1.35%	0.25%	81.7%	2.39%	0.25%
All portfolios (average)	1.77%	0.31%	82.6%	-	-

Note: Spread Std Dev and Q Std Dev are the standard deviations of the realized weekly return spread and the exponentially weighted view vector Q respectively, computed over the full sample period. Vol Reduction is computed as 1 - $\sigma(Q)/\sigma(\text{Spread})$, expressed as a percentage. Spread Avg and Q Avg are the time-series means.

D. Model Inputs & Diagnostics

Table D.1 - Covariance Matrix Diagnostics

Diagnostic	Average	Minimum	Maximum	% PD
Condition number	309.6	180.97	419.82	-
Min eigenvalue	0.000640	0.000566	-	100%
Avg Variance	0.001623	-	-	-
Avg weekly Std Dev	4.02%	-	-	-
Positive definite	100%	-	-	-

Note: All diagnostics are computed at each weekly rebalancing date using a 52-week rolling window of excess returns and the Analytical Nonlinear Shrinkage estimator of Ledoit and Wolf (2020). The condition number is the ratio of the largest to the smallest eigenvalue. The minimum eigenvalue and positive definiteness (% PD) confirm the numerical stability of the matrices. Avg variance is the average of diagonal elements of the covariance matrix. Avg weekly Std Dev is the square root of the average variance.

Table D.2 - Portfolio Composition Statistics

Portfolio	Avg N stocks	HHI	Wtd avg beta	Wtd avg mcap ($bn)
RQ1 portfolios				
10% cutoff	352	0.177	0.86	787
20% cutoff	335	0.167	0.92	1,058
30% cutoff	355	0.151	0.87	1,193
RQ2 - Agreement				
Positive sentiment	308	0.178	0.99	1,646
Negative sentiment	347	0.053	0.63	871
RQ2 - Sign (10%)				
Positive sentiment	353	0.198	0.78	791
Negative sentiment	310	0.078	0.85	1,296
RQ2 - Sign (20%)				
Positive sentiment	334	0.164	0.85	1,036
Negative sentiment	316	0.087	0.83	1,465
RQ2 - Sign (30%)				
Positive sentiment	346	0.139	0.80	1,149
Negative sentiment	331	0.093	0.82	1,499

Note: Average portfolio characteristics computed over the full sample period. Statistics are based on the original view specification ($\lambda = 0.95$, 52-week window). Avg N stocks is the average number of assets receiving non-zero portfolio weight. HHI is the Herfindahl-Hirschman Index of portfolio concentration, computed as Σw_i^2 at each date and averaged over time. Weighted average beta is the portfolio-weight-average of individual stock market betas estimated from a full-sample OLS regression of weekly excess returns on S&P 500 excess returns. Weighted average market capitalization is reported in USD billions.

Table D.3 - Average and Median Correlations Between Posterior Expected Returns (μBL) and Prior Equilibrium Returns (Π)

Portfolio	Avg correlation	Median correlation
RQ1 portfolios		
10% cutoff	0.309	0.705
20% cutoff	0.309	0.680
30% cutoff	0.309	0.643
RQ2 - Agreement		
Positive sentiment	0.477	0.826
Negative sentiment	0.510	0.709
RQ2 - Sign (10%)		
Positive sentiment	0.202	0.454
Negative sentiment	0.473	0.755
RQ2 - Sign (20%)		
Positive sentiment	0.296	0.736
Negative sentiment	0.404	0.701
RQ2 - Sign (30%)		
Positive sentiment	0.304	0.681
Negative sentiment	0.410	0.714

Note: Correlations are computed weekly between the vector of posterior expected returns μ_{BL} and the vector of prior equilibrium returns Π, and averaged over the full sample period. A correlation below one indicates that the sentiment views move the posterior away from the prior. Lower values indicate greater divergence between posterior and prior.

E. Full Performance Results

Table E.1 - RQ1 Full Portfolio Performance Across All Tau Values

τ	Avg Ret	Std Dev	Sharpe	Sortino	Max DD	Turnover	IR
Benchmark	**0.32%**	**2.07%**	**1.123**	**1.500**	**-17.47%**	**-**	**-**
			10% cutoff				
0.01	0.30%	2.09%	1.034	1.427	-19.04%	32.5%	-0.391
0.05	0.26%	2.12%	0.872	1.252	-17.50%	70.2%	-0.423
0.10	0.27%	2.24%	0.853	1.346	-14.89%	78.3%	-0.260
0.20	0.29%	2.47%	0.858	1.473	-15.52%	80.2%	-0.104
1.00	0.33%	2.72%	0.875	1.690	-16.68%	80.4%	0.024
			20% cutoff				
0.01	0.32%	2.09%	1.104	1.539	-19.09%	30.8%	-0.039
0.05	0.31%	2.13%	1.057	1.464	-16.40%	65.2%	-0.067
0.10	0.29%	2.30%	0.918	1.285	-15.48%	71.6%	-0.129
0.20	0.25%	2.53%	0.708	0.994	-17.42%	73.1%	-0.258
1.00	0.23%	2.77%	0.587	0.825	-20.51%	72.9%	-0.290
			30% cutoff				
0.01	0.30%	2.07%	1.039	1.435	-18.22%	31.1%	-0.444
0.05	0.30%	2.12%	1.021	1.425	-15.97%	64.5%	-0.147
0.10	0.29%	2.24%	0.937	1.336	-14.78%	70.0%	-0.139
0.20	0.28%	2.41%	0.845	1.225	-15.42%	70.9%	-0.151
1.00	0.28%	2.57%	0.773	1.144	-17.54%	70.0%	-0.154

