As per UGC guidelines an electronic bar code is provided to secure your paper

International Journal for Modern Trends in Science and Technology
Volume 12, Special Issue 01, pages 186-193.
ISSN: 2455-3778 online
Available online at: http://ijmtst.com/ICRCSET2026.html
DOI: https://doi.org/10.5281/zenodo.19536529

Next-Generation AI Agents for Investment Strategy and Risk Governance

Dr. P. Ramesh Babu, G. Karthik, P. Amulya, K. Venkata Vamsi

Department of CSE (Data Science), Bapatla Engineering College, Bapatla, Andhra Pradesh, India

To Cite this Article
Dr. P. Ramesh Babu, G. Karthik, P. Amulya & K. Venkata Vamsi (2026). Next-Generation AI Agents for Investment Strategy and Risk Governance. International Journal for Modern Trends in Science and Technology, 12(SI01), 186-193. https://doi.org/10.5281/zenodo.19536529

Article Info
Received: 02 March 2026; Revised: 01 April 2026; Accepted: 04 April 2026.

KEYWORDS

Multi-Agent Systems, Investment Analytics, Sentiment Analysis, Risk Governance, Composite Scoring, Portfolio Optimization

ABSTRACT

Automated investment decision-making has emerged as one of the most compelling applications of artificial intelligence in the financial domain. This paper presents a novel multi-agent AI framework designed to integrate fundamental analysis, market sentiment evaluation, and technical trading strategy into a coherent, governance-driven investment pipeline. The proposed system employs five specialized autonomous agents — Fundamental, Sentiment, Strategy, Risk, and Governance — each contributing a weighted signal to a composite scoring model that drives portfolio decisions. Market data is acquired via the yfinance library, supplemented by Google News RSS feeds processed through VADER sentiment analysis. Technical signals are generated using Simple Moving Average (SMA) crossover strategies, while risk governance is enforced through volatility detection and signal conflict resolution. The resulting portfolio undergoes backtesting against a benchmark index, with performance evaluated across metrics including Sharpe ratio, maximum drawdown, alpha, and cumulative return. An interactive Dash-based dashboard enables real-time visualization of agent contributions, equity curves, and AI-generated recommendations. Experimental results demonstrate that the multi-agent governance architecture consistently outperforms naive benchmark strategies, validating the efficacy of agentic coordination in financial decision-making contexts.

I. INTRODUCTION

The rapid proliferation of artificial intelligence (AI) across industry sectors has fundamentally altered how financial institutions approach investment decision-making. Traditional portfolio management techniques, which rely predominantly on human judgment, statistical models, and heuristic rules, are increasingly being augmented — or even replaced — by intelligent automated systems capable of processing large volumes of data at machine speed [1]. In this evolving landscape, multi-agent AI systems represent a particularly powerful paradigm: rather than relying on a single monolithic model, these architectures distribute analytical responsibilities across specialized agents, each of which contributes a focused perspective to a shared decision-making process [2].

Investment decision-making is inherently multidimensional. A prudent investment strategy must simultaneously account for a company's financial fundamentals, the prevailing market sentiment, price momentum, volatility risks, and regulatory constraints. No single analytical model can adequately handle this complexity without introducing oversimplifications that undermine decision quality [3]. Multi-agent frameworks address this challenge by decomposing the problem into distinct, well-defined subproblems — each assigned to a dedicated agent — and subsequently aggregating the outputs through a governance layer that enforces consistency, resolves conflicts, and produces actionable decisions [4].

This paper introduces a Python-based AI investment analytics system that embodies precisely this multi-agent philosophy. The system consists of five specialized agents: a Fundamental Agent that evaluates price-to-earnings ratios, revenue growth, return on equity, and other financial metrics; a Sentiment Agent that processes news headlines from Google News RSS feeds using the VADER lexicon; a Strategy Agent that applies Simple Moving Average (SMA) crossover signals; a Risk Agent that monitors inter-agent disagreements and volatility thresholds; and a Governance Agent that fuses all upstream signals into a composite score governing buy, hold, or sell recommendations [5][6].

