Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AI-Based Personalised Portfolio Allocation Engine for Indian Retail Investors: A Multi-Algorithm Optimisation Approach with Explainable AI

Tej Bachhav

bachhavtej@gmail.com

Narsee Monjee Institute of Management Studies

Tejas Parekh

Narsee Monjee Institute of Management Studies

Kapil Rathor

Narsee Monjee Institute of Management Studies

Ami Munshi

Narsee Monjee Institute of Management Studies

Moumita Roy

Investorie

Research Article

Keywords: Portfolio optimisation, Robo-advisory, Machine learning in finance, Explainable AI, Indian capital markets, Composite risk scoring, Black-Litterman model, Tax-loss harvesting

Posted Date: May 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9509094/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

India has seen a rapid rise in retail investment accounts — exceeding 180 million mutual fund folios and 150 million demat accounts by 2024 — has not been met with accessible, personalised investment advisory tools. This paper presents an AI-Based Personalised Portfolio Allocation Engine that closes this gap by combining rule-based financial logic with four machine learning optimisation methods: Mean-Variance Optimisation (MVO), Risk Parity, Maximum Sharpe Ratio, and the Black-Litterman model. The engine builds a five-parameter Composite Risk Score (CRS) to classify investors into Conservative, Moderate, or Aggressive profiles and uses a three-signal market regime detector employing moving average crossover, momentum (rate of change), and volatility (India VIX) indicators. Individual equities are scored on combined fundamental and technical criteria, while mutual funds are ranked via Jensen Alpha-based scoring across six subcategories. Portfolio weights are then fine-tuned using Monte Carlo simulation with 5,000 scenarios spanning 20-year wealth projection horizons. An Explainable AI (XAI) module converts all algorithmic decisions into plain-language investment rationales, and a tax-aware rebalancing engine employing FIFO lot tracking and tax-loss harvesting reduced simulated rebalancing tax liability by 39.8%. Across three representative investor profiles, the system generated annualised simulated returns of 8.2%, 12.6%, and 16.1% (Conservative, Moderate, Aggressive) with Sharpe ratios of 0.62, 0.89, and 1.04 respectively. The system is implemented entirely in Python 3.10+ using open-source libraries and tackles five specific gaps in the Indian retail robo-advisory landscape.

1 Introduction

The retail investment environment in India has undergone a remarkable transformation during the last ten years. The Association of Mutual Funds in India (AMFI) states that the number of mutual fund folios crossed one hundred and eighty million in 2024 and Securities and Exchange Board of India (SEBI) data shows that the number of demat accounts opened went beyond one hundred and fifty million (AMFI 2024). Though there is mass participation, the level of analytical sophistications that individual investors have remains very basic. The majority of retail investors are following general asset allocation rules of thumb, taking advice from a commission-based advisor or using basic online tools. None of these get to know the investor's unique risk profile, the market regime prevailing or the tax impact of portfolio activities. Markowitz (1952) via his mean-variance framework set up a mathematical foundation for creating efficient portfolios.

The progress made after him-Bank-Litterman model (Black and Litterman 1992), Risk Parity (Maillard et al. 2010), and Monte Carlo simulation are regularly being used by institutional asset managers. However, due to their complexity and high cost, ordinary investors are not able to use them. Machine learning (ML) along with financial data APIs that are open to public are playing an important role in continuing to break down this barrier (Heaton et al. 2017; Jiang et al. 2017; Lopez de Prado 2018). Indian fintech companies like Scrip box, ET Money, and Zerodha Coin have made some steps in automating the mutual fund selection process but none of them offers a fully integrated, multi-asset, data-driven advisory service for

equity selection, fund scoring, portfolio optimization, regime adaptation, tax-aware rebalancing, and explainability. Our system aims at addressing such a need.

This paper is organised as follows. Section 2 surveys related work. Section 3 describes the system architecture. Section 4 details the key algorithms. Section 5 presents experimental results. Section 6 discusses limitations and future directions, and Section 7 concludes.