Note: The table reports full-sample performance of RQ1 ranking-based portfolios across all five τ. Average return and standard deviation are weekly figures. Sharpe and Sortino ratios are annualized by $\sqrt{52}$, assuming independent weekly returns. Maximum drawdown is the largest peak-to-trough cumulative loss over the full sample. Turnover is the average one-way weekly portfolio turnover. IR is the annualized information ratio. The parameter τ controls the weight assigned to views relative to the equilibrium prior in the Black-Litterman posterior. For each cutoff specification, the long leg comprises the high-sentiment group and the short leg the low-sentiment group, with weights proportional to market capitalization within each leg.

Table E.2 - RQ2 Full Portfolio Performance Across All Tau Values

Panel A – Agreement-based Portfolios

τ	Avg Ret	Std Dev	Sharpe	Sortino	Max DD	Turnover	IR
			Agreement positive				
0.01	0.33%	2.07%	1.146	1.602	-18.35%	23.0%	0.159
0.05	0.29%	2.20%	0.965	1.490	-17.83%	60.1%	-0.184
0.10	0.27%	2.31%	0.844	1.375	-19.51%	62.0%	-0.260
0.20	0.25%	2.40%	0.766	1.268	-22.86%	61.4%	-0.288
1.00	0.19%	2.46%	0.571	0.951	-26.22%	61.1%	-0.486
			Agreement negative				
0.01	0.33%	2.07%	1.157	1.578	-18.06%	14.9%	0.328
0.05	0.26%	2.01%	0.945	1.294	-18.91%	42.9%	-0.743
0.10	0.22%	1.89%	0.829	1.126	-17.71%	51.9%	-0.937
0.20	0.18%	1.82%	0.705	0.977	-16.40%	57.4%	-0.973
1.00	0.11%	1.78%	0.442	0.669	-17.16%	62.8%	-1.078

Panel B – Sign-based positive Portfolios

τ	Avg Ret	Std Dev	Sharpe	Sortino	Max DD	Turnover	IR
			Sign positive 10%				
0.01	0.32%	2.01%	1.152	1.575	-16.76%	29.3%	-0.036
0.05	0.29%	2.01%	1.027	1.519	-12.90%	79.4%	-0.209
0.10	0.32%	2.16%	1.052	1.689	-12.10%	79.2%	-0.033
0.20	0.36%	2.33%	1.110	1.853	-13.52%	76.8%	0.127
1.00	0.39%	2.57%	1.103	1.879	-15.02%	73.6%	0.216
			Sign positive 20%				
0.01	0.35%	2.03%	1.250	1.730	-16.92%	26.1%	0.662
0.05	0.34%	2.16%	1.130	1.623	-13.59%	68.7%	0.085
0.10	0.32%	2.32%	0.992	1.388	-13.93%	69.3%	-0.014
0.20	0.30%	2.45%	0.871	1.192	-15.88%	68.0%	-0.095
1.00	0.23%	2.59%	0.644	0.873	-18.78%	65.7%	-0.290
			Sign positive 30%				
0.01	0.34%	2.04%	1.205	1.650	-17.26%	25.7%	0.418
0.05	0.31%	2.12%	1.064	1.606	-15.72%	64.2%	-0.061
0.10	0.32%	2.18%	1.059	1.658	-13.89%	64.6%	-0.010
0.20	0.29%	2.24%	0.924	1.465	-15.14%	63.2%	-0.144
1.00	0.21%	2.31%	0.668	1.069	-16.51%	62.4%	-0.397