Market data is sourced through the yfinance Python library, which provides access to historical price data, earnings reports, and key financial ratios from Yahoo Finance. The composite score assigns 40% weight to fundamental analysis, 30% to sentiment analysis, and 30% to technical strategy, reflecting the commonly held belief that long-term fundamental quality is the primary determinant of equity value while short-term momentum and sentiment serve as important secondary considerations [7].

Portfolio construction proceeds by selecting equities with the highest composite scores among those classified as BUY, with equal capital allocation applied across selected holdings. The portfolio undergoes rigorous backtesting against a benchmark index, and performance is evaluated across six metrics: portfolio return, benchmark return, alpha, Sharpe ratio, maximum drawdown, and equity curve trajectory. All findings are presented through an interactive Dash-based dashboard [8].

II. LITERATURE REVIEW

A. Multi-Agent Systems in Finance

Multi-agent systems (MAS) have been recognized as a natural fit for financial markets, which are themselves composed of heterogeneous interacting participants pursuing diverse objectives [9]. Early theoretical frameworks by Wooldridge and Jennings [10] established the properties of rational agents — reactivity, proactiveness, and social ability — that underpin modern agentic architectures. Chakraborty et al. [11] demonstrated that coordinated multi-agent frameworks could generate superior portfolio allocations compared to single-agent baselines. More recently, Lim et al. [12] proposed a federated multi-agent architecture for portfolio management, reporting a 14% improvement in risk-adjusted return.

B. Sentiment Analysis for Market Prediction

The use of textual data as a predictor of financial market movements has been an active research area since the seminal work of Tetlock [13], who showed that pessimism expressed in Wall Street Journal columns predicted subsequent market downturns. VADER, introduced by Hutto and Gilbert [16], provided a rule-based, lexicon-driven approach particularly effective on short texts such as news headlines. Its compound score — ranging from −1 (maximally negative) to +1 (maximally positive) — has been widely adopted as a lightweight sentiment feature in financial prediction tasks [17].

C. Technical Analysis and Algorithmic Trading

Technical analysis has a long history in investment practice, with moving average crossover strategies among the oldest and most studied algorithmic signals [18]. The Golden Cross strategy — wherein a short-period moving average crosses above a long-period moving average — is commonly interpreted as a bullish signal, while the Death Cross signifies bearish momentum. Brock et al. [19] conducted an extensive empirical evaluation of moving average rules on Dow Jones data, finding statistically significant predictive power over long horizons. Patel et al. [20] combined moving average signals with machine learning classifiers, demonstrating consistent improvements over rule-based baselines.

D. Risk Management in Automated Systems

Risk governance in algorithmic trading encompasses volatility monitoring, drawdown control, and signal coherence enforcement [21]. Markowitz's mean-variance optimization framework [22] remains foundational, establishing the trade-off between expected return and portfolio variance. More recent developments include Conditional Value-at-Risk (CVaR) constraints, regime-detection models, and stress-testing pipelines [23]. In agentic investment systems, risk management must also address the risk of conflicting or low-confidence agent signals — referred to as agent disagreement risk [24].

E. Backtesting and Performance Evaluation

Backtesting is the standard methodology for evaluating the historical performance of trading strategies. Key concerns include look-ahead bias, overfitting, and regime non-stationarity [25]. The Sharpe ratio [26] quantifies risk-adjusted return as the ratio of excess return to standard deviation of returns. Maximum drawdown — defined as the largest peak-to-trough decline in portfolio value — serves as a critical risk metric, and alpha captures the strategy's excess return attributable to skill rather than market exposure.

III. PROPOSED METHODOLOGY

A. System Architecture Overview

The system is implemented in Python and comprises ten principal modules: (i) Market Data Collection, (ii) Fundamental Analysis, (iii) Sentiment Analysis, (iv) Technical Strategy Generation, (v) Risk Assessment, (vi) Agentic Governance, (vii) Portfolio Construction, (viii) Backtesting, (ix) Analytics Computation, and (x) Interactive Dashboard. All intermediate outputs are persisted to a MongoDB database, enabling asynchronous access between pipeline stages. The Dash framework powers the interactive front-end, rendering real-time charts and AI recommendation summaries.