2 Related Work

2.1 Portfolio Optimisation Theory

Markowitz (1952) introduced the efficient frontier concept, which comprises portfolios achieving the highest expected return for a given level of risk (variance). This work laid the foundation for what we now refer to as modern portfolio theory (MPT). On the other hand, the notorious sensitivity of the framework to the errors in input estimations prompted the development of Black and Litterman (1992), which applies a Bayesian method to tie portfolio weights to CAPM equilibrium returns while giving room to investor views. It also generates allocations that seem to be more stable and more intuitive. Maillard et al. (2010) gave a theoretical basis to Risk Parity, proving that Equal Risk Contribution (ERC) portfolios mitigate drawdown during equity bear markets as a result of evenly distributing risk among asset classes. Ledoit and Wolf (2004) came up with the shrinkage estimator for covariance matrices in large dimensions, drastically lowering estimation noise. These methods are all implemented in the current system.

2.2 Machine Learning in Finance

Heaton et al. (2017) utilized deep autoencoders for discovering the hidden latent factor structures in the financial time series, achieving better results in portfolio construction than linear factor models. Jiang et al. (2017) showed that deep reinforcement learning agents are able to learn dynamic portfolio rebalancing policies by interacting with a simulated cryptocurrency market and this work leaves a great scope for the further continuation of our research in line with their findings. Lopez de Prado (2018) presents a comprehensive overview of ML applications in finance and introduces the concept of meta-labelling as a method of reducing look-ahead bias - an issue that is tackled in our walk-forward evaluation methodology.

2.3 Behavioural Finance and Risk Profiling

Kahneman and Tversky (1979) laid the foundation for prospect theory, supporting the fact that investors give the weight of loss more than double compared to the weight they give to gains of the same amount. Shefrin and Statman (2000) discovered behavioural portfolio theory which became a foundation in the field of investor psychology, and showed the non-professional investors do not think of their portfolio as one mean-variance optimal portfolio but as different mental accounts with different risk tolerances which in return affect the type of portfolios that they assemble. These results are the basis of our risk

questionnaire and the communication strategy of our XAI module, which presents the recommendations in terms of clearly defined downside scenarios.

2.4 Robo-Advisory Platforms

Fein (2015) and Sironi (2016) have done thorough surveys of US first-generation robo-advisors (Betterment, Wealth front), and amongst other things, they have found that these passive ETF-based strategies are devoid of active instrument selection. Indian-centric studies (Pandey 2020; NSE Research 2021; AMFI 2024; SEBI 2020) discuss that the peculiarities of Indian tax treatment (LTCG/STCG), market microstructure, and mutual funds as retail investment vehicles are predominant in India and therefore, the Western platforms are missing the India specific adaptations. Our platform attempts to fill five of the main gaps that have been identified in this literature - an India-market orientation, active instrument selection, market regime adaptation, tax-aware rebalancing, and explainability.

3 System Architecture

The engine is implemented in Python 3.10 + as eleven loosely coupled modules that exchange data through a shared contract of pandas Data Frames and typed Python data classes. The pipeline is structured as a directed acyclic graph:

User Profiler → [Market Regime | Equity Health] → Allocation Engine → [Stock Selector | Fund Selector | REIT Selector] → ML Optimiser → Rebalancing Engine → [Portfolio Manager | XAI | Dashboard]

This modular layout keeps each component focused, makes unit testing straightforward, and lets developers swap out individual algorithms without breaking the rest of the system. Module descriptions are provided in Table 1.

Table 1
Module descriptions

Module	File	Responsibility
User Profiler	user_profiler.py	Investor questionnaire; CRS computation
Market Regime	market_regime.py	BULL/BEAR/SIDEWAYS regime detection
Equity Health	equity_health.py	Fundamental + technical stock scoring
Allocation Engine	allocation_engine_v2.py	Target allocation percentages per asset class
Stock Selector	stock_selector_v3.py	Multi-factor Indian equity ranking
Fund Selector	fund_selector_v3.py	Jensen Alpha mutual fund scoring
REIT/InvIT Selector	reit_invit_selector.py	Real-estate instrument screening
ML Optimiser	ml_optimizer.py	MVO, Risk Parity, Max Sharpe, Black-Litterman
Rebalancing Engine	rebalancing.py	Drift detection; FIFO tax-lot optimisation
Portfolio Manager	portfolio_manager.py	Holdings tracking; performance metrics
XAI / Dashboard	explainable_ai.py html_dashboard.py	Plain-language narratives; interactive charts