Panel C – Sign-based negative Portfolios

τ	Avg Ret	Std Dev	Sharpe	Sortino	Max DD	Turnover	IR
			Sign negative 10%				
0.01	0.32%	2.13%	1.096	1.518	-18.97%	16.1%	0.018
0.05	0.23%	2.18%	0.766	1.051	-20.62%	40.3%	-0.917
0.10	0.19%	2.21%	0.633	0.855	-20.10%	48.3%	-0.924
0.20	0.18%	2.22%	0.586	0.804	-18.68%	54.3%	-0.801
1.00	0.16%	2.26%	0.510	0.748	-16.36%	61.5%	-0.698
			Sign negative 20%				
0.01	0.32%	2.14%	1.073	1.484	-19.09%	19.6%	-0.082
0.05	0.27%	2.16%	0.902	1.260	-19.15%	41.4%	-0.482
0.10	0.24%	2.16%	0.815	1.173	-17.92%	47.6%	-0.518
0.20	0.20%	2.16%	0.682	0.979	-18.84%	52.5%	-0.638
1.00	0.15%	2.12%	0.523	0.800	-18.83%	58.7%	-0.759
			Sign negative 30%				
0.01	0.29%	2.11%	0.982	1.375	-19.28%	23.3%	-0.706
0.05	0.22%	2.12%	0.745	1.101	-20.31%	48.5%	-0.868
0.10	0.17%	2.13%	0.566	0.834	-21.69%	53.9%	-0.962
0.20	0.12%	2.13%	0.419	0.622	-23.32%	57.4%	-1.022
1.00	0.09%	2.09%	0.313	0.475	-24.43%	61.4%	-1.010

Note: The table reports full-sample performance of RQ2 sentiment-directional portfolios across all five τ vaues (January 2024 – December 2025, 104 weekly rebalancing dates). Average return and standard deviation are weekly figures. Sharpe and Sortino ratios are annualized by √52, assuming independent weekly returns. Maximum drawdown is the largest peak-to-trough cumulative loss over the full sample. Turnover is the average one-way weekly portfolio turnover. IR is the annualized information ratio relative to the S&P 500 benchmark. Agreement-based portfolios classify firms as positive (negative) only when both Bloomberg News and Twitter sentiment indicators are simultaneously positive (negative). Sign-based portfolios apply a percentile cutoff with an additional constraint requiring the composite sentiment score to be strictly positive (negative). For positive sentiment portfolios the long leg comprises the sentiment-classified group and the short leg the remaining universe; for negative sentiment portfolios this is reversed. Weights within each leg are proportional to market capitalization. Main body Table 5 reports results at $\tau = 0.01$ and $\tau = 1.0$ only.

F. Results: Subperiod Analysis

Table F.1 - Subperiod Analysis: RQ1 Portfolio Performance by Year, All τ Values

Portfolio		2024				2025			
	τ	***Sharpe***	**Sortino**	**IR**	**Turnover**	**Sharpe**	**Sortino**	**IR**	**Turnover**
Benchmark		1.691	2.558	-	-	0.742	0.965	-	-
Ranking-Based (10% cutoff)									
	0.01	1.621	2.253	-0.144	36.7%	0.630	0.895	-0.655	28.4%
	0.05	1.409	1.885	0.006	70.2%	0.373	0.561	-0.987	70.1%
	0.10	1.443	2.170	0.329	76.2%	0.181	0.289	-0.977	80.3%
	0.20	1.573	2.697	0.668	78.3%	-0.045	-0.074	-1.057	82.2%
	1.00	1.697	3.458	0.915	78.3%	-0.210	-0.365	-1.111	82.5%
Ranking-Based (20% cutoff)									
	0.01	1.680	2.441	-0.077	32.5%	0.729	1.038	-0.001	29.2%
	0.05	1.489	1.925	0.123	63.6%	0.657	0.951	-0.310	66.9%
	0.10	1.336	1.726	0.210	67.8%	0.443	0.660	-0.539	75.4%
	0.20	1.244	1.572	0.275	68.9%	0.036	0.055	-0.922	77.2%
	1.00	1.253	1.567	0.391	68.7%	-0.284	-0.443	-1.142	77.0%
Ranking-Based (30% cutoff)									
	0.01	1.491	2.156	-0.788	34.0%	0.732	1.013	-0.079	28.3%
	0.05	1.289	1.737	-0.191	63.6%	0.765	1.089	-0.094	65.4%
	0.10	1.277	1.714	0.091	66.9%	0.555	0.822	-0.401	73.0%
	0.20	1.324	1.798	0.292	67.0%	0.271	0.408	-0.660	74.8%
	1.00	1.407	1.984	0.476	65.5%	-0.020	-0.030	-0.883	74.5%

Note: Sharpe and Sortino ratios annualized by √52. Max DD is the largest peak-to-trough loss within the subperiod. Turnover is average one-way weekly portfolio turnover. IR is the annualized information ratio relative to the S&P 500 benchmark. The first subperiod (2024) was characterized by a sustained bull market; the second (2025) included the sharp market correction of April 2025 followed by a rapid recovery. See Table F.2 for the corresponding RQ2 subperiod results.

