

LEVERAGING LLMs FOR TOP-DOWN SECTOR ALLOCATION IN AUTOMATED TRADING

Ryan Quek Wei Heng¹ Edoardo Vittori² Keane Ong^{1,3} Rui Mao⁴

Erik Cambria⁴ Gianmarco Mengaldo^{1,3,5,6}

¹College of Design and Engineering, National University of Singapore

²CVA Management and A.I. Investments, IMI Corporate and Investment Banking, Intesa Sanpaolo

³Asian Institute of Digital Finance, National University of Singapore

⁴College of Computing and Data Science, Nanyang Technological University

⁵Sustainable and Green Finance Institute, National University of Singapore

⁶Honorary Research Fellow, Imperial College London

ryanquekweiheng@u.nus.edu

edoardo.vittori@intesasanpaolo.com

keane.ongweiyang@u.nus.edu

rui.mao@ntu.edu.sg

cambria@ntu.edu.sg

mpegim@nus.edu.sg

Preprint accepted to ICLR Workshop Advances in Financial AI: Opportunities, Innovations, and Responsible AI on March 5, 2025

ABSTRACT

This paper introduces a methodology leveraging Large Language Models (LLMs) for sector-level portfolio allocation through systematic analysis of macroeconomic conditions and market sentiment. Our framework emphasizes top-down sector allocation by processing multiple data streams simultaneously, including policy documents, economic indicators, and sentiment patterns. Empirical results demonstrate superior risk-adjusted returns compared to traditional cross-momentum strategies, achieving a Sharpe ratio of 2.51 and portfolio return of 8.79% versus -0.61 and -1.39% respectively. These results suggest that LLM-based systematic macro analysis presents a viable approach for enhancing automated portfolio allocation decisions at the sector level.

1 INTRODUCTION

The rapid advancement of Large Language Models (LLMs) has catalyzed significant innovations in algorithmic trading and investment management, with numerous studies exploring their potential for stock selection and portfolio optimization (Ding et al., 2024). However, while existing research has predominantly focused on bottom-up security selection using company-specific metrics (see e.g., Yu et al. (2024); Zhang et al. (2024); Wang et al. (2024)), there remains a crucial gap in understanding how LLMs can enhance top-down investment strategies, particularly in the context of sector allocation within a portfolio.

Macroeconomic factors fundamentally shape investment decision-making, influencing stock valuations through both economic channels and market sentiment (Jareño and Negrut, 2016; Jabeen et al., 2022; Ma et al., 2024). A top-down investment approach, based on macroeconomic analysis, enables investors to anticipate broad market trends, optimize portfolio allocation across economic cycles, and identify systematic risks before they manifest in individual equity prices. These macroeconomic forces generate heterogeneous sectoral responses; for instance, inflationary pressures typically constrain manufacturing growth while benefiting services and agricultural sectors (Chaudhry et al., 2013). Moreover, hawkish Federal Reserve communications not only elevate borrowing costs and alter corporate decision-making but also intensify investor concerns about financially vulnerable firms, amplifying negative sentiment (Cieslak and McMahon, 2023; Xing et al., 2020).

This sentiment-driven reaction often precipitates market movements that exceed what fundamentals alone would otherwise suggest. Key macroeconomic variables—including monetary policy, inflation dynamics, and economic growth trajectories—propagate through the financial system, influencing both firms’ cost of capital and aggregate consumption patterns. Given that the task of constructing a portfolio based on sector allocation is one that requires piecing together information from different sources, this paper seeks to address the following **Research Question** – How can LLMs be leveraged to enhance top-down sector allocation strategies by integrating macroeconomic analysis with market sentiment for automated portfolio construction?

To respond to this question, we introduce a novel LLM-based methodology that emphasizes top-down sector allocation through integrated analysis of macroeconomic and market sentiment data. While existing approaches incorporate various data sources including sentiment and fundamentals, they typically focus on bottom-up analysis of individual securities. Our framework leverages LLMs to systematically process and synthesize multiple data streams simultaneously (including policy documents, economic indicators, and sentiment patterns), enabling dynamic adjustment of sector allocations based on market conditions. By automating the extraction and interpretation of these macro financial relationships, our framework enhances the responsiveness of sector allocation strategies through a predominantly top-down lens. This provides a more systematic approach to capturing sentiment-driven price movements, offering new insights into sector-level portfolio construction that complement traditional security-level analysis.

Results obtained through backtesting indicate that our sector allocation strategy significantly outperforms traditional cross-momentum approaches. Over the testing period, our methodology achieved a positive return of 8.79% and a Sharpe ratio of 2.51, compared to a loss of 1.39% and a Sharpe ratio of -0.61 for the cross-momentum strategy. These results suggest that our top-down LLM-based approach effectively captures sector-level opportunities while maintaining strong risk-adjusted performance, demonstrating the potential value of incorporating systematic macro and sentiment analysis in portfolio allocation decisions.

2 RELATED WORKS

In this section, we cover related literature on the following four topics: cross-sectional momentum investment, top-down investment, usage of LLMs as trading agents, and finance specific Aspect based Sentiment Analysis.

2.1 CROSS-SECTIONAL MOMENTUM INVESTMENT

Traditional cross-sectional momentum strategies rank assets based on their relative performance, taking long positions in top performers and short positions in underperformers. The classical approach introduced by Jegadeesh and Titman (1993) ranks stocks based on past returns, while modern variations incorporate volatility normalization and machine learning models trained to minimize mean squared error before ranking assets to construct portfolios.