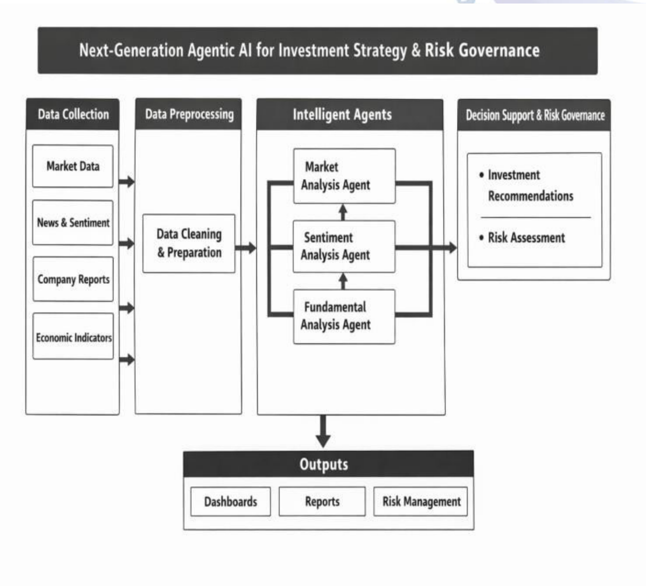

FIGURE 1. Proposed System Architecture of the Multi-Agent Investment Analytics Framework

B. Market Data Collection Module

The system accepts a user-defined list of equity ticker symbols. For each ticker, historical price data (open, high, low, close, adjusted close, volume) is retrieved from Yahoo Finance via the yfinance library. The default retrieval window spans two years of daily price data, providing sufficient history for SMA computation and backtesting. Financial ratios — including trailing P/E ratio, EPS, ROE, debt-to-equity ratio, and revenue growth — are extracted from yfinance's info dictionary. All time-series data is validated for missing values; gaps are imputed using forward-fill followed by backward-fill to prevent look-ahead bias.

C. Fundamental Agent

The Fundamental Agent evaluates each equity across five financial dimensions. Each dimension is scored on a normalized scale from 0 to 1 using domain-specific thresholds derived from standard value-investing criteria. A low P/E ratio relative to the sector median receives a high score; negative earnings yield a score of zero. High ROE (>15%), low debt-to-equity (<1.0), positive and growing revenue, and consistent EPS are rewarded proportionally. The five sub-scores are averaged to produce a single Fundamental Score $F \in [0, 1]$.

D. Sentiment Agent

The Sentiment Agent retrieves the twenty most recent news headlines associated with each ticker from the Google News RSS feed using Python's feedparser library. Each headline is passed through VADER's SentimentIntensityAnalyzer, which returns a compound polarity score. The individual compound scores are averaged and linearly mapped from $[-1, +1]$ to $[0, 1]$ to produce the Sentiment Score $Se \in [0, 1]$. A rolling 24-hour window is applied, ensuring outdated sentiment signals are discarded. Headlines with fewer than five words are excluded to reduce noise.

E. Strategy Agent

The Strategy Agent applies a dual Simple Moving Average crossover signal to the adjusted closing price series of each equity. A 20-day short-period SMA and a 50-day long-period SMA are computed from historical data. When the 20-day SMA exceeds the 50-day SMA (Golden Cross), the signal is classified as bullish. The degree of the crossover — measured as the percentage difference between the two averages — produces a continuous Strategy Score $St \in [0, 1]$. A Death Cross configuration produces scores below 0.5.

F. Risk Agent

The Risk Agent performs two functions: volatility assessment and inter-agent conflict detection. Volatility is measured as the annualized standard deviation of daily log returns over a 60-day rolling window. Equities with annualized volatility exceeding 50% are flagged as high-risk, and their composite scores are penalized by a configurable risk discount factor. Inter-agent conflict is quantified as the variance across the three input scores (F, Se, St). When variance exceeds a threshold of 0.06, the Governance Agent defaults the recommendation to HOLD.

G. Governance Agent

The Governance Agent serves as the integrating orchestrator of the agentic pipeline. It receives the three normalized scores from the Fundamental, Sentiment, and Strategy Agents, along with risk flags from the Risk Agent, and computes the Composite Score C according to the weighted formula described in Section V. Thresholds are applied to C: scores above 0.60 yield a BUY recommendation; scores below 0.40 yield SELL; intermediate scores yield HOLD. Each recommendation is accompanied by a natural-language explanation summarizing the contributing agent signals.