Table F.2 - Subperiod Analysis: RQ2 Portfolio Performance by Year, All τ Values

Panel A – Agreement Sentiment Portfolios

Portfolio		2024				2025		
τ	***Sharpe***	**Sortino**	**IR**	**Turnover**	**Sharpe**	**Sortino**	**IR**	**Turnover**
Bench-mark	1.691	2.558	-	-	0.742	0.965	-	-
				Agreement Positive				
0.01	1.726	2.666	0.417	25.2%	0.735	1.007	-0.145	20.9%
0.05	1.603	2.821	0.405	60.3%	0.382	0.547	-0.944	59.9%
0.10	1.444	2.603	0.322	61.3%	0.222	0.327	-0.987	62.6%
0.20	1.375	2.489	0.304	59.6%	0.102	0.152	-1.008	63.1%
1.00	1.340	2.499	0.278	58.8%	-0.268	-0.396	-1.348	63.3%
				Agreement Negative				
0.01	1.639	2.383	-0.381	14.9%	0.835	1.139	0.928	15.0%
0.05	1.242	1.744	-1.187	43.0%	0.725	0.992	-0.253	42.9%
0.10	0.950	1.375	-1.544	51.3%	0.723	0.949	-0.369	52.5%
0.20	0.648	0.932	-1.698	56.5%	0.753	1.011	-0.336	58.4%
1.00	0.338	0.503	-1.753	60.8%	0.551	0.850	-0.535	64.8%

Panel B – Signed-baes Sentiment Portfolios (10%)

Portfolio	2024				2025			
τ	***Sharpe***	**Sortino**	**IR**	**Turnover**	**Sharpe**	**Sortino**	**IR**	**Turnover**
Bench-mark	1.691	2.558	-	-	0.742	0.965	-	-
				Positive				
0.01	1.731	2.559	-0.022	32.0%	0.760	1.047	-0.049	26.5%
0.05	1.723	2.701	0.200	77.9%	0.459	0.669	-0.606	80.9%
0.10	1.876	3.823	0.645	76.3%	0.261	0.372	-0.708	82.1%
0.20	2.075	4.583	1.026	73.3%	0.085	0.117	-0.799	80.2%
1.00	2.224	5.575	1.302	70.9%	-0.170	-0.225	-0.996	76.3%
				Negative				
0.01	1.639	2.322	0.094	16.9%	0.720	1.015	-0.060	15.2%
0.05	1.301	1.657	-0.389	38.3%	0.326	0.474	-1.586	42.2%
0.10	1.121	1.384	-0.465	43.6%	0.185	0.265	-1.452	52.9%
0.20	1.019	1.257	-0.489	46.5%	0.154	0.229	-1.094	61.9%
1.00	0.895	1.119	-0.572	51.1%	0.122	0.210	-0.800	71.8%

Panel C – Signed-baes Sentiment Portfolios (20%)

Portfolio	2024				2025			
τ	***Sharpe***	**Sortino**	**IR**	**Turnover**	**Sharpe**	**Sortino**	**IR**	**Turnover**
Bench-mark	1.691	2.558	-	-	0.742	0.965	-	-
				Positive				
0.01	1.805	2.532	0.504	28.2%	0.869	1.229	0.853	24.1%
0.05	1.597	2.228	0.348	67.9%	0.679	0.994	-0.255	69.4%
0.10	1.446	2.068	0.322	67.4%	0.504	0.676	-0.400	71.2%
0.20	1.340	1.913	0.306	65.6%	0.326	0.415	-0.560	70.3%
1.00	1.194	1.665	0.212	64.3%	-0.021	-0.026	-0.867	67.1%
				Negative				
0.01	1.527	2.227	-0.599	19.9%	0.776	1.084	0.405	19.3%
0.05	1.307	1.777	-0.460	38.7%	0.583	0.836	-0.501	44.1%
0.10	1.276	1.781	-0.242	42.6%	0.401	0.587	-0.798	52.5%
0.20	1.227	1.736	-0.223	44.6%	0.162	0.233	-1.023	60.4%
1.00	1.120	1.602	-0.355	47.8%	-0.064	-0.103	-1.095	69.4%

Panel D – Sentiment-based Portfolios (30% cutoff)

Portfolio	2024				2025			
τ	***Sharpe***	**Sortino**	**IR**	**Turnover**	**Sharpe**	**Sortino**	**IR**	**Turnover**
Bench-mark	1.691	2.558	-	-	0.742	0.965	-	-
				Positive				
0.01	1.738	2.469	0.436	27.8%	0.820	1.124	0.394	23.7%
0.05	1.515	2.380	0.188	63.4%	0.641	0.932	-0.339	64.9%
0.10	1.503	2.449	0.272	62.7%	0.609	0.913	-0.286	66.4%
0.20	1.420	2.292	0.231	61.8%	0.396	0.607	-0.506	64.7%
1.00	1.328	2.130	0.167	61.2%	-0.051	-0.081	-0.924	63.6%
				Negative				
0.01	1.439	2.183	-1.074	25.0%	0.681	0.953	-0.310	21.7%
0.05	1.155	1.751	-0.789	47.7%	0.427	0.630	-0.950	49.2%
0.10	1.150	1.708	-0.504	50.9%	0.067	0.098	-1.450	56.8%
0.20	1.289	1.933	-0.233	51.8%	-0.335	-0.497	-1.793	62.8%
1.00	1.370	2.011	-0.148	53.7%	-0.602	-0.932	-1.782	69.0%

Note: Sharpe and Sortino ratios annualized by √52. Max DD is the largest peak-to-trough loss within the subperiod. Turnover is average one-way weekly portfolio turnover. IR is the annualized information ratio relative to the S&P 500 benchmark. Agreement-based portfolios require both Bloomberg News and Twitter indicators to be simultaneously positive (negative). Sign-based portfolios apply a percentile cutoff requiring the composite score to be strictly positive (negative). The first subperiod (2024) was characterized by a sustained bull market; the second (2025) included the sharp market correction of April 2025 followed by a rapid recovery. See Table F.1 for the corresponding RQ1 subperiod results.