Learning to Rank (LTR) algorithms further improve upon traditional approaches by explicitly optimizing for ranking quality rather than prediction accuracy (Poh et al., 2020). These algorithms, which include pairwise methods like RankNet (Burges et al., 2005) and LambdaMART (Wu et al., 2010) and listwise methods like ListNet (Cao et al., 2007), learn the relative ordering between instruments directly. When applied to cross-sectional momentum strategies, LTR methods demonstrate significant improvements over traditional approaches, achieving approximately three times higher Sharpe ratios.

2.2 TOP-DOWN INVESTMENT

The relationship between macroeconomic indicators and stock markets has been widely studied, demonstrating their significant impact on equity values.

Early research on Sri Lanka’s stock market (Gunasekarage et al., 2004) found that interest rates (measured by the treasury bill rate) and inflation (measured by the consumer price index) strongly influence stock prices, with the treasury bill rate exerting the most significant effect over a 17-year period (1985–2001).

More recent studies leverage macroeconomic regime modeling to enhance equity factor investing strategies (Nuriyev et al., 2024). Techniques such as Hidden Markov Models and meta-correlation-based clustering have outperformed traditional strategies by dynamically adjusting portfolio weights based on economic conditions, favoring quality factors during crises and momentum factors during growth phases. This further underscores the critical role of macroeconomic indicators in shaping stock market behavior.

2.3 LLM TRADING AGENTS

Previous works (Ding et al., 2024) on LLM-based trading agents can be broadly categorized into (i) LLM as a Trader and (ii) LLM as an Alpha Miner. This literature review focuses on the former, examining works in News-Driven and Reflection-Driven LLMs as Traders.

2.3.1 LLM AS A TRADER - NEWS DRIVEN

News-driven LLMs integrate stock news and macroeconomic updates into their prompt context before instructing the LLMs to predict stock price movements.

Studies (Lopez-Lira and Tang, 2023; Wu, 2024) show that both proprietary and open-source LLMs, even without specialized financial training, can effectively analyze news sentiment for stock market prediction. Backtesting a simple long-short strategy based on sentiment scores from news headlines has demonstrated promising results.

A more advanced approach Fatouros et al. (2024) enhances the approach by summarizing and refining news data while reasoning about its relationship with stock price movements. A memory module stores these summaries, which are later retrieved as contextual "recommendations" for trading decisions. By leveraging progressive daily news summaries, macroeconomic insights, and stock price momentum summaries, the MS-Top10-Cap-GPT strategy outperformed the S&P 100 index by approximately 30%.

2.3.2 LLM AS A TRADER - REFLECTION DRIVEN

Apart from memories, reflections can be used in LLM decision making. In this context, Reflections (Park et al., 2023) are defined as high-level knowledge and insights progressively aggregated from raw memories and observations. The inclusion of memory and reflection in LLM-based algorithms offers significant benefits such as mitigating the risk of hallucinations (Ji et al., 2023) and obtaining high-level understanding of the environment (Park et al., 2023).

FinMem (Yu et al., 2024) introduces a trading agent with layered memorization and characteristics. The raw inputs, such as daily news and financial reports, are summarized into memories. Upon the arrival of new observations, the relevant memories are retrieved and integrated with these observations to produce reflections. Both memories and reflections are stored in a layered memory bucket. During the trading phase, these memories and reflections are retrieved and utilized by the decision-making module to generate the final trading decisions. The retrieval method considers the recency, relevancy, and importance of the information.

Similarly, FinAgent (Zhang et al., 2024) extends this approach with a multimodal module that processes numeric, text, and image data. It enhances decision making by incorporating technical indicators like MACD and RSI, along with analyst guidance, demonstrating superior backtesting performance over FinMem.

2.4 FINANCE SPECIFIC ASPECT BASED SENTIMENT ANALYSIS

Traditional sentiment analysis in finance has historically focused on broad positive/negative classifications of text towards stocks, but markets and financial decisions require more nuanced understanding. Aspect based Sentiment Analysis (ABSA) has emerged as a crucial advancement by enabling the extraction and analysis of sentiment toward specific aspects or features within financial text Du et al. (2024); Malandri et al. (2018); Ma et al. (2023). Previous works of finance-specific ABSA (Ong et al., 2023) highlight the separation of sentiment across different financial aspects for explainable sentiment, allowing for assessment of sentiment towards specific risk factors or market aspects.

On the other hand, LLMs without finetuning were used for the task on finance specific ABSA. Through the Heterogeneous multi-agent Discussion (HAD) framework (Xing, 2024), the accuracy and F1 scores obtained on established datasets like FiQA (Maia et al., 2018), demonstrate the potential of LLMs to perform finance specific ABSA without specific financial training.

3 DATA SOURCE

This section provides an overview of the data sources utilized for this project. The backtesting period was limited to January 2019 through June 2019 due to constraints on computational resources and cost.

Our backtesting framework focused exclusively on S&P 500 constituents to maintain a well-defined set of liquid, large-cap stocks. To mitigate survivorship bias, we dynamically updated our universe to reflect real-time changes in the index composition. For instance, when First Republic Bank (FRC) replaced SCANA Corporation (SCG) following its acquisition by Dominion Energy on 2nd January 2019 (S&P Global, 2018), our set of stock used for backtesting was adjusted accordingly. Historical market data, including OHLC (open, high, low, close) prices and trading volumes, was sourced from the Alpha Vantage (2025) API.