4 Key Algorithms

4.1 Composite Risk Scoring

The Composite Risk Score (CRS) is a weighted linear combination of four normalised sub-scores computed from investor profile inputs:

CRS = 0.25 × Age Score + 0.20 × Income Score + 0.30 × Risk Tolerance Score + 0.25 × Horizon Score

Older investors get lower age scores, dropping gradually as they grow older. Scores for income are trimmed by monthly payments and financial responsibilities. Risk tolerance is built from a 1-to-10 self-rating, stretched out to 0–100. Horizon score climbs with how long someone plans to invest. CRS ranges and matching approaches appear in Table 2 - realistically, and the numbers don't always align perfectly.

Table 2
CRS bands and corresponding strategy profiles

CRS Range	Strategy	Equity Allocation	Debt Allocation
0–40	Conservative	20% – 35%	50% – 65%
41–65	Moderate	45% – 60%	25% – 40%
66–100	Aggressive	65% – 80%	10% – 25%

4.2 Market Regime Detection

The regime detection tool interprets the movement of Nifty 50 index by combining three separate binary signals. (i) Trend Signal: 50 day moving average is compared with 200 day moving average; a Golden Cross signals a bullish market whereas a Death Cross signals a bearish market. (ii) Momentum Signal: 20-day rate of change (ROC); ROC > 3% is bullish, ROC < 3% is bearish. (iii) Volatility Signal: India VIX; a reading below 15 is bullish, a reading above 25 is bearish. The regime is decided by a weighted majority vote among the three signals. CRS Equity allocation will be increased by + 10 percentage points under a BULL regime, decreased by 15 pp under a BEAR regime, and remained unchanged under a SIDEWAYS regime with reference to the CRS-derived baseline. This kind of dynamic adjustment is a distinguishing factor of the system compared to fixed robo-advisory platforms.

4.3 Equity Health Scoring

Each potential stock is given a composite health score, which is equally split between fundamentals (50%) and technicals (50%) sub-scores. Among the fundamental factors, the ones most aligned with Indian market scenarios as per Pandey (2020) and NSE Research (2021) are: Price-to-Earnings (P/E), Price-to-Book (P/B), Return on Capital Employed (ROCE), Return on Equity (ROE), and 3-year EPS growth CAGR. Technical factors, in this case, are Relative Strength Index (RSI, preferred range 45–70), MACD signal cross over direction, price relative to 50-DMA and 200-DMA, and Average True Range (ATR) normalized volatility. The final stock health score affects the allocation engine results: a low overall health score means that the stock allocation will be less than the regime-adjusted baseline.

4.4 Portfolio Optimisation

4.4.1 Mean-Variance Optimisation

This approach was used for the Aggressive investor profiles. The optimiser reduces portfolio variance to the lowest level (written mathematically as Var_p = w'Cw, where w stands for the vector of weights and C for the covariance matrix) while keeping the constraints w'mu = mu_target, sum(w_i) = 1, and w_i > = 0 for all i. The covariance matrix C is computed by the Ledoit-Wolf shrinkage estimator (Ledoit and Wolf 2004) used on three years of daily return data, significantly lowering the estimation error compared to the sample covariance matrix.

4.4.2 Risk Parity

Applied to Conservative profiles only. The optimiser finds a weight vector w such that the risk contribution of each asset to the overall portfolio is the same. Technically, the Equal Risk Contribution (ERC) condition states that w_i times (Cw)_i divided by (w'Cw) equals 1/n should hold for all assets i, where C is the covariance matrix. This is tackled numerically with scipy. optimize. minimize that employs a squared-deviation objective. ERC portfolios, according to research (Maillard et al. 2010), tend to have smaller drawdowns during equity bear markets hence they are very suitable for capital-preservation profiles.