G. Results: Robustness

Table G.1 - Sharpe and Sortino ratios for Alternative View Construction

	Sharpe Ratio (annualized)						Sortino Ratio (annualized)					
	$\tau = 1.0$			$\tau = 0.01$			$\tau = 1.0$			$\tau = 0.01$		
	Original	Short-window	Delta-sent	Original	Short-window	Delta-sent	Original	Short-window	Delta-sent	Original	Short-window	Delta-sent
Benchmark	1.123	1.123	1.123	1.123	1.123	1.123	1.500	1.500	1.500	1.500	1.500	1.500
RQ1 Portfolios												
10% cutoff	0.875	0.870	0.879	1.034	1.037	0.845	1.690	1.664	1.670	1.427	1.415	1.156
20% cutoff	0.587	0.699	0.529	1.104	1.095	1.054	0.825	1.016	0.749	1.539	1.519	1.481
30% cutoff	0.773	0.838	0.715	1.039	1.076	1.046	1.144	1.232	1.070	1.435	1.477	1.496
RQ2 – Agreement-based Portfolios												
Positive	0.571	0.606	0.450	1.146	1.193	0.998	0.951	0.980	0.724	1.602	1.683	1.425
Negative	0.442	0.617	0.308	1.157	1.130	1.013	0.669	0.910	0.479	1.578	1.529	1.372
RQ2 – Positive Sign-based Portfolios												
10% cutoff	1.103	1.072	1.089	1.152	1.170	0.923	1.879	1.793	1.901	1.575	1.582	1.216
20% cutoff	0.644	0.794	0.581	1.250	1.327	1.172	0.873	1.112	0.792	1.730	1.852	1.625
30% cutoff	0.668	0.857	0.573	1.205	1.323	1.058	1.069	1.376	0.920	1.650	1.817	1.469
RQ2 – Negative Sign-based Portfolios												
10% cutoff	0.510	0.610	0.452	1.096	1.106	0.817	0.748	0.892	0.687	1.518	1.524	1.148
20% cutoff	0.523	0.592	0.514	1.073	1.063	0.995	0.800	0.885	0.817	1.484	1.462	1.427
30% cutoff	0.313	0.305	0.277	0.982	0.920	0.916	0.475	0.456	0.425	1.375	1.284	1.378

Note: All results are full-sample (January 2024 – December 2025). Sharpe and Sortino ratios annualized by √52. Turnover is average one-way weekly portfolio turnover. Original specification: 52-week exponentially weighted window, λ = 0.95. Short window specification: 4-week exponentially weighted window, λ = 0.50. Delta-sentiment specification: week-on-week change in the sentiment-based return spread as the view signal Q. τ controls the weight assigned to views relative to the equilibrium prior.

Table G.2 - View Vector Predictability Regressions

Portfolio	α	$\beta(Q)$	t-stat	p-value	R^2
		RQ1 portfolios			
10% cutoff	0.097	-0.738	-1.505	0.132	0.016
20% cutoff	0.061	-0.395	-0.796	0.426	0.006
30% cutoff	0.067	-0.850	-1.757	0.079[5]	0.020
		RQ2 - Agreement			
Positive sentiment	0.056	-1.095	-1.713	0.087[5]	0.019
Negative sentiment	0.016	0.174	0.503	0.615	0.001
		RQ2 - Sign (10%)			
Positive sentiment	0.065	-0.901	-1.619	0.106	0.019
Negative sentiment	0.045	-0.658	-1.299	0.194	0.011
		RQ2 - Sign (20%)			
Positive sentiment	0.049	-0.713	-1.393	0.164	0.015
Negative sentiment	0.030	-0.226	-0.493	0.622	0.002
		RQ2 - Sign (30%)			
Positive sentiment	0.056	-1.131	-2.294	0.022[6]	0.028
Negative sentiment	0.032	-0.342	-0.779	0.436	0.004

Note: The table reports OLS estimates of $rt + 1 = \alpha + \beta \cdot Qt + \varepsilon t + 1$, where $rt + 1$ is the realized return differential of the portfolio in week $t + 1$ and Q_t is the view vector constructed at the end of week t. t-statistics are computed using Newey-West standard errors with 4 lags to account for potential autocorrelation in residuals.

[5] denotes significance at the 10% level.

[6] denotes significance at the 5% level.