The following Macro data were selected to be representative of the macroeconomic variables-interest rates, inflation, employment and economic growth-that represent the distinct economic categories as outlined by Nuriyev et al. (2024).

- Inflation:
 - Consumer Price Index (CPI) (U.S. Bureau of Labor Statistics, 2025a)
 - Producer Price Index (PPI) (U.S. Bureau of Labor Statistics, 2025c)
 - Personal Consumption Expenditures (PCE) (U.S. Bureau of Economic Analysis, 2025)
- Employment:
 - Non-Farm Payrolls (NFP) (U.S. Bureau of Labor Statistics, 2025b)
- Economic Growth:
 - Purchasing Managers’ Index (PMI) (Investing.com, 2025)
- Interest Rates:
 - Federal Open Market Committee (FOMC) minutes (Board of Governors of the Federal Reserve System, 2025)

Notably, the PCE data for 2018-Dec, 2019-Jan and 2019-Feb were released off schedule due to the partial federal government shutdown. As such, the latest available data from previous months was used instead.

By using keywords such as stock-specific company names and global market identifiers, news articles were scraped from the web via NewsAPI (2025) to be used as sentiment data. A total of 300000 articles were used to form the backtesting window of 6 months, 2019-Jan to 2019-Jun. Each news article consisted of the published date, title, description and content.

4 DESIGN PROCESS

Top-down sector allocation investing requires a comprehensive understanding of both macroeconomic conditions and company-specific sentiments. As shown in our architecture diagram (refer to Fig 1), we address this through a dual-stream processing pipeline that analyzes news articles and macroeconomic data in parallel. The system integrates these inputs through a Ranking Agent which generates stock rankings that inform sector-level investment decisions. The architecture incorporates memory and reflection modules (represented as circles in the diagram) as discussed in Section 2.3.2 to help the system develop a high-level understanding of the macroeconomic environment.

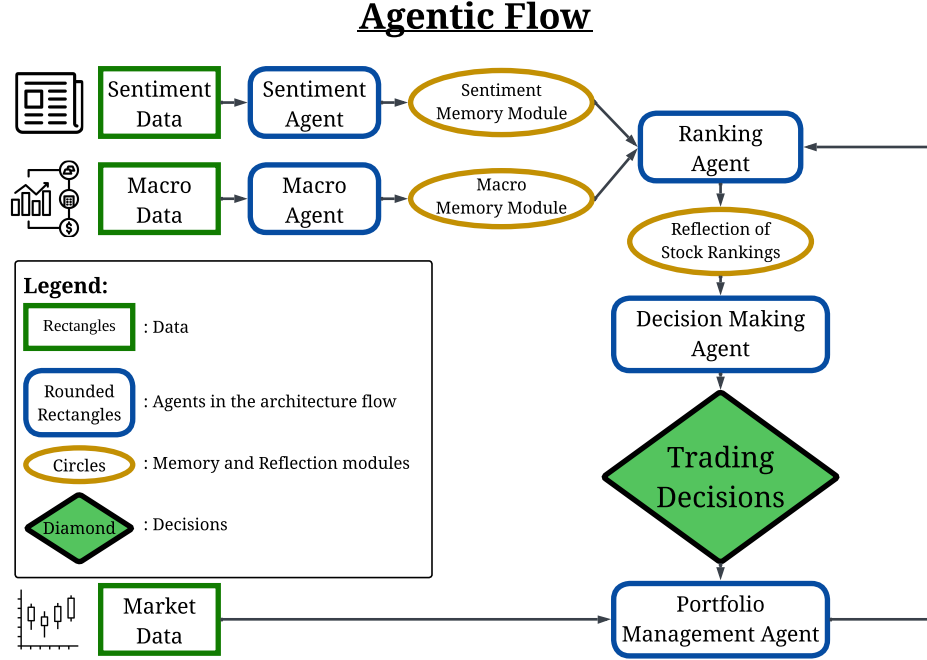


Figure 1: Overview of agentic flow designed. Icons obtained from Flaticon.com

As outlined in Section 2.4, recent research has demonstrated that LLMs can perform financial sentiment analysis without domain-specific training, guiding our design of the sentiment analysis pipeline. The Sentiment Agent, powered by the DeepSeek-7B-Chat model (DeepSeek-AI, 2025), processes news articles through two critical analyses: Named Entity Recognition (NER) and ABSA. NER is essential because our keyword search encompasses both stock-specific and broader market-relevant terms, enabling the identification of articles pertinent to target stocks that might be missed by simple keyword matching. The ABSA component analyzes article descriptions and content to identify multiple aspects of sentiment, providing the Ranking Agent with granular information about how macroeconomic conditions may amplify or diminish stock-specific sentiments through various channels. The resulting sentiment data (stock ticker and aspect sentiment pairs) is maintained in a dedicated Sentiment Memory Module, which serves as a persistent repository for historical sentiment information. The DeepSeek-Chat-7B model, selected for three key advantages: its open-source nature, efficient inference speed, and cost-effective operation. The Sentiment Agent executes the analysis through a specialized prompt template (detailed in Appendix, Section 9.1).

