



# Leveraging AI for Personalized Financial Client Advisory Services: A Modular, Human-Centric Framework Based on Markowitz Portfolio Theory

Kaushik Khare<sup>1,\*</sup>, Manisha Dhage<sup>2</sup>, Eshan Deosatwar<sup>3</sup>, Yash Bankar<sup>4</sup>, Sahil Rodge<sup>5</sup>

<sup>1,\*</sup>Marathwada Mitra Mandal's College of Engineering, Karvenagar, Pune, Maharashtra, India 41105 (kaushikkhare26@gmail.com)

<sup>2</sup>Head of AI & Data Science Department, Marathwada Mitra Mandal's College of Engineering, Karvenagar, Pune, Maharashtra, India 41105 (kaushikkhare26@gmail.com)

<sup>3</sup>Marathwada Mitra Mandal's College of Engineering, Karvenagar, Pune, Maharashtra, India 41105 (deosatwareshan@gmail.com)

<sup>4</sup>Marathwada Mitra Mandal's College of Engineering, Karvenagar, Pune, Maharashtra, India 41105 (yashbankar789@gmail.com)

<sup>5</sup>Marathwada Mitra Mandal's College of Engineering, Karvenagar, Pune, Maharashtra, India 41105 (sahilrodge4@gmail.com)

**Abstract.** This research presents an AI-powered financial advisory system that delivers personalized portfolio recommendations from individual investor risk profiles. The system comprises three interconnected modules. The first identifies top-performing Nifty50 stocks through financial modelling; the second allocates capital across five asset classes; stocks, mutual funds, gold, cryptocurrency, and fixed deposits; guided by risk tolerance captured in a registration questionnaire; and the third manages risk by simulating market scenarios and applying trend-based strategies to limit losses. Outputs are presented in an interactive dashboard showing asset allocation, expected performance, and risk measures. The work demonstrates how combining AI with established financial principles yields personalized, accessible investment tools for novice and experienced investors alike.

**Keywords:** Artificial Intelligence · Financial Advisory System · Portfolio Optimization · Risk Management · Asset Allocation

## 1 Introduction

The financial advisory industry is shifting from advisor-dependent models toward AI-driven, scalable solutions. Traditional human consultants and inflexible digital tools struggle with scalability, human bias, slow market responsiveness, and limited personalization—weaknesses that intensify as financial ecosystems grow more dynamic and data-centric. Investor profiles vary widely across risk appetite, financial literacy, objectives, and life stage, demanding intelligent systems that continuously adapt to changing user needs and market conditions rather than applying standardized strategies.

© The Author(s) 2026

W. Srimaharaj (ed.), *Proceedings of the 2nd International Conference on Lifespan Innovation (ICLI 2026)*, Advances in Social Science, Education and Humanities Research 1035,

[https://doi.org/10.2991/978-94-6239-797-2\\_34](https://doi.org/10.2991/978-94-6239-797-2_34)

This research introduces a modular, AI-enabled advisory architecture for scalable personalization, built from three synergistic modules. The first applies the Black–Litterman methodology to determine optimal weight distributions across Nifty50 equities, ranking the top ten securities by fusing market-neutral baselines with subjective inputs, echoing explainable financial AI [4, 11]. The second uses Modern Portfolio Theory (MPT) to allocate capital across five asset classes—equities, mutual funds, cryptocurrencies, gold, and fixed deposits—according to a risk profile quantified through a structured behavioural questionnaire [12, 13]. The third minimizes risk via 100,000-iteration Monte Carlo simulation to locate optimal return–risk trade-offs [10] and a rule-based sell trigger activated when a stock's 50-day moving average falls below its 200-day counterpart [2, 21]. The system further evolves with user behaviour and external variables, integrating behavioural biometrics [13], hybrid decision-making [11], and life-phase modelling [12], while embedding algorithmic transparency [4], data protection [18], and fairness [15] for a responsible, regulation-aligned advisory experience [20, 22].