H. Transaction Costs Sensitivity

Table H.1 - Transaction Cost Sensitivity Analysis (Weekly Average Returns)

τ	Avg Turnover	Pre-Cost Returns	After-Cost Returns			Benchmark
			5bps	10bps	20bps	
RQ1 10% cutoff						
0.01	32.5%	0.30%	0.27%	0.24%	0.17%	0.32%
0.05	70.2%	0.26%	0.20%	0.15%	0.04%	0.32%
0.10	78.3%	0.27%	0.19%	0.11%	-0.04%	0.32%
0.20	80.2%	0.29%	0.21%	0.13%	-0.03%	0.32%
1.00	80.4%	0.33%	0.25%	0.17%	0.01%	0.32%
RQ1 20% cutoff						
0.01	30.8%	0.32%	0.30%	0.27%	0.22%	0.32%
0.05	65.2%	0.31%	0.24%	0.16%	0.02%	0.32%
0.10	71.6%	0.29%	0.21%	0.13%	-0.02%	0.32%
0.20	73.1%	0.25%	0.17%	0.09%	-0.05%	0.32%
1.00	72.9%	0.23%	0.15%	0.08%	-0.07%	0.32%
RQ1 30% cutoff						
0.01	31.1%	0.30%	0.27%	0.24%	0.18%	0.32%
0.05	64.5%	0.30%	0.23%	0.17%	0.04%	0.32%
0.10	70.0%	0.29%	0.22%	0.14%	0.00%	0.32%
0.20	70.9%	0.28%	0.21%	0.13%	-0.01%	0.32%
1.00	70.0%	0.28%	0.21%	0.14%	-0.01%	0.32%

Note: Returns are reported as weekly averages. Transaction costs are computed as: cost = avg_turnover × 2 × bps, where the factor of 2 accounts for the round-trip nature of each rebalancing transaction. The benchmark return of 0.32% per week is reported for reference and assumes zero transaction costs.

References

Aggarwal, D. (2022). The role of sentiments in equity stock market behavior: A critical review. *Qualitative Research in Financial Markets*, *14*(2), 270–288.

Allen, D. E., McAleer, M., & Singh, A. K. (2019). Daily market news sentiment and stock prices. *Applied Economics*, *51*(30), 3212–3235.

Artzner, P., Delbaen, F., Eber, J. M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, *9*(3), 203–228.

Baker, M., & Wurgler, J. (2006). Investor sentiment and the cross-section of stock returns. *The Journal of Finance*, *61*(4), 1645–1680.

Baker, M., & Wurgler, J. (2007). Investor sentiment in the stock market. *Journal of Economic Perspectives*, *21*(2), 129–151.

Bandopadhyaya, A., & Jones, A. L. (2006). Measuring investor sentiment in equity markets. *Journal of Asset Management*, *7*(3–4), 208–215.

Barberis, N., Greenwood, R., Jin, L., & Shleifer, A. (2018). Extrapolation and bubbles. *Journal of Financial Economics*, *129*(2), 203–227.

Barberis, N., Shleifer, A., & Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, *49*(3), 307–343.

Barua, R., & Sharma, A. K. (2023). Fear and greed indicator for ETF return prediction and portfolio optimization using a machine learning-based Black–Litterman model. *Finance Research Letters*, *58*, 104515.

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, *5*(4), 323–370.

Bessler, W., Opfer, H., & Wolff, D. (2017). Multi-asset portfolio optimization and out-of-sample performance: An evaluation of Black–Litterman, mean-variance, and naïve diversification approaches. *The European Journal of Finance*, *23*(1), 1–30.

Black, F., & Litterman, R. (1992). Global portfolio optimization. *Financial Analysts Journal*, *48*(5), 28–43.

Blamont, D., & Firoozye, N. (2003). Black Litterman Asset Allocation. Deutsche Bank Global Markets Research, Fixed Income Weekly.

Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, *2*(1), 1–8.

Brandt, M. W. (2010). Portfolio choice problems. In Y. Aït-Sahalia & L. P. Hansen (Eds.), *Handbook of Financial Econometrics: Tools and Techniques* (Vol. 1, pp. 269–336). Elsevier.

Britten-Jones, M. (1999). The sampling error in estimates of mean-variance efficient portfolio weights. *The Journal of Finance*, *54*(2), 655–671.

Brown, G. W., & Cliff, M. T. (2005). Investor sentiment and asset valuation. *The Journal of Business*, *78*(2), 405–440.

Campbell, J. Y., & Viceira, L. M. (2002). *Strategic asset allocation: Portfolio choice for long-term investors*. Oxford University Press.

Chopra, V. K., & Ziemba, W. T. (1993). The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management*, *19*(2), 6–11.

Choueifaty, Y., & Coignard, Y. (2008). Towards maximum diversification. *The Journal of Portfolio Management*, *35*(1), 40–51.