Parallel to the sentiment pipeline, the macroeconomic data stream is processed by a Macro Agent. Since raw macroeconomic values lack contextual meaning for LLMs, this component calculates month-over-month percentage changes for key economic indicators including CPI, PPI, PCE, NFP, and PMI. These relative changes provide meaningful representations of macroeconomic trends that LLMs can effectively interpret and analyze. Additionally, the Macro Agent processes FOMC meeting minutes using DeepSeek-Chat-7B to generate concise summaries to be stored as memories in the Macro Memory Module for subsequent retrieval. For more information about the prompts used, please refer to the Appendix, Section 9.2.

The core of the agentic architecture is the Ranking agent, which retrieves memories from the News and Macro modules, as well as current portfolio positions from the Portfolio Management Agent, to generate a Stock Ranking reflection for a single day. To prevent look-ahead bias, the agent operates under strict temporal constraints, where it only processes news memories published on the previous trading day and macro memories based on their official release dates, using the latest available data points for each economic indicator. This temporal filtering mechanism ensures that all trading decisions are made using only information that would have been available to market participants at the time of trading, maintaining the practical applicability of our backtesting results.

The Ranking Agent implements a hierarchical top-down sector allocation strategy through three sequential analytical phases: (i) macroeconomic analysis, (ii) sector allocation, and (iii) sentiment integration. In the initial phase, the agent analyzes inflation trends, evaluates economic strength indicators, and assess monetary policy trajectories. This macroeconomic assessment helps identify which sectors perform the best in the given macroeconomic environment. The second phase focuses on portfolio optimization, where the agent determines optimal allocations across the 11 GICS sectors (S&P Global Market Intelligence, 2016) and evaluates necessary portfolio adjustments by comparing target allocations against current sector exposures. In the final phase, leverages sentiment-aspect pairs from the Sentiment Module to select stocks within each targeted sector. This refinement process evaluates how prevailing macroeconomic conditions may amplify positive sentiments or mitigate negative ones for specific business aspects, enabling the selection of stocks best positioned for both long and short positions within their respective sectors. More information on the prompt used to generate the stock ranking reflections is available in Appendix, Section 9.3.

We employ the Llama-3.3-70B-Instruct model (Meta AI, 2024a) as the foundation for our Ranking Agent. This model selection was motivated by several key factors: its open-source nature, its extensive training on instruction-following tasks using publicly available datasets and over 25M synthetically generated examples. These characteristics make it particularly well-suited for interpreting and executing the complex, multi-phase analysis required for our sector allocation strategy. The daily stock ranking reflection is processed by a Decision-making Agent, implemented using Llama-3.2-3B (Meta AI, 2024b), which converts qualitative rankings into structured JSON trading decisions. Given the standardized format of the reflections generated by the Ranking Agent, we opted for a lighter-weight instruction-following model to efficiently parse and transform these reflections into actionable trade signals. These decisions are then routed to the Portfolio Management Agent, which interfaces with real-time market data to optimize trade execution.

The Portfolio Management Agent implements a conservative capital management strategy designed to balance portfolio performance with risk mitigation. To maintain sufficient liquidity for opportunistic trading while preventing overexposure, the agent enforces a 90% maximum capital utilization threshold. During the execution of trades, new positions are initiated only when they do not create conflicts with existing holdings, whereas positions are fully liquidated when new rankings indicate a shift in stock sentiment outlook. Furthermore, to maintain alignment with our investment scope, positions in stocks removed from the S&P 500 index during the backtesting window are automatically liquidated.

5 METHODOLOGY

To evaluate the effectiveness of this architecture, we implement two distinct methodologies: our proposed top-down sector-based allocation strategy that leverages on LLM-driven macroeconomic analysis, and a baseline cross-momentum strategy that build a portfolio purely based on momentum factors. Compared to the top-down sector allocation strategy, the ranking agent in the cross-momentum strategy would only carry out step (iii), integrating sentiment and macroeconomic information, to generate the reflection of stock rankings. Also, the size of the position of each trade is fixed at 1% of the portfolio. More information is available in Section 9.4 of the Appendix.

An initial capital of USD100M was used to simulate real-world institutional trading conditions and provide sufficient scale to meaningfully assess the impact of transaction costs on strategy performance. Trading costs were tracked through a comprehensive framework that accounts for both explicit and implicit costs. A commission rate of 10 basis points per trade was applied to reflect standard institutional brokerage fees, while market impact costs of 10 basis points were included to account for price slippage during execution.

Table 1: Backtesting results from Jan-2019 to Jun-2019

Strategy	PCT Change in Portfolio	Sharpe Ratio
Cross-Momentum	-1.39%	-0.61
Sector-Allocation	8.79%	2.51

6 RESULTS

The portfolio percentage change and Sharpe ratio were selected as the primary metrics for evaluating the strategy, as they provide complementary perspectives on both performance and risk-adjusted returns. Portfolio percentage change offers a normalized view of performance, allowing for direct comparisons across different portfolio sizes and time periods by expressing gains and losses as a proportion of the initial investment.

This facilitates a standardized assessment of the strategy’s performance. In contrast, the Sharpe ratio, which gauges excess returns relative to volatility, was chosen to evaluate the risk-adjusted efficiency of the strategy. Together, these metrics offer a comprehensive evaluation of the strategy’s profitability and its ability to generate sustainable returns.