## 2 Methodology

This framework builds an AI-driven advisory system centred on portfolio optimization for Nifty50 equities, combining the Black–Litterman approach with Modern Portfolio Theory (MPT) [21]. It integrates financial indicators, quantified risk parameters, machine-learning logic, and optimization routines to produce high-Sharpe-ratio portfolios aligned to individual investor profiles [4, 11, 17].

### 2.1 Data Acquisition

Historical data for Nifty50 constituents was sourced via the yfinance Python package, covering adjusted closing prices from January 2020 to January 2025 for equities including Reliance Industries (RELIANCE.NS), Tata Consultancy Services (TCS.NS), and ICICI Bank (ICICIBANK.NS). Nifty50 index-level data was fetched concurrently as a market benchmark and neutral prior for Black–Litterman calibration.

### 2.2 Preprocessing and Return Engineering

Raw price series were converted to daily logarithmic returns:

$$\log \text{ return} = \ln(P_t / P_{t-1}) \quad (1)$$

where  $P_t$  and  $P_{t-1}$  are the closing prices at times  $t$  and  $t-1$ . These returns were annualised to estimate per-equity expected returns and volatilities, after cleaning missing values and synchronizing the time series across assets [19].

### 2.3 Black–Litterman Model Integration

The Black–Litterman model refined baseline expectations by adjusting equilibrium returns for investor sentiment, market risk premium, and the prevailing risk-free rate, incorporating volatility and inter-asset correlation matrices to sharpen asset-level expected returns [2]. Adjusted return per stock was computed as:

Adjusted Return = Base Return  $\times$  (1 + Correlation Adjustment)  $\times$  (1 – Volatility Adjustment)

To limit overexposure to tightly correlated assets, an average pairwise correlation matrix penalized strongly correlated assets through a scaling factor that reduced their effective weight; each asset's normalized volatility similarly down-weighted highly volatile stocks, containing systemic risk exposure.

### 2.4 Portfolio Optimization Strategy

The optimization engine combined Monte Carlo simulation with quadratic programming to find weightings maximizing the Sharpe ratio (excess return over volatility). Per-stock weights were constrained to between 1% and 20%, preserving diversification and avoiding concentration risk while targeting maximum risk-adjusted return.

### 2.5 Application of Modern Portfolio Theory

Modern Portfolio Theory (MPT), introduced by Harry Markowitz in 1952, constructs portfolios that maximize expected return for a given level of risk, or minimize risk for a target return [23]. Complementing Black–Litterman, MPT built an efficient frontier from the covariance matrix of historical returns, deriving weights that balance individual asset volatility against correlation structure.

The total risk of a portfolio is measured by variance:

$$\sigma_p^2 = \sum \sum w_i \times w_j \times \sigma_{ij} \quad (i, j = 1 \dots n) \quad (2)$$

where  $\sigma_p^2$  is portfolio variance,  $w_i$  and  $w_j$  are asset weights, and  $\sigma_{ij}$  the covariance between assets  $i$  and  $j$ .

The total portfolio expected return is given by:

$$E(R_p) = \sum w_i \times E(R_i) \quad (i = 1 \dots n) \quad (3)$$

Portfolio efficiency is assessed via the Sharpe Ratio:

$$\text{Sharpe Ratio} = (E(R_p) - R_f) / \sigma_p \quad (4)$$

where  $R_f$  is the risk-free rate of return. The optimization objective is:

$$\text{maximize } (w^T \mu - R_f) / \sqrt{(w^T \Sigma w)} \quad \text{subject to: } \sum w_i = 1, w_i \geq 0 \quad (5)$$

Final allocations followed investor-specific risk profiles captured through the behavioural questionnaire, aligning portfolios with personal financial goals and risk tolerance [23].