Clarke, R., de Silva, H., & Thorley, S. (2011). Minimum-variance portfolio composition. *The Journal of Portfolio Management*, *37*(Winter), 31–45.

Colasanto, F., Grilli, L., Santoro, D., & Villani, G. (2022). BERT's sentiment score for portfolio optimisation: A fine-tuned view in Black and Litterman model. *Neural Computing and Applications*, *34*(20), 17507–17521.

Consoli, S., Barbaglia, L., & Manzan, S. (2022). Fine-grained, aspect-based sentiment analysis on economic and financial lexicon. *Knowledge-Based Systems*, *247*, 108781.

Cookson, J. A., & Niessner, M. (2020). Why don't we agree? Evidence from a social network of investors. *The Journal of Finance*, *75*(1), 173–228.

Da, Z., Engelberg, J., & Gao, P. (2011). In search of attention. *The Journal of Finance*, *66*(5), 1461–1499.

De Bondt, W. F. M., & Thaler, R. (1985). Does the stock market overreact? *The Journal of Finance*, *40*(3), 793–805.

DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The Review of Financial Studies*, *22*(5), 1915–1953.

Dennis, P., & Mayhew, S. (2002). Risk-neutral skewness: Evidence from stock options. *The Journal of Financial and Quantitative Analysis*, *37*(3), 471–493.

Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *The Journal of Finance*, *47*(2), 427–465.

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, *33*(1), 3–56.

Feng, Y., Zhang, Y., & Wang, Y. (2024). Sustainable portfolio construction via machine learning: ESG, SDG and sentiment. *European Financial Management*, *31*(3), 1148–1169.

Fisher, K. L., & Statman, M. (2000). Cognitive biases in market forecasts. *The Journal of Portfolio Management*, *27*(1), 1–10.

Frazzini, A., Israel, R., & Moskowitz, T. J. (2014). *Trading costs of asset pricing anomalies* (Chicago Booth Research Paper No. 14-05). SSRN.

Giacometti, R., Bertocchi, M., Rachev, S. T., & Fabozzi, F. (2007). Stable distributions in the Black–Litterman approach to asset allocation. *Quantitative Finance*, *7*(4), 423–433.

Grinold, R. C., & Kahn, R. N. (1999). *Active portfolio management: A quantitative approach for producing superior returns and controlling risk* (2nd ed.). McGraw-Hill.

Harvey, C. R., & Siddique, A. (2000). Conditional skewness in asset pricing tests. *The Journal of Finance*, *55*(3), 1263–1295.

He, G., & Litterman, R. (1999). *The intuition behind Black-Litterman model portfolios*. Goldman Sachs Investment Management Research. Goldman Sachs & Company.

Hung, M.-C., Hsia, P.-H., Kuang, X.-J., & Lin, S.-K. (2024). Intelligent portfolio construction via news sentiment analysis. *International Review of Economics and Finance*, *89*, 605-617.

Idzorek, T. M. (2005). *A step-by-step guide to the Black-Litterman model: Incorporating user-specified confidence levels*. Ibbotson Associates.

Jiang, H., He, P., Chen, W., Liu, X., Gao, J., & Zhao, T. (2020). SMART: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 2177–2190).

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.

Kara, M., Ulucan, A., & Atici, K. B. (2019). A hybrid approach for generating investor views in Black–Litterman model. *Expert Systems with Applications*, *128*, 256–270.

Ke, Z. T., Kelly, B., & Xiu, D. (2023). Predicting returns with text data. *Journal of Business & Economic Statistics*, *41*(4), 1019–1035.

Klein, R. W., & Bawa, V. S. (1976). The effect of estimation risk on optimal portfolio choice. *Journal of Financial Economics*, *3*(3), 215–231.

Kolm, P. N., Ritter, G., & Simonian, J. (2021). Black-Litterman and beyond: The Bayesian paradigm in investment management. *Journal of Portfolio Management*, *47*(5), 11–22.

Konno, H., & Yamazaki, H. (1991). Mean-absolute deviation portfolio optimization model and its applications to Tokyo stock market. *Management Science*, *37*(5), 519–531.

Konstantinidis, T., Golan, M., Xu, Y. L., Constantinides, T. G., & Mandic, D. P. (2023). Text mining for sentiment analysis in bond portfolio construction. In *2023 24th International Conference on Digital Signal Processing (DSP)* (pp. 1–5). IEEE.

Koratamaddi, P., Wadhwani, K., Gupta, M., & Sanjeevi, S. G. (2021). Market sentiment-aware deep reinforcement learning approach for stock portfolio allocation. *Engineering Science and Technology, an International Journal*, *24*(4), 848–859.

Kraus, A., & Litzenberger, R. H. (1976). Skewness preference and the valuation of risk assets. *The Journal of Finance*, *31*(4), 1085–1100.

Kresta, A., Xiong, J., & Maidiya, B. (2024). Sentiment and stock characteristics: Comprehensive study of individual investor influence on returns, volatility, and trading volumes. *Business Systems Research*, *15*(2), 67–82.