The results presented in the Table 1 are calculated as follows: Portfolio percentage change is derived from the final portfolio value, which includes unrealized profits, minus trading costs. The Sharpe ratio is computed by dividing the daily excess returns by the standard deviation of the portfolio’s returns, providing a measure of risk-adjusted performance.

The results indicate that the Sector-Allocation strategy significantly outperformed the Cross-Momentum approach across both metrics. While the Cross-Momentum strategy resulted in a negative portfolio return of -1.39% and a Sharpe ratio of -0.61, suggesting poor risk-adjusted performance, the Sector-Allocation strategy delivered superior results with an 8.79% portfolio gain and a strong Sharpe ratio of 2.51. The positive Sharpe ratio for the Sector-Allocation strategy indicates that it not only generated better absolute returns but also achieved this performance with favorable risk-adjusted characteristics. This suggests that the sector-based approach was more effective at capturing market opportunities while managing volatility. The stark contrast in performance metrics between the two strategies highlights the benefits of incorporating top-down sector allocation in trading strategies.

7 LIMITATIONS AND FUTURE DIRECTIONS

While our results demonstrate promising performance, several limitations and opportunities for future research warrant discussion. A primary limitation is the relatively short backtesting period from 2019 to 2024, constrained by computational resources and costs. Extending this period would provide more robust validation across different market cycles and economic conditions. Additionally, our current implementation could benefit from larger language models, which have demonstrated superior capabilities in complex Natural Language Processing tasks. Specifically, using models with more parameters could enhance the accuracy of ABSA, NER, and the interpretation of FOMC minutes by capturing more subtle linguistic nuances and relationships.

Several promising directions exist for future research. First, expanding the model’s data inputs to include company fundamental data, particularly quarterly earnings reports, could reveal whether incorporating bottom-up financial metrics alongside our top-down approach yields meaningful improvements in portfolio performance (Ong et al., 2025).

Second, the ranking mechanism could be enhanced through the implementation of LTR algorithms as mentioned in Section 2.1 or Reinforcement Learning techniques. These approaches have demonstrated success in improving cross-sectional strategies, suggesting potential benefits for sector allocation. By incorporating historical ranking performance as feedback, these methods could dynamically optimize allocation decisions and enhance the model’s adaptability to changing market conditions.

Additionally, integrating FINMEM’s (Yu et al., 2024) architecture could enhance the model’s ability to process multi-timeframe financial data while maintaining important information over extended periods, potentially capturing longer-term patterns and relationships that human traders might overlook.

Lastly, financial Explainable Artificial Intelligence techniques (Yeo et al., 2025) such as counterfactual explanations, can be used to evaluate performance in alternative scenarios. These methods provide insight into LLM decision making and support regulatory compliance.

8 CONCLUSION

In this paper, we introduced a novel methodology leveraging Large Language Models for top-down sector allocation through integrated analysis of macroeconomic and market sentiment data. Our approach demonstrates that LLMs can effectively process and synthesize multiple data streams simultaneously, enabling dynamic sector allocation adjustments based on market conditions. The empirical results validate the effectiveness of this approach, with our strategy achieving an 8.79% return and a Sharpe ratio of 2.51, significantly outperforming traditional cross-momentum approaches which recorded a -1.39% return and -0.61 Sharpe ratio during the same period.

These results suggest that systematic top-down analysis using LLMs offers a viable approach for enhancing portfolio allocation decisions. The framework’s ability to process unstructured data and identify macro financial relationships represents a step forward in automating complex market analysis tasks. As the field of quantitative investing continues to evolve, our methodology demonstrates the potential for LLMs to complement existing investment strategies by providing a more systematic approach to sector-level portfolio construction.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the MOE Tier 1 Startup project, under grant number #22-3565-A0001-1. The research is also supported by MOE Academic Research Fund Tier 2 (STEM RIE2025 Award MOE-T2EP20123-0005) and RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026), administered by A*STAR, as well as supported by Alibaba Group and NTU Singapore.

REFERENCES

- Alpha Vantage. Alpha vantage api, 2025. URL <https://www.alphavantage.co/>. Accessed: 2025-02-06.
- Board of Governors of the Federal Reserve System. Fomc meeting calendars, statements, and minutes, 2025. URL <https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm>. Accessed: 2025-02-06.
- Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hultender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pages 89–96, 2005.
- Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*, pages 129–136, 2007.
- Imran Sharif Chaudhry, Muhammad Ayyoub, and Fatima Imran. Does inflation matter for sectoral growth in pakistan? an empirical analysis. *Pakistan Economic and Social Review*, pages 71–92, 2013.
- Anna Cieslak and Michael McMahon. Tough talk: The fed and the risk premium. *Available at SSRN*, 2023.
- DeepSeek-AI. Deepseek LLM 7b chat, 2025. URL <https://huggingface.co/deepseek-ai/deepseek-llm-7b-chat>. Accessed: 2025-02-10.
- Han Ding, Yinheng Li, Junhao Wang, and Hang Chen. Large language model agent in financial trading: A survey. *arXiv preprint arXiv:2408.06361*, 2024.
- Kelvin Du, Frank Xing, Rui Mao, and Erik Cambria. Financial sentiment analysis: Techniques and applications. *ACM Computing Surveys*, 56(9):1–42, 2024. doi: 10.1145/3649451.
- Georgios Fatouros, Konstantinos Metaxas, John Soldatos, and Dimosthenis Kyriazis. Can large language models beat wall street? unveiling the potential of ai in stock selection. *arXiv preprint arXiv:2401.03737*, 2024.