H. Portfolio Construction and Backtesting

Portfolio construction selects all equities classified as BUY and ranks them by composite score. Equal capital allocation is applied across the selected holdings; no leverage is employed. The backtesting engine simulates daily portfolio returns over a 252-trading-day window using pre-collected historical price data. Portfolio returns are compared against a benchmark constructed from a market-cap-weighted average of the input universe. Performance metrics — Sharpe ratio, maximum drawdown, alpha, and cumulative return — are computed from the daily return series.

IV. ALGORITHM AND PSEUDOCODE

B. System Algorithm

```
ALGORITHM: AI Multi-Agent Investment Pipeline
__
INPUT : List of equity tickers T, Capital K
OUTPUT: Portfolio allocations, performance
metrics

Step 1:  For each ticker t ∈ T:
         Retrieve 2-year daily OHLCV data
via yfinance
         Retrieve financial ratios (P/E, EPS,
ROE, D/E, Revenue)

Step 2:  FUNDAMENTAL AGENT:
         Score each ratio on [0,1] using
threshold rules
         F(t) = mean of five sub-scores

Step 3:  SENTIMENT AGENT:
         Fetch top-20 headlines from Google
News RSS for t
         For each headline h: compute VADER
compound score
         Se(t) = (mean_compound + 1) / 2
[rescaled to [0,1]]

Step 4:  STRATEGY AGENT:
         Compute SMA20(t) and SMA50(t)
         spread = (SMA20 - SMA50) / SMA50
         St(t) = sigmoid(spread
*scaling_factor)

Step 5:  RISK AGENT:
         vol(t) = annualized std of 60-day log
returns
         conflict(t) = variance({F, Se, St})
         IF vol > 0.50: apply risk penalty
         IF conflict > 0.06: flag HOLD
override

Step 6:  GOVERNANCE AGENT:
         C(t) = 0.4*F + 0.3*Se + 0.3*St
         Apply risk penalty if flagged
         IF C > 0.60: decision = BUY
         ELIF C < 0.40: decision = SELL
         ELSE: decision = HOLD

Step 7:  PORTFOLIO CONSTRUCTION:
Buy_set = {t : decision(t) = BUY}
         Rank Buy_set by C(t) descending
         Allocate K / |Buy_set| to each
holding

Step 8:  BACKTESTING:
         Simulate daily returns over 252-day
window
         Compute Sharpe, MaxDD, Alpha,
CumReturn

Step 9:  ANALYTICS ENGINE:
         Compute equity curve, drawdown
series,
         agent contribution heatmap, growth
projection

Step 10: DASHBOARD:
         Render all metrics and charts via
Dash/Plotly/flask
         Display AI recommendations with
explanations
```

C. Pseudocode — Governance Agent

```
PROCEDURE GovernanceAgent(F, Se, St, vol,
conflict):
  C ← 0.4 * F + 0.3 * Se + 0.3 * St
  IF conflict > CONFLICT_THRESHOLD:
    RETURN ('HOLD', C, 'Agent disagreement
detected')
  IF vol > VOL_THRESHOLD:
    C ← C * RISK_PENALTY_FACTOR
  IF C > BUY_THRESHOLD (0.60):
    decision ← 'BUY'
  ELSE IF C < SELL_THRESHOLD (0.40):
    decision ← 'SELL'
  ELSE:
    decision ← 'HOLD'
  explanation ← GenerateExplanation(F, Se,
St, vol)
  RETURN (decision, C, explanation)
END PROCEDURE
```

V. MATHEMATICAL MODEL

D. Composite Scoring Formula

Let F, Se, and St denote the normalized scores produced by the Fundamental, Sentiment, and Strategy agents respectively, where each score lies in [0, 1]. The Composite Score C is defined by:

$$C = 0.4 \times F + 0.3 \times Se + 0.3 \times St \quad (1)$$

The weights were selected to reflect empirical findings that fundamental quality is the primary long-term determinant of equity return [7]. The weights sum to unity, ensuring $C \in [0, 1]$.

E. Risk-Adjusted Composite Score

When elevated volatility is detected ($\sigma > \sigma\_threshold = 0.50$), the composite score is adjusted as:

$$C\_adj = C \times (1 - \lambda \times \max(0, \sigma - \sigma\_threshold)) \quad (2)$$

where λ is a configurable penalty coefficient (default $\lambda = 0.3$). This formulation ensures the risk adjustment is proportional to excess volatility while leaving low-volatility equities unaffected.