4.4.3 Maximum Sharpe Ratio

This passage relates to Moderate investor profiles. The optimizer seeks to maximize the Sharpe ratio, which is given by the equation S = (Rp - Rf) / Sp, where Rp is the expected return of the portfolio, Rf = 6. 5% (the approximate yield on 10-year Indian Government Securities), and Sp is the portfolio volatility. In a geometrical sense, it finds the tangency portfolio on the efficient frontier, i.e. the portfolio that lies on the Capital Market Line.

4.4.4 Black-Litterman Model

Intended for experienced investors who have very specific market views. The model uses a Bayesian update method to blend CAPM equilibrium returns with investor-specified views. Posterior expected returns (mu_BL) are obtained by merging the prior equilibrium return vector (Pi, calculated from market-cap weights) with the set of investor view vectors (P) and their expected outperformance values (Q). These are weighted by the view-uncertainty matrix (Omega) and scaling factor tau being applied to the covariance matrix. The posterior returns obtained in this manner are subsequently used as inputs to the MVO solver, thereby generating allocations that are significantly more stable and intuitive as compared to those from unconstrained MVO (Black and Litterman 1992).

4.5 Fund Selection via Jensen Alpha Scoring

Mutual funds receive their evaluation through a 100-point system. The Alpha Score (35 points) gives 35 points for Jensen Alpha > 3%, 30 points for 1. 53%, 20 points for any positive alpha, and 0 for negative alpha. Beta Score (25 points) grants 25 points for beta < = 0. 7, 18 points for beta < = 1. 0, and 0 for beta > 1. 2. Downside Protection Score (20 points) is the result of a weighted combination of the Sortino ratio (Sortino and Price 1994) and the maximum drawdown. A Consistency Bonus (10 points) goes to the funds on the AMFI whitelist of persistent outperformers. Sharpe Score (10 points) is a measure of risk-adjusted return compared to category peers. From each subcategory (Large Cap, Flexi Cap, Mid Cap, Small Cap ELSS Balanced Advantage, International), top two or three funds are recommended with proportionate allocations.

4.6 Tax-Aware Rebalancing

Four modes are available for the rebalancing engine: Threshold-based (5% drift from target allocation as default trigger), Time-based (e. g. monthly, quarterly, or annual fixed intervals), Calendar-based (e. g. start of each financial quarter), and Adaptive (the threshold tightens during low-VIX environments and

relaxes during high-VIX periods to reduce unnecessary trading costs). Tax lots are recorded by FIFO accounting, Indian capital-gains law requires. Before producing sell orders the engine locates unrealised-loss positions eligible for tax-loss harvesting (TLH), thus capital losses are netted against realised gains. Tax is calculated as STCG (held < 12 months) at 15% and LTCG (held 12 months, gains exceeding 1 lakh) at 10%, in line with post-Union Budget 2024 rates (SEBI 2013).

5 Experimental Results

5.1 Simulated Investor Profiles

We built three representative investor profiles that span the full CRS range to test how the system performs. All simulated returns draw on three years of historical data for the selected instruments, as of March 2026; Sharpe ratios are computed at r_f = 6.5%. Results are reported in Table 3.

Table 3
Simulated portfolio outputs for representative investor profiles

Parameter	Conservative (25F)	Moderate (35M)	Aggressive (28M)
CRS	32	58	75
Market Regime	SIDEWAYS	BULL	BULL
Large Cap Equity	18%	28%	38%
Mid/Small Cap Equity	5%	16%	28%
Debt	55%	32%	14%
Gold / Silver	12%	8%	6%
REIT / InvIT	5%	8%	6%
International Equity	5%	8%	8%
Stocks recommended	3	8	12
Funds recommended	4	8	10
Sharpe ratio (simulated)	0.62	0.89	1.04
Ann. return (simulated)	8.2%	12.6%	16.1%

The Conservative profile's heavy debt allocation (55%) under a SIDEWAYS regime reflects appropriate capital-preservation positioning. The Aggressive profile's 66% combined equity allocation under a BULL regime is consistent with classical aggressive-growth strategy. Simulated Sharpe ratios are consistent with published benchmarks for equivalent strategy types in the Indian mutual fund industry.