- Abeyratna Gunasekarage, Anirut Pisedtasalasai, and David M Power. Macroeconomic influence on the stock market: evidence from an emerging market in south asia. *Journal of Emerging Market Finance*, 3(3):285–304, 2004.
- Investing.com. Ism manufacturing pmi economic calendar, 2025. URL <https://www.investing.com/economic-calendar/ism-manufacturing-pmi-173>. Accessed: 2025-02-06.
- Ayesha Jabeen, Muhammad Yasir, Yasmeen Ansari, Sadaf Yasmin, Jihoon Moon, and Seungmin Rho. An empirical study of macroeconomic factors and stock returns in the context of economic uncertainty news sentiment using machine learning. *Complexity*, 2022(1):4646733, 2022.
- Francisco Jareño and Loredana Negrut. Us stock market and macroeconomic factors. *Journal of Applied Business Research*, 32(1):325, 2016.
- Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of finance*, 48(1):65–91, 1993.
- Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. Towards mitigating hallucination in large language models via self-reflection. *arXiv preprint arXiv:2310.06271*, 2023.
- Alejandro Lopez-Lira and Yuehua Tang. Can chatgpt forecast stock price movements? return predictability and large language models. *arXiv preprint arXiv:2304.07619*, 2023.
- Yu Ma, Rui Mao, Qika Lin, Peng Wu, and Erik Cambria. Multi-source aggregated classification for stock price movement prediction. *Information Fusion*, 91:515–528, 2023.
- Yu Ma, Rui Mao, Qika Lin, Peng Wu, and Erik Cambria. Quantitative stock portfolio optimization by multi-task learning risk and return. *Information Fusion*, 104:102165, 2024. ISSN 1566-2535. doi: 10.1016/j.inffus.2023.102165. URL <https://www.sciencedirect.com/science/article/pii/S1566253523004815>.
- Macedo Maia, Siegfried Handschuh, André Freitas, Brian Davis, Ross McDermott, Manel Zarrouk, and Alexandra Balahur. Wwv’18 open challenge: financial opinion mining and question answering. In *Companion proceedings of the the web conference 2018*, pages 1941–1942, 2018.
- Lorenzo Malandri, Frank Xing, Carlotta Orsenigo, Carlo Vercellis, and Erik Cambria. Public mood-driven asset allocation: the importance of financial sentiment in portfolio management. *Cognitive Computation*, 10(6):1167–1176, 2018.
- Meta AI. Llama 3.3-70b instruct, 2024a. URL <https://huggingface.co/meta-llama/Llama-3.3-70B-Instruct>. Accessed: 2025-02-10.
- Meta AI. Llama 3.2-3b: A foundational large language model, 2024b. URL <https://huggingface.co/meta-llama/Llama-3.2-3B>. Accessed: 2025-02-10.
- NewsAPI. Newsapi - get news from multiple sources with a single request, 2025. URL <https://newsapi.org>. Accessed: 2025-02-10.
- Dmitriy Nuriyev, Songyun Duan, and Lingjie Yi. Augmenting equity factor investing with global macro regimes. In *Proceedings of the 5th ACM International Conference on AI in Finance*, pages 445–452, 2024.
- Keane Ong, Wihan van der Heever, Ranjan Satapathy, Erik Cambria, and Gianmarco Mengaldo. FinXABSA: Explainable finance through aspect-based sentiment analysis. In *IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 773–782. IEEE, 2023.
- Keane Ong, Rui Mao, Ranjan Satapathy, Ricardo Shirota Filho, Erik Cambria, Johan Sulaeman, and Gianmarco Mengaldo. Explainable natural language processing for corporate sustainability analysis. *Information Fusion*, 115:102726, 2025. ISSN 1566-2535. doi: 10.1016/j.inffus.2024.102726. URL <https://www.sciencedirect.com/science/article/pii/S1566253524005049>.

- Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22, 2023.
- Daniel Poh, Bryan Lim, Stefan Zohren, and Stephen Roberts. Building cross-sectional systematic strategies by learning to rank. *arXiv preprint arXiv:2012.07149*, 2020.
- S&P Global. First republic bank set to join s&p 500, 2018. URL <https://press.spglobal.com/2018-12-27-First-Republic-Bank-Set-to-Join-S-P-500>. Accessed: 2025-02-09.
- S&P Global Market Intelligence. Gics® mapbook brochure, 2016. URL <https://www.spglobal.com/marketintelligence/en/documents/gics-mapbook-brochure.pdf>. Accessed: 2025-02-09.
- U.S. Bureau of Economic Analysis. Table 2.8.7. percent change from preceding period in prices for personal consumption expenditures by major type of product, monthly, 2025. URL <https://apps.bea.gov/iTable/reqid=19&step=3&isuri=1&1921=survey&1903=84>. Accessed: 2025-02-06.
- U.S. Bureau of Labor Statistics. Consumer price index (cpi) for all urban consumers (cpi-u), u.s. city average, all items, seasonally adjusted, 2025a. URL <https://data.bls.gov/dataViewer/view/timeseries/CUSR0000SA0>. Accessed: 2025-02-06.
- U.S. Bureau of Labor Statistics. All employees, total nonfarm (payems), 2025b. URL <https://fred.stlouisfed.org/series/PAYEMS>. Retrieved from FRED, Federal Reserve Bank of St. Louis. Accessed: 2025-02-06.
- U.S. Bureau of Labor Statistics. Producer price index by commodity: Final demand, 4-digit level, 2025c. URL <https://data.bls.gov/dataViewer/view/timeseries/WPSFD4>. Accessed: 2025-02-06.
- Meiyun Wang, Kiyoshi Izumi, and Hiroki Sakaji. Llmfactor: Extracting profitable factors through prompts for explainable stock movement prediction. *arXiv preprint arXiv:2406.10811*, 2024.
- Qiang Wu, Christopher JC Burges, Krysta M Svore, and Jianfeng Gao. Adapting boosting for information retrieval measures. *Information Retrieval*, 13:254–270, 2010.
- Ruoxu Wu. Portfolio performance based on llm news scores and related economical analysis. *Available at SSRN 4709617*, 2024.
- Frank Xing. Designing heterogeneous llm agents for financial sentiment analysis. *ACM Transactions on Management Information Systems*, 2024.
- Frank Xing, Lorenzo Malandri, Yue Zhang, and Erik Cambria. Financial sentiment analysis: An investigation into common mistakes and silver bullets. In *COLING*, pages 978–987, 2020.
- Wei Jie Yeo, Wihan Van Der Heever, Rui Mao, Erik Cambria, Ranjan Satapathy, and Gianmarco Mengaldo. A comprehensive review on financial explainable AI. *Artificial Intelligence Review*, 58:19, 2025.
- Yangyang Yu, Haohang Li, Zhi Chen, Yuechen Jiang, Yang Li, Denghui Zhang, Rong Liu, Jordan W Suchow, and Khaldoun Khashanah. Finmem: A performance-enhanced llm trading agent with layered memory and character design. In *Proceedings of the AAAI Symposium Series*, volume 3, pages 595–597, 2024.
- Wentao Zhang, Lingxuan Zhao, Haochong Xia, Shuo Sun, Jiaze Sun, Molei Qin, Xinyi Li, Yuqing Zhao, Yilei Zhao, Xinyu Cai, et al. Finagent: A multimodal foundation agent for financial trading: Tool-augmented, diversified, and generalist. *arXiv preprint arXiv:2402.18485*, 2024.

9 APPENDIX

9.1 PROMPT FOR SENTIMENT PIPELINE

You are a financial sentiment analyzer. Your task is to
 ↳ analyze news articles about companies and extract
 ↳ sentiment information about different aspects of the
 ↳ company mentioned in the article. Respond ONLY with a JSON
 ↳ object, no additional text or markdown.

Instructions:

1. Analyze the provided news article's title, description,
 ↳ and content.
2. Identify the main stock/company being discussed.
3. Extract 3 to 5 key aspects discussed in the article
 ↳ (e.g., earnings, revenue, management, products, market
 ↳ position, growth, competition).
4. For each aspect, determine the sentiment on a scale:
 - positive (1)
 - neutral (0)
 - negative (-1)
5. Return the analysis in the following JSON format
 ↳ exactly, replacing the example values with your
 ↳ analysis.

```
{ "stock": "AAPL",
  "aspect_sentiment_pairs": [
    [ "revenue", 1 ],
    [ "product_performance", -1 ],
    [ "services", 1 ]
  ] }
```

Rules:

- Only include information that is explicitly discussed in
 ↳ the article
- Only include aspects that belong to the list of relevant
 ↳ aspects to look out for
- Base sentiment strictly on the article's content, not
 ↳ external knowledge
- Be consistent with aspect naming (e.g., always use
 ↳ "revenue" instead of mixing "revenue" and "sales")
- Don't include duplicate aspects
- Limit to the most significant 3-5 aspects mentioned
- Use the most commonly known stock ticker
- If no clear stock ticker is mentioned, use the company
 ↳ name in the stock field,

List of relevant aspects to look for:

- revenue/sales
- earnings/profit
- market_share
- product_performance
- management
- growth
- competition
- regulatory
- innovation

- customer_demand
- operational_efficiency
- partnerships
- risk
- strategy

Example Analysis:

Input:

Title: EGG Reports Record Q4 Revenue Despite Product Sales
→ Miss

Description: EGG posts strong services growth but flagship
→ product disappoints

Content: EGG Inc. reported its highest-ever fourth-quarter
→ revenue of \$89.5 billion, though Product sales fell
→ short of analyst expectations. The company's services
→ division saw remarkable growth, up 16% year-over-year,
→ helping offset the weaker hardware performance. CEO
→ Tim Cook expressed confidence in the company's product
→ pipeline but acknowledged supply chain challenges.

Output:

```
{"stock": "EGG",  
"aspect_sentiment_pairs": [  
  ["revenue", 1],  
  ["product_performance", -1],  
  ["services", 1],  
  ["supply_chain", -1]  
]}
```

End of example""

9.2 PROMPT FOR SUMMARIZING FOMC MINUTES

Please analyze the following FOMC meeting minutes and provide
→ a structured analysis focusing on these key aspects:

1. Interest Rate Policy and Outlook
 - Identify explicit statements about current interest rate
→ decisions
 - Extract any forward guidance or projections about future
→ rate movements
 - Note any dissenting views or alternative scenarios
→ discussed
2. Economic Assessment
 - Summarize the Committee's view on:
 - * GDP growth and economic activity
 - * Labor market conditions
 - * Inflation rates and price stability
 - * Financial market conditions
3. Risk Analysis
 - List major risks to the economic outlook
 - Detail both upside and downside risks

Meeting Minutes:
{text}

Provide analysis in a clear, structured format.