2.6 Efficient Frontier and Portfolio Metrics

Monte Carlo simulation visualised the efficient frontier—the set of portfolios offering the highest expected return for a given level of risk. Once optimised, expected return, portfolio volatility, and Sharpe ratio were computed to evaluate the optimization [23].




Fig. 1 System Architecture

3 System Implementation

The framework is realized across three integrated modules—algorithmic trading, portfolio management, and risk identification and mitigation—described in the following subsections.

3.1 Module I: Algorithmic Trading Recommendations

Module I couples algorithmic trading with Black–Litterman optimization to generate actionable recommendations for the Nifty 50 index, using historical data adjusted for risk and correlation to maximise returns while minimising risk [1, 3]. Here algorithmic trading optimises portfolio composition from the Black–Litterman output and issues real-time buy/sell/hold recommendations [2, 5].




Fig. 2 Correlation Heatmap

**Role of the Black–Litterman Model.** Adjusting market equilibrium returns, the model yields more reliable expected-return estimates by integrating investor views with volatility and correlation data, producing a more diversified, optimised portfolio [4, 6].

**Adjusted Returns and Portfolio Construction.** The algorithm weights stocks by risk-adjusted return, enforces diversification through correlation-balanced weights, and minimises exposure to high-volatility assets [1, 2].

**Buy-Sell-Hold Signal Generation.** For each constituent stock the algorithm emits Buy (strong expected return, low volatility), Sell (low return or high volatility), or Hold (moderate return, acceptable risk) signals [3, 5].




Fig. 3 Risk vs. Return Analysis




Fig. 4 Black–Litterman Adjusted Returns

**Weight Adjustments.** Allocations are continuously rebalanced against updated market conditions and Black–Litterman outputs, keeping the portfolio optimised in real time [1, 6].

**Risk Management and Position Sizing.** The algorithm maximises the Sharpe ratio, sizes each position by risk tolerance and expected return, and applies stop-loss and take-profit levels to bound downside risk [2, 4, 5].

**Backtesting and Evaluation.** Strategy performance is backtested on historical data against the Nifty 50 benchmark using Total Return, Maximum Drawdown, Sharpe Ratio, and Win Rate [6, 8].

Once optimised, the algorithm can deploy on a live trading platform via a brokerage API, transmitting orders directly to financial markets to execute trades in real time [1, 4].

3.2 Module II: Portfolio Management and Optimization

Module II enhances portfolio performance through quantitative techniques centred on the Black–Litterman model, which fuses market information with investor insight to build resilient allocations matched to client risk preferences and financial goals [4, 13].




Fig. 5 Portfolio Optimization

**Allocation and Efficient Frontier.** Built on MPT and enhanced by Black–Litterman, the module selects and adjusts assets for the best return within acceptable risk, first estimating equilibrium returns then adjusting them by the investor's absolute or relative views [3, 4, 5]. Mean-variance optimization sets allocations from the return covariance matrix [2, 6, 8], Black–Litterman adds flexibility [3, 6, 19], and the resulting efficient frontier, benchmarked against conventional approaches, matches portfolios to investor risk tolerance [4, 5].

**Optimization, Risk Management, and Deployment.** Performance is improved by maximising risk-adjusted measures such as the Sharpe ratio across market-wide and asset-specific risk [1, 2, 11], with risk mitigated through diversification, risk parity, and drawdown controls such as stop-loss orders and dynamic reallocation during volatility [3, 5, 6]. The model is backtested using Total Return, Maximum Drawdown, Sharpe Ratio, and Alpha/Beta against a benchmark [2, 4], then deployed with Black–Litterman-driven dynamic rebalancing as conditions change [5, 6, 19].

### 3.3 Module III: Risk Identification and Mitigation

Module III addresses predictive risk profiling, stress testing, and tail-risk analysis—techniques central to identifying and mitigating risk in volatile financial markets.