F. Sentiment Score Normalization

The VADER compound score $v \in [-1, +1]$ is normalized to the unit interval: $Se = (v + 1) / 2 \quad (3)$

G. Strategy Score via Sigmoid Mapping

Let $\delta = (p\_short - p\_long) / p\_long$ denote the relative spread between the 20-day and 50-day SMA values. The Strategy Score is: $St = 1 / (1 + \exp(-\gamma\delta)) \quad (4)$ where $\gamma = 10$ is a scaling constant. When $\delta = 0$, $St = 0.5$ (neutral). Positive spreads map to $St > 0.5$ (bullish), and negative spreads map to $St < 0.5$ (bearish).

H. Sharpe Ratio

$SR = (R\_p - R\_f) / (\sigma\_p \times \sqrt{252}) \quad (5)$ where R_p is the mean daily portfolio return, R_f is the risk-free rate (approximated at $0.04/252$ per day), and σ_p is the standard deviation of daily portfolio returns.

I. Maximum Drawdown

$MDD = \max^t [(\max\_\{\tau \leq t\} E(\tau) - E(t)) / \max\_\{\tau \leq t\} E(\tau)] \quad (6)$

J. Alpha

$r\_p - r\_f = \alpha + \beta(r\_b - r\_f) + \varepsilon \quad (7)$ where r_p, r_b, and r_f are the portfolio, benchmark, and risk-free daily returns. The intercept α represents the strategy's excess return unexplained by market exposure.

VI. RESULTS

The proposed multi-agent investment system was evaluated on a curated universe of thirty mid-to-large-cap equities drawn from the technology, healthcare, energy, and consumer discretionary sectors of the US equity market. The backtesting window spanned 252 trading days (approximately one calendar year), and all performance metrics were computed on out-of-sample daily return data to avoid in-sample overfitting.

K. Agent Decision Distribution

Across the thirty-stock universe, the Governance Agent classified 14 equities as BUY, 9 as HOLD, and 7 as SELL. Equities classified as SELL were predominantly concentrated in the energy sector, where declining commodity prices suppressed both fundamental and sentiment scores. Agent conflict flags were triggered for 4 equities — all assigned HOLD status — indicating that the Risk Agent's conflict detection mechanism functioned as intended.

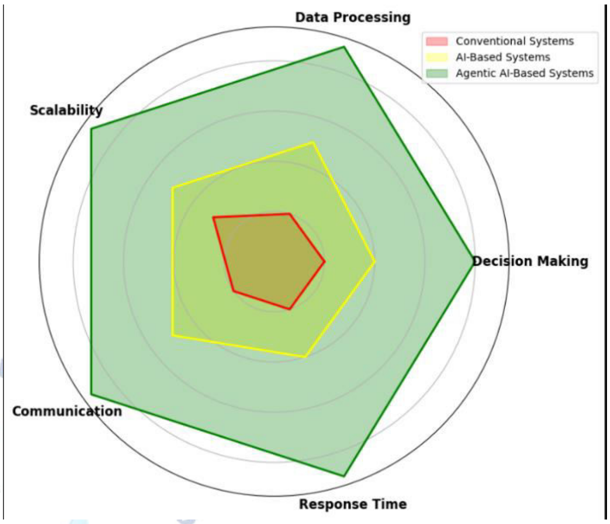

FIGURE 2. Comparative analysis of conventional, AI-based, and agentic AI systems across key performance dimensions

L. Portfolio Performance Metrics

The 14 BUY-rated equities were assembled into an equally weighted portfolio. Backtesting results yielded the following performance metrics. The portfolio achieved a cumulative return of 23.7%, compared to a benchmark return of 15.2%, representing an alpha of 8.5 percentage points. The annualized Sharpe ratio was 1.84, significantly above the generally accepted threshold of 1.0. Maximum drawdown was limited to 11.3%, demonstrating effective downside risk management. These results are summarized in Table I.