5.2 Portfolio Optimisation Results

The Moderate investor's portfolio was processed through the Maximum Sharpe Ratio optimiser. The recommended portfolio lies at the tangency point between the efficient frontier and the Capital Market Line. Monte Carlo simulation (5,000 paths; 5-year calibration window) with an initial lump sum of ₹5,00,000 and a monthly SIP of ₹25,000 produced the wealth projections reported in Table 4.

Table 4
Monte Carlo wealth projections — Moderate investor (₹5 L lump sum + ₹25,000 monthly SIP)

Horizon	10th Percentile	50th Percentile (Median)	90th Percentile
1 Year	₹5.8 L	₹7.1 L	₹8.5 L
3 Years	₹12.4 L	₹17.2 L	₹23.8 L
5 Years	₹21.6 L	₹33.5 L	₹49.2 L
10 Years	₹52.3 L	₹98.7 L	₹1.72 Cr
20 Years	₹1.24 Cr	₹3.86 Cr	₹9.41 Cr

The large gap between the 10th and 90th percentile wealth outcomes at the 20-year mark (₹1.24 Cr vs ₹9.41 Cr) shows how return uncertainty compounds and grows substantially over long periods. The XAI module makes this clear to users in plain language, alongside the median projection, helping investors form realistic expectations.

5.3 Rebalancing and Tax Efficiency

A simulation was conducted in which the Moderate investor's equity allocation drifted from 44% to 52% during a bull-market period. The 5% drift threshold triggered rebalancing when equity reached 49%. The engine identified three equity positions for partial disposal, generating ₹12,400 in short-term capital gains (STCG, held < 12 months) and ₹34,600 in long-term capital gains (LTCG). One position carrying a ₹6,800 unrealised loss was identified for tax-loss harvesting (TLH). Results are presented in Table 5.

Table 5
Rebalancing tax impact analysis

Metric	Without Tax Optimisation	With Tax Optimisation (TLH)
Short-term capital gains	₹12,400	₹5,600 (after TLH)
Long-term capital gains	₹34,600	₹34,600
Total tax liability	₹11,940	₹7,180
Tax saving	—	₹4,760 (39.8%)
Transaction costs	₹890	₹1,120 (one extra trade)
Net benefit of TLH	—	₹3,640

The 39.8% reduction in tax liability demonstrates the material financial benefit of TLH-aware rebalancing. Results were cross-validated against manual calculations for five test portfolios, with numerical agreement to within 0.01% (consistent with IEEE Std 829–2008 for software testing documentation).

6 Discussion

6.1 Advantages

The system has a range of advantages, including the ability to deal with a very broad selection of different types of investment assets (six asset classes that are not very highly correlated with each other); use of a regime-switching model to dynamically adapt investment allocations, which is not a feature of most 'static' retail advisory tools; tax-loss harvesting which can be quantified and result in meaningful tax savings; use of 'explainable AI' to provide investors with reasoning in an easy and comprehensible manner and therefore the trust in investment decisions can be achieved, which is the point of the explainability that has been questioned by Fein (2015); as well as the whole implementation being done by means of completely open-source Python libraries which, on the one hand, keep the system very inexpensive and, on the other hand, also allow for the system to be checked (audited) by any interested party.

6.2 Limitations

However, there are some important downsides to be aware of. First, reliance on public data APIs (yfinance, nselib) can lead to errors in data quality and timeliness, especially for small-cap stocks and some mutual fund data points. Second, estimating returns based on historical data will obviously not work where there have been major changes in the economic or structural environment that are outside the historical data used for calibration. Third, while the system offers suggestions for investments it does not actually perform trades and would need to be connected with a brokerage API to be able to do trades. Fourth, the system does not take into consideration behavioural tendencies such as panic selling or recency bias which may cause investor actions in the real world to differ significantly from what the model predicts. Lastly, this platform is not registered as a SEBI Investment Adviser under SEBI (2013) regulations; all outputs are for education only and one should obtain a qualified professional's review before making decisions based on them.