**Fig. 6** Monte Carlo Simulation

**Predictive Risk Profiling.** Machine-learning models anticipate asset-level risk, surfacing portfolio vulnerabilities; model validation ensures the accuracy and interpretability needed for confidence in AI-driven tools [5].

**AI in Fraud Detection and Mitigation.** AI flags suspicious activity with fewer false alarms and lower error rates than traditional methods, enabling faster, more precise fraud detection that strengthens portfolio security [4].

**ML-Enhanced Monte Carlo Stress Testing.** ML-augmented Monte Carlo simulation strengthens stress testing and tail-risk analysis by capturing complex patterns conventional models miss [3], across scenarios such as market downturns (crises, political instability, asset bubbles) and inflation shocks affecting fixed income, equity, and commodity markets.

Tail-risk analysis further prepares portfolios for rare but potentially devastating events, ensuring adequate mitigation measures are in place.

## 4 Results

The Black–Litterman model generated adjusted returns that were optimized into the portfolio with the highest Sharpe ratio. The final portfolio held top Nifty 50 stocks weighted by risk-return profile and adjusted for correlation and volatility, achieving an optimal balance of risk and return that informs personalized investment strategies. The results confirm that combining machine-learning techniques with traditional financial models supports dynamic, AI-driven advisory services for investors [8, 11, 22].

## 5 Discussion

Modular AI has transformed financial services by enabling scalable personalization. Discrete components handle specialized tasks—portfolio management, risk assessment, client interaction—and collaboratively process large datasets, letting advisors and institutions deliver customized services efficiently at scale [1, 3, 13].

**Advantages in Adaptability, Transparency, and Inclusivity.** Modular systems adapt quickly to market shifts, regulatory updates, and changing client preferences [2]; their independently observable modules simplify auditing of decision-making [4]; and they democratise access by serving segments from high-net-worth to smaller investors [5, 15].

**Challenges in Trust, Interpretability, and Regulation.** Clients must trust that recommendations serve their best interests [6], yet model complexity often obscures decision rationale [7], and the growing role of AI in finance requires regulators to enforce ethical and legal standards covering data privacy, algorithmic bias, and fairness [8].

## 6 Challenges and Limitations

Several challenges and limitations must be addressed for effective real-world deployment of modular AI in financial environments.

**Data Security and Explainability.** Sensitive financial data demands strict regulatory compliance and advanced cybersecurity [9], while deep-learning “black boxes” limit explainability, eroding client trust and slowing adoption [10].

**Algorithmic Overreach and Client Over-Reliance.** Treating AI as an autonomous decision-maker rather than a decision-support tool [11] risks flawed choices in volatile or unprecedented conditions; balanced human oversight remains essential [12].

**Model Overfitting and Dynamic Market Behaviour.** Models risk overfitting historical data and misreading evolving market dynamics [13]; continuous retraining is required to track shifting market regimes [14].

**Cost and Complexity of Real-World Deployment.** Integrating AI with legacy infrastructure, ensuring compliance, and securing data are costly [15], and ongoing maintenance, monitoring, and upgrades demand sustained investment in infrastructure and skilled personnel [16].

## 7 Future Work

As AI continues to advance in finance, several avenues merit further research and development.

**Life-Event-Based Self-Adjusting Agents.** Agents that dynamically adapt strategies to milestones—marriage, childbirth, retirement, career change—would enable more proactive, relevant planning [17].

**Integration with Web3 for Decentralised Finance.** Pairing modular AI with blockchain could strengthen transparency, security, and trust in decentralised finance by uniting immutable ledgers with predictive analytics [18].

**Enhanced Human-AI Collaboration for Emotional Support.** Advances in natural language processing and affective computing could let systems offer real-time emotional support during stressful decisions, easing client anxiety through empathetic, personalised guidance [19].

**Extended Personalization for Neurodiverse Clients.** Tailored communication and decision support aligned with autism spectrum or ADHD profiles would broaden inclusivity and accessibility [20].