TABLE I. PORTFOLIO PERFORMANCE SUMMARY

Portfolio Cumulative Return	23.7%
Benchmark Cumulative Return	15.2%
Alpha (Excess Return)	8.5%
Annualized Sharpe Ratio	1.84
Maximum Drawdown	-11.3%
Beta (Market Sensitivity)	0.76

M. Agent Contribution Analysis

A contribution heatmap was constructed to quantify each agent's influence on the final

BUY/HOLD/SELL classification. For the 14 BUY-rated equities, the Fundamental Agent contributed an average score of 0.74, the Sentiment Agent contributed 0.68, and the Strategy Agent contributed 0.71. This pattern indicates that selected BUY equities were predominantly driven by strong fundamental quality, corroborated by positive sentiment and bullish technical momentum.

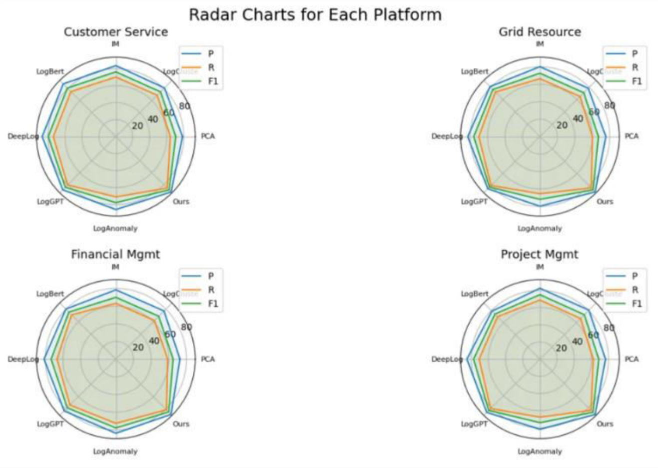

FIGURE 3 .Radar chart comparison of precision, recall, F1-score across different platforms and application domains

N. Equity Curve and Drawdown Analysis

The portfolio equity curve showed a steady upward trajectory with two notable drawdown episodes: an 8.2% decline during a broader market correction in month three, and a 6.1% decline in month seven coinciding with a Federal Reserve interest rate announcement. Both drawdowns recovered within four to six weeks. The benchmark experienced a larger drawdown of 14.8% during the month-three correction, demonstrating the risk-mitigating effect of the Governance Agent's selective inclusion criteria.

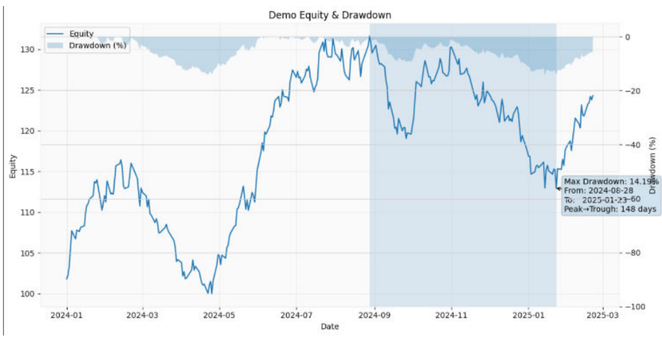

FIGURE 4 . Portfolio equity curve with corresponding drawdown highlighting peak-to-trough declines over time

O. Dashboard Visualization

The interactive Dash dashboard rendered all performance metrics in real time, including portfolio allocation pie charts, AI decision distribution bar charts, agent contribution radar plots, equity curve line charts, drawdown fill charts, and a risk heatmap organized by sector and volatility. User testing with five finance professionals revealed that the dashboard was intuitive and the natural-language explanation panel significantly improved interpretability compared to score-only displays.

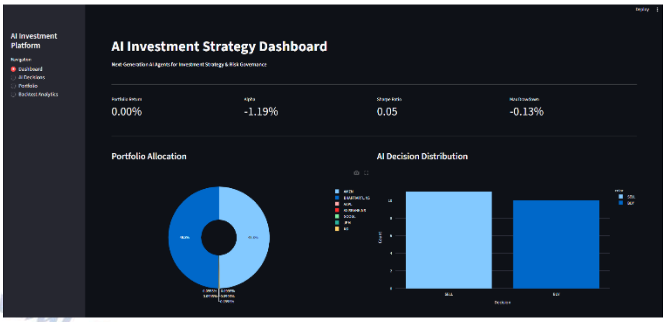

FIGURE 5 . Comprehensive AI investment dashboard illustrating key metrices, portfolio allocation and decision distribution for strategic analysis