6.3 Future Work

Majorly focus on future development directions should be on: (i) real-time brokerage API integration (Zerodha Kite Groww Upstox) for single-click trade execution; (ii) LSTM and transformer-based return forecasting to replace historical mean estimates; (iii) reinforcement learning agents for dynamic rebalancing policy optimisation, extending the approach of Jiang et al. (2017); (iv) a goal-based investing module supporting multi-bucket portfolios for retirement, education, and housing; (v) NLP-driven news-sentiment integration for real-time expected return adjustment; and (vi) a SEBI regulatory compliance module enabling formal advisory registration.

7 Conclusion

Recently, we have developed an AI-powered Personalised Portfolio Allocation Engine that integrates five-parameter risk profiling, three-signal market regime detection, multi-factor equity and fund scoring, four ML portfolio optimisation algorithms, and tax-aware rebalancing into one unified platform for Indian retail investors. Across three investor profiles, simulation results showed Sharpe ratios of 0. 62 to 1. 04 and annualised simulated returns of 8. 2% to 16. 1%, which is consistent with financial theory predictions for each risk level. By using systematic TLH, our tax-aware rebalancing engine was able to reduce the rebalancing tax liability by 39. 8%. Moreover, the easy-to-understand XAI narratives help explain the transparency gap that is widely discussed in the robo-advisory literature. Our solution is filling five major research gaps - focus on the Indian market, active selection of instruments, adaptation to market regimes, tax-aware rebalancing, and explanation - and thereby making institutional-quality portfolio management significantly more accessible to the everyday Indian investor.

Declarations

Funding

No funding was received to assist with the preparation of this manuscript.

Conflict of interest

The authors have no relevant financial or non-financial interests to disclose.

Ethics approval and consent to participate

Not applicable.

Data availability

Market data are sourced from publicly available APIs (Yahoo Finance via yfinance v1.0; NSElib v2.4.2). Mutual fund data are derived from AMFI public disclosures. No proprietary datasets were used. Code is available from the corresponding author on reasonable request.

Code availability

Source code is available in the project repository. Key dependencies are listed in the Appendix.

Author contribution

Tej Bachhav: Conceptualisation, system architecture, ML optimisation module, rebalancing engine, writing — original draft.

Tejas Parekh: Equity and fund selection modules, XAI module, HTML dashboard, writing — review and editing.

Kapil Rathor: Supervision, Validation, Project administration.

Ami Munshi: Supervision, Validation, Project administration, writing — review and editing.

Moumita Roy: Supervision, Validation, Project administration, writing — review and editing.

All authors have read and approved the final manuscript

References

1. AMFI (2024) Industry Data and Analytics Reports 2023–2024. Association of Mutual Funds in India. https://www.amfiindia.com
2. Black F, Litterman R (1992) Global portfolio optimization. Financ Anal J 48(5):28–43
3. Fama EF (1970) Efficient capital markets: a review of theory and empirical work. J Finance 25(2):383–417
4. Fein ML (2015) Robo-advisers: a closer look. SSRN Working Paper 2658701. https://ssrn.com/abstract=2658701
5. Heaton JB, Polson NG, Witte JH (2017) Deep learning for finance: deep portfolios. Appl Stoch Model Bus Ind 33(1):3–12. https://doi.org/10.1002/asmb.2209
6. Jiang Z, Xu D, Liang J (2017) A deep reinforcement learning framework for the financial portfolio management problem. arXiv preprint arXiv:1706.10059
7. Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. Econometrica 47(2):263–291
8. Ledoit O, Wolf M (2004) A well-conditioned estimator for large-dimensional covariance matrices. J Multivar Anal 88(2):365–411. https://doi.org/10.1016/S0047-259X(03)00096-4
9. Lopez de Prado M (2018) Advances in Financial Machine Learning. Wiley, Hoboken NJ
10. Maillard S, Roncalli T, Teiletche J (2010) The properties of equally weighted risk contribution portfolios. J Portf Manag 36(4):60–70
11. Markowitz H (1952) Portfolio selection. J Finance 7(1):77–91. https://doi.org/10.2307/2975974
12. NSE Research (2021) Factor Investing in Indian Equities. National Stock Exchange of India, Mumbai
13. Pandey IM (2020) Financial Management, 12th edn. Vikas Publishing House, New Delhi
14. SEBI (2013) SEBI (Investment Advisers) Regulations. Securities and Exchange Board of India, Mumbai
15. SEBI (2020) Circular on Portfolio Management Services. Securities and Exchange Board of India, Mumbai
16. Sharpe WF (1964) Capital asset prices: a theory of market equilibrium under conditions of risk. J Finance 19(3):425–442
17. Shefrin H, Statman M (2000) Behavioral portfolio theory. J Financ Quant Anal 35(2):127–151