9.3 PROMPT FOR RANKING AGENT IN TOP-DOWN STRATEGY

You are a quantitative macro strategist specializing in top-down
→ allocation strategies. Your task is to analyze macroeconomic
→ conditions first, then use sentiment data to select stocks
→ that align with the macro outlook.

Given the following inputs:

[MACRO DATA TRENDS]

- Latest trend readings: CPI {value}, PPI {value}, PCE
→ {value}, NFP {value}, PMI {value}

[FOMC MINUTES SUMMARY]

- Recent FOMC minutes summary: {key_points}

[STOCK UNIVERSE]

- List of S&P 500 stocks with their sentiment data
- Format: date|ticker|aspect_sentiment_pairs

[CURRENT PORTFOLIO]

- List of current positions:
 - Long positions: [(TICKER, weight)]
 - Short positions: [(TICKER, weight)]

For sector analysis, use the 11 GICS sectors:

- Information Technology
- Financials
- Healthcare
- Consumer Discretionary
- Consumer Staples
- Industrials
- Energy
- Materials
- Communication Services
- Utilities
- Real Estate

Analysis Process:

1. Macro Environment Assessment:
 - Analyze inflation trends (CPI, PPI, PCE)
 - Evaluate economic strength (PMI, NFP)
 - Consider monetary policy outlook (FOMC)
 - Identify which sectors should perform best in this
→ environment
2. Sector-Level Analysis:
 - Determine sector overweight/underweight based on
→ macro
 - Compare current sector exposure vs target allocation
 - Identify sectors requiring position changes
3. Stock Selection Within Sectors:
 - Prioritize stocks in preferred sectors
 - Use sentiment data to rank within sectors
 - Consider existing positions (avoid unnecessary
→ turnover)

Provide response in format:

MACRO ENVIRONMENT:

- Current economic conditions
- Key drivers
- Sector implications

SECTOR VIEWS:

- Overweight sectors: [list with rationale]
- Underweight sectors: [list with rationale]
- Current vs Target exposure

PORTFOLIO RECOMMENDATIONS:

Positions to Long:

1. [TICKER] (Sector: X)
 - Macro alignment: [explanation]
 - Sector view: [explanation]
 - Supporting sentiment: [relevant aspects]
 - Position size recommendation: [X%]

Positions to Short:

[Same format as above]

TURNOVER ANALYSIS:

- Summary of recommended changes
- Rationale for maintaining existing positions

9.4 PROMPT FOR RANKING AGENT IN CROSS-SECTIONAL STRATEGY

You are a quantitative analyst specializing in

- ↪ sentiment-driven trading strategies. Your task is to
- ↪ analyze and rerank stocks for a long-short strategy based
- ↪ on sentiment data and macroeconomic context.

Given the following inputs:

[LIST OF STOCKS]

- List of stocks that can be included in the portfolio
- All stocks are assumed to start with the same score

[SENTIMENT DATA]

- List of stocks with sentiment-aspect pairs from news
 - ↪ articles
- Each pair contains: stock ticker, date, specific aspect
 - ↪ (e.g., "management", "financial performance"), and
 - ↪ sentiment score (-1, 0, 1)
- Sample format: 2024-01-15|AAPL|[[management,
 - ↪ 1],[revenue, -1]]
 - ↪ (published_date|stock|aspect_sentiment_pairs)

[MACRO DATA TRENDS]

- Latest trend readings: CPI {value}, PPI {value}, PCE
 - ↪ {value}, NFP {value}, PMI {value}

[FOMC MINUTES SUMMARY]

- Recent FOMC minutes summary: {key_points}

Analyze how these macro trends might intensify or mitigate

- ↪ the sentiment towards different aspects. For example:
- An increase in inflation might make cost-related
 - ↪ sentiments more impactful

- An increase in PMI readings might reduce supply chain
 - ↪ concern impacts
- Trends for Employment data might affect consumer demand
 - ↪ sentiment importance

For each stock, please:

1. Evaluate each sentiment-aspect pair
2. Adjust the importance of each aspect based on current
 - ↪ macro conditions
3. Assign a composite score that considers:
 - Sentiment scores
 - Macro-influenced aspect weights

Then:

1. Rank the stocks from highest to lowest composite scores
2. Split into long candidates (positive scores) and short
 - ↪ candidates (negative scores)
3. Explain your reasoning for the each long and short
 - ↪ picks
4. Check thorough the long and short candidates. If there
 - ↪ are duplicates, review the composite scores and keep
 - ↪ only the position with a higher composite score.

Provide your response in this format:

LONG CANDIDATES:

1. [TICKER] - Score: [X]
 - Key aspects: [list most influential aspects]
 - Macro amplifiers: [which macro factors strengthened the
 - ↪ case]
2. [Continue for long picks...]

SHORT CANDIDATES:

[Same format as above]

MACRO ANALYSIS:

- Brief explanation of how macro conditions influenced the
 - ↪ rankings
- Which factors were most decisive

Only provide factual analysis based on the data given.