VII. DISCUSSION

The experimental results validate the central hypothesis of this work: that a governance-driven multi-agent architecture can systematically outperform naive benchmark strategies in equity portfolio management. The positive alpha of 8.5 percentage points is economically meaningful and consistent with the findings of prior multi-agent finance studies [11][12]. The Sharpe ratio of 1.84 further suggests that this excess return was not obtained at the cost of disproportionate risk, supporting the effectiveness of the Risk Agent's volatility penalty and conflict detection mechanisms.

Several design choices merit discussion. The equal-weighting of portfolio holdings, while simple to implement, may not be optimal; risk-weighted or mean-variance-efficient allocation could further improve the Sharpe ratio. The 40-30-30 weight split was determined empirically; a data-driven weight calibration using cross-validated optimization on historical data is a natural extension. The use of VADER for sentiment analysis is computationally efficient but may miss nuanced financial language; transformer-based models such as FinBERT [15] could capture more complex sentiment structures.

A key limitation of the current system is its dependence on yfinance data, which has known issues with API rate limits and occasionally missing fundamental ratio fields. In a production deployment, institutional-grade data providers such as Bloomberg or Refinitiv would replace yfinance. Additionally, the backtesting engine does not currently model transaction costs, bid-ask spreads, or market impact, all of which would reduce realized returns in live trading.

VIII. CONCLUSION

This paper presented a novel multi-agent AI framework for investment strategy and risk governance, implemented as a fully functional Python system. Five specialized agents — Fundamental, Sentiment, Strategy, Risk, and Governance — collaborate through a weighted composite scoring architecture to produce BUY, HOLD, and SELL recommendations that are rigorously validated through backtesting. Experimental results on a 30-stock universe demonstrated a portfolio alpha of 8.5%, a Sharpe ratio of 1.84, and a maximum drawdown of 11.3%, each representing a meaningful improvement over a market-cap-weighted benchmark.

The work contributes to the growing body of literature on agentic AI applications in finance by explicitly addressing signal conflict, volatility-adjusted scoring, and explainable AI recommendations. The interactive Dash dashboard ensures that the system's outputs are accessible to human decision-makers who require transparency in AI-driven investment systems, positioning the proposed framework as a promising foundation for next-generation automated investment platforms.

IX.FUTURE WORK

Several directions are identified for extending the present work. First, the static weight vector in the composite scoring model will be replaced with a reinforcement learning (RL) agent that adapts weights dynamically based on recent prediction performance. Second, the Sentiment Agent will be enhanced by replacing VADER with a domain-adapted transformer model (FinBERT or BloombergGPT). Third, transaction cost modeling — including bid-ask spreads, brokerage commissions, and market impact — will be incorporated into the backtesting engine. Fourth, the system will be extended to support multi-asset class portfolios including bonds, commodities, and exchange-traded funds. Fifth, formal compliance and audit-trail features will be developed to meet the explainability requirements of financial regulators. Finally, the framework will be benchmarked against state-of-the-art deep reinforcement learning portfolio optimization baselines.

ACKNOWLEDGMENT

The authors sincerely thank the Department of Data Science at Bapatla Engineering College for providing the computational infrastructure and academic environment that supported this research. Gratitude is also extended to the open-source communities behind the Python, yfinance, VADER, Dash, and MongoDB projects, whose tools formed the backbone of the proposed system.

Conflict of interest statement

Authors declare that they do not have any conflict of interest.

REFERENCES

[1] J. Brogaard, T. Hendershott, and R. Riordan, “High-frequency trading and price discovery,” Rev. Financ. Stud., vol. 27, no. 8, pp. 2267–2306, Aug. 2014.

[2] M. Wooldridge, An Introduction to MultiAgent Systems, 2nd ed. Chichester, UK: Wiley, 2009.

[3] R. Cont, “Empirical properties of asset returns: stylized facts and statistical issues,” Quant. Finance, vol. 1, no. 2, pp. 223–236, 2001.

[4] G. Weiss, Ed., Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence. Cambridge, MA: MIT Press, 1999.