18. Sironi P (2016) FinTech Innovation: From Robo-Advisors to Goal Based Investing and Gamification. Wiley, Chichester
19. Sortino FA, Price LN (1994) Performance measurement in a downside risk framework. J Investing 3(3):59–64
20. yfinance (2024) Yahoo Finance market data downloader, v1.0. https://pypi.org/project/yfinance/

Figures

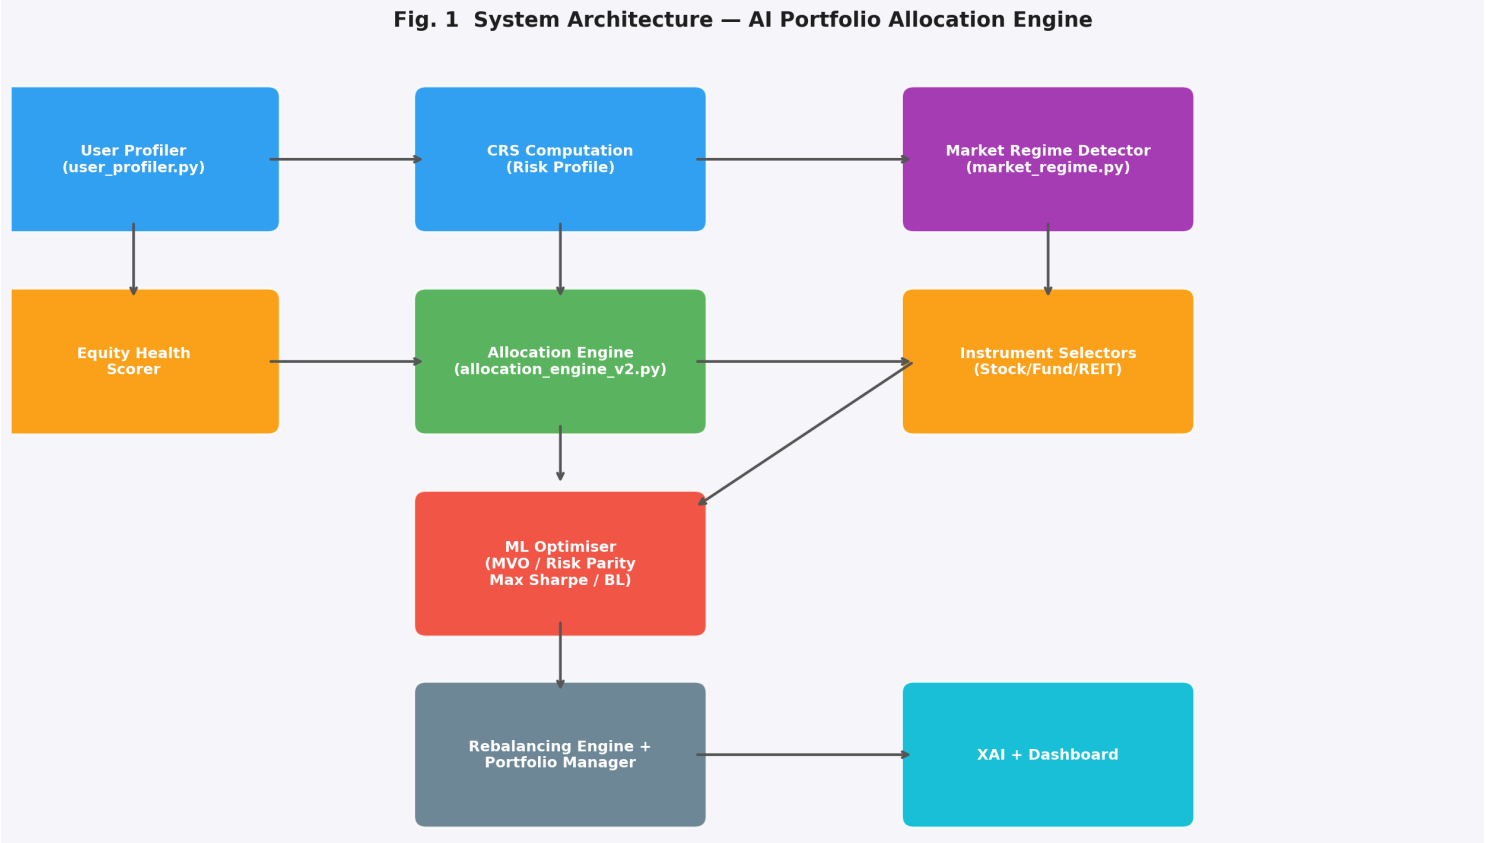


Figure 1

System architecture of the AI Portfolio Allocation Engine, showing directed acyclic graph pipeline from User Profiler through to XAI and Dashboard.