Kucukaslan, B., & Tas, O. (2024). Twitter sentiment analysis for optimal portfolio construction. *PressAcademia Procedia*, *20*, 17–23.

Kumar, M. S., & Persaud, A. (2002). Pure contagion and investors' shifting risk appetite: Analytical issues and empirical evidence. *International Finance*, *5*(3), 401–436.

Ledoit, O., & Wolf, M. (2020). Analytical nonlinear shrinkage of large-dimensional covariance matrices. *Annals of Statistics*, *48*(5), 3043–3065.

Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics*, *47*(1), 13–37.

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, *66*(1), 35–65.

Maillard, S., Roncalli, T., & Teïletche, J. (2010). The properties of equally weighted risk contribution portfolios. *The Journal of Portfolio Management*, *36*(4), 60–70.

Malandri, L., Xing, F. Z., Orsenigo, C., Vercellis, C., & Cambria, E. (2018). Public mood-driven asset allocation: The importance of financial sentiment in portfolio management. *Cognitive Computation*, *10*(6), 1167–1176.

Mantshimuli, L. A., & Mwamba, J. W. M. (2025). Large language models in financial sentiment analysis: Evidence from the South African stock market. *Investment Management and Financial Innovations*, *22*(3), 214–227.

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, *7*(1), 77–91.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, *51*(3), 247–257.

Meucci, A. (2008). Fully flexible views: Theory and practice. *Risk*, *21*(10), 97–102.

Michaud, R. O. (1989). The Markowitz optimization enigma: Is "optimized" optimal? *Financial Analysts Journal*, *45*(1), 31–42.

Molk, P., & Partnoy, F. (2019). Institutional investors as short sellers? *Boston University Law Review*, *99*(3), 837–865.

Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica*, *34*(4), 768–783.

Muthivhi, M., & van Zyl, T. L. (2021). Fusion of sentiment and asset price predictions for portfolio optimization. In *2021 IEEE 24th International Conference on Information Fusion (FUSION)* (pp. 1–8).

Novy-Marx, R., & Velikov, M. (2016). A taxonomy of anomalies and their trading costs. *Review of Financial Studies*, *29*(1), 104–147.

Reis, P., Serra, A. P., & Gama, J. (2026). A deep learning framework for forecasting medium-term covariance in multi-asset portfolios. *Journal of Forecasting*, 1–32.

Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, *2*(3), 21–41.

Roll, R. (1977). A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory. *Journal of Financial Economics*, *4*(2), 129–176.

Rosenman, R., Tennekoon, V., & Hill, L. G. (2011). Measuring bias in self-reported data. *International Journal of Behavioural and Healthcare Research*, *2*(4), 320–332.

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, *13*, 341–360.

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, *1*(5), 206–215.

Satchell, S., & Scowcroft, A. (2000). A demystification of the Black–Litterman model: Managing quantitative and traditional portfolio construction. *Journal of Asset Management*, *1*(2), 138–150.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, *19*(3), 425–442.

Sharpe, W. F. (1966). Mutual fund performance. *The Journal of Business*, *39*, 119–138.

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *The American Economic Review*, *71*(3), 421–436.

Shleifer, A., & Vishny, R. W. (1997). The limits of arbitrage. *The Journal of Finance*, *52*(1), 35–55.

Sortino, F. A., & Price, L. N. (1994). Performance measurement in a downside risk framework. *The Journal of Investing*, *3*(3), 59–64.

Stambaugh, R. F., Yu, J., & Yuan, Y. (2012). The short of it: Investor sentiment and anomalies. *Journal of Financial Economics*, *104*(2), 288–302.

Sun, L., Najand, M., & Shen, J. (2016). Stock return predictability and investor sentiment: A high-frequency perspective. *Journal of Banking & Finance*, *73*, 147–164.

Taheripour, E., Sadjadi, S. J., & Amiri, B. (2025). A novel approach to portfolio construction: An application of FinBERT sentiment analysis and credibilistic CVaR criterion. *IEEE Access*, *13*, 76775–76795.

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance*, *62*(3), 1139–1168.

Tobin, J. (1958). Liquidity preference as behavior towards risk. *The Review of Economic Studies*, *25*(2), 65–86.

Todd, A., Bowden, J., & Moshfeghi, Y. (2024). Text-based sentiment analysis in finance: Synthesising the existing literature and exploring future directions. *Intelligent Systems in Accounting, Finance and Management*, *31*(1), e1549.

Treynor, J. L. (1962). Toward a theory of market value of risky assets. In R. A. Korajczyk (Ed.), *Asset pricing and portfolio performance: Models, strategy and performance metrics* (pp. 15–22). Risk Books.

Whaley, R. E. (2000). The investor fear gauge. *The Journal of Portfolio Management*, *26*(3), 12–17.

Yu, J., & Yuan, Y. (2011). Investor sentiment and the mean-variance relation. *Journal of Financial Economics*, *100*(2), 367–381.