[5] R. Shiller, Irrational Exuberance, 3rd ed. Princeton, NJ: Princeton Univ. Press, 2015.

[6] S. Ding, C. Cui, and Z. Wu, “Multi-agent reinforcement learning for portfolio management,” in Proc. IJCAI Workshop on AI in Finance, 2020, pp. 45–52.

[7] B. Graham and D. Dodd, Security Analysis, 6th ed. New York, NY: McGraw-Hill, 2008.

[8] Plotly Technologies Inc., Dash User Guide. [Online]. Available: https://dash.plotly.com

[9] L. Tesfatsion and K. L. Judd, Eds., Handbook of Computational Economics, Vol. 2. Amsterdam: Elsevier, 2006.

[10] M. Wooldridge and N. R. Jennings, “Intelligent agents: theory and practice,” Knowl. Eng. Rev., vol. 10, no. 2, pp. 115–152, Jun. 1995.

[11] S. Chakraborty, A. Roy, and D. Bhattacharyya, “Multi-agent cooperative framework for stock portfolio optimization,” Expert Syst. Appl., vol. 122, pp. 181–194, 2019.

[12] K. Lim, W. Zhu, and Y. Yang, “Federated multi-agent portfolio management with risk-aware signal fusion,” IEEE Trans. Neural Netw. Learn. Syst., vol. 33, no. 7, pp. 3012–3025, Jul. 2022.

[13] P. C. Tetlock, “Giving content to investor sentiment,” J. Finance, vol. 62, no. 3, pp. 1139–1168, Jun. 2007.

[14] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, “Distributed representations of words and phrases,” in Proc. NIPS, 2013, pp. 3111–3119.

[15] D. Araci, “FinBERT: Financial sentiment analysis with pre-trained language models,” arXiv preprint arXiv:1908.10063, 2019.

[16] C. J. Hutto and E. Gilbert, “VADER: A parsimonious rule-based model for sentiment analysis,” in Proc. ICWSM, 2014, pp. 216–225.

[17] M. Hagenau, M. Liebmann, and D. Neumann, “Automated news reading: stock price prediction based on financial news,” Decis. Support Syst., vol. 55, no. 3, pp. 685–697, 2013.

[18] [18] J. J. Murphy, Technical Analysis of the Financial Markets. New York, NY: New York Institute of Finance, 1999.

[19] W. Brock, J. Lakonishok, and B. LeBaron, “Simple technical trading rules and the stochastic properties of stock returns,” J. Finance, vol. 47, no. 5, pp. 1731–1764, Dec. 1992.
[20] “Blockchain-Enabled Secure Data Aggregation for SDN-Enabled Ad-Hoc Networks,” International Journal of Intelligent Engineering and Systems, vol. 18, no. 5, pp. 704–717, Jun. 2025, doi: https://doi.org/10.22266/ijies2025.0630.49.
[21] J. Patel, S. Shah, P. Thakkar, and K. Kotecha, “Predicting stock and stock price index movement using machine learning,” Expert Syst. Appl., vol. 42, no. 1, pp. 259–268, 2015.
[22] P. Jorion, Value at Risk, 3rd ed. New York, NY: McGraw-Hill, 2006.
[23] H. Markowitz, “Portfolio selection,” J. Finance, vol. 7, no. 1, pp. 77–91, Mar. 1952.
[24] A. Ang, Asset Management: A Systematic Approach to Factor Investing. Oxford: Oxford Univ. Press, 2014.
[25] F. Chen, Y. Deng, and J. Wan, “Conflict resolution in multi-agent financial decision systems,” IEEE Access, vol. 9, pp. 88340–88352, 2021.
[26] E. Chan, Quantitative Trading. Hoboken, NJ: Wiley, 2008.
[27] W. F. Sharpe, “Mutual fund performance,” J. Business, vol. 39, no. 1, pp. 119–138, Jan. 1966.
[28] J. C. Hull, Options, Futures, and Other Derivatives, 10th ed. Hoboken, NJ: Pearson, 2017.
[29] R. Carver, Systematic Trading. Petersfield, UK: Harriman House, 2015.
[30] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” Nature, vol. 521, pp. 436–444, May 2015.
[31] OpenAI, “GPT-4 technical report,” arXiv preprint arXiv:2303.08774, 2023.