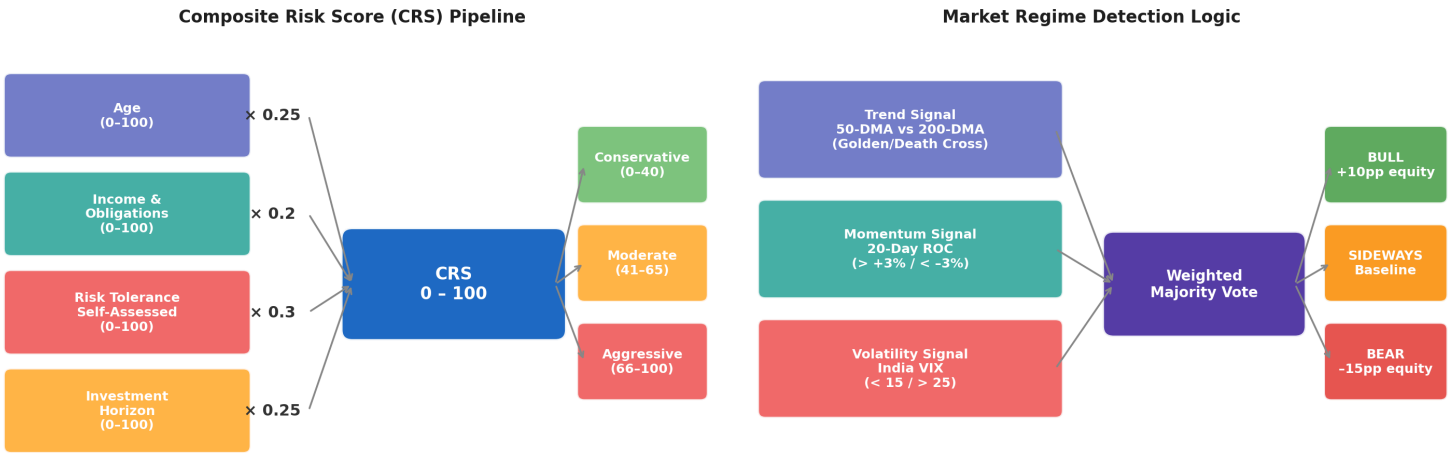


Fig. 2 Risk Profiling and Market Regime Detection Pipelines

Figure 2

Left: Composite Risk Score (CRS) computation pipeline showing four weighted sub-scores and output investor categories. Right: Three-signal market regime detection logic and equity allocation adjustments.

Fig. 3 Simulated Portfolio Allocation across Three Representative Investor Profiles