## 8 Conclusion

This paper examined modular AI across three domains—algorithmic trading, portfolio management, and risk identification and mitigation—demonstrating its potential to deliver personalised, scalable, and efficient financial services through predictive analytics, continuous strategy adaptation, and proactive risk forecasting. Challenges persist in trust, interpretability, data security, and regulatory compliance, since model complexity can undermine transparency; explainability and rigorous compliance are therefore vital to mainstream adoption. With advances such as life-event-based agents, Web3-enabled decentralised finance, human-AI emotional support, and solutions for

neurodiverse clients, modular AI represents a paradigm shift poised to make financial advisory more inclusive, transparent, and adaptive.

## References

1. Al-Dhamari, A., Rahman, M., Zhang, L.: Systematic review of AI techniques in financial trading markets. *IEEE Access* 11, 1–20 (2023)
2. Li, Y., Gao, J., Chen, T.: Reinforcement learning for dynamic trading strategies. In: *Proc. IEEE Conf. Computational Intelligence for Financial Engineering & Economics* (2020)
3. Ta, Q., Liu, L., Nguyen, M.: AI-driven market sentiment analysis for trading. *J. Comput. Finance* 22(4), 55–74 (2018)
4. *Frontiers in Artificial Intelligence: Explainable AI for portfolio construction*. *Front. Artif. Intell.* 7 (2024)
5. Bartram, S., Wilkens, D., Taylor, A.: Multi-objective evolutionary algorithms for asset allocation. *Eur. J. Oper. Res.* 291(2), 670–685 (2021)
6. Nuzzo, A., Morone, A.: Deep learning for risk-adjusted portfolio returns. *Quant. Finance Lett.* 6, 81–92 (2017)
7. CFA Institute: Real-time portfolio rebalancing with AI. *CFA Institute Research Reports* (2020)
8. Neupane, B., Sharma, R.: AI model risk management in financial institutions. *J. Risk Manage. AI Syst.* 14(1), 34–47 (2025)
9. Forbes Finance Council: Predictive analytics for fraud detection. *Forbes* (March 2024). <https://www.forbes.com>
10. ScienceDirect: AI-enhanced stress testing for market shocks. *Appl. Artif. Intell. Finance* 16(2), 101–118 (2024)
11. Nguyen, M.: AI-powered financial advisory services. *J. Financ. Innov.* 12, 221–238 (2024)
12. Tres Astronauts: Goal-based financial planning with AI. *J. Comput. Wealth Manage.* 5(1), 29–43 (2024)
13. World Economic Forum: Behavioural biometrics for client risk profiling. *WEF White Paper* (January 2025)
14. *Journal of AI-Assisted Scientific Discovery: Low-cost robo-advisors for mass market adoption*. *JAASD* 3(2), 76–90 (2024)
15. *Frontiers in AI: Bias mitigation in AI-driven advice*. *Front. AI* 5(3), 149–160 (2024)
16. World Economic Forum: Quantum computing for portfolio optimization. *WEF Briefing Paper* (February 2025)
17. Chartis Research: Federated learning for privacy-preserving data analysis. *Chartis Insights Report* (March 2025)
18. Forbes: Generative AI for synthetic financial scenarios. *Forbes* (April 2024). <https://www.forbes.com>
19. World Economic Forum: AI adoption in banking: a 10-year outlook. *WEF Foresight Series* (February 2025)
20. Chartis Research: Regulatory challenges in AI advisory services. *Chartis Research Reports* (2025)
21. Bartram, S., Harford, R., Jin, T.: Machine learning in financial markets: a review. *J. Financ. Data Sci.* 2(1), 1–23 (2020)
22. ScienceDirect: Ethical AI in wealth management. *AI Ethics Financ. Plan.* 13(2), 87–102 (2024)
23. Markowitz, H.: Portfolio selection. *J. Finance* 7(1), 77–91 (1952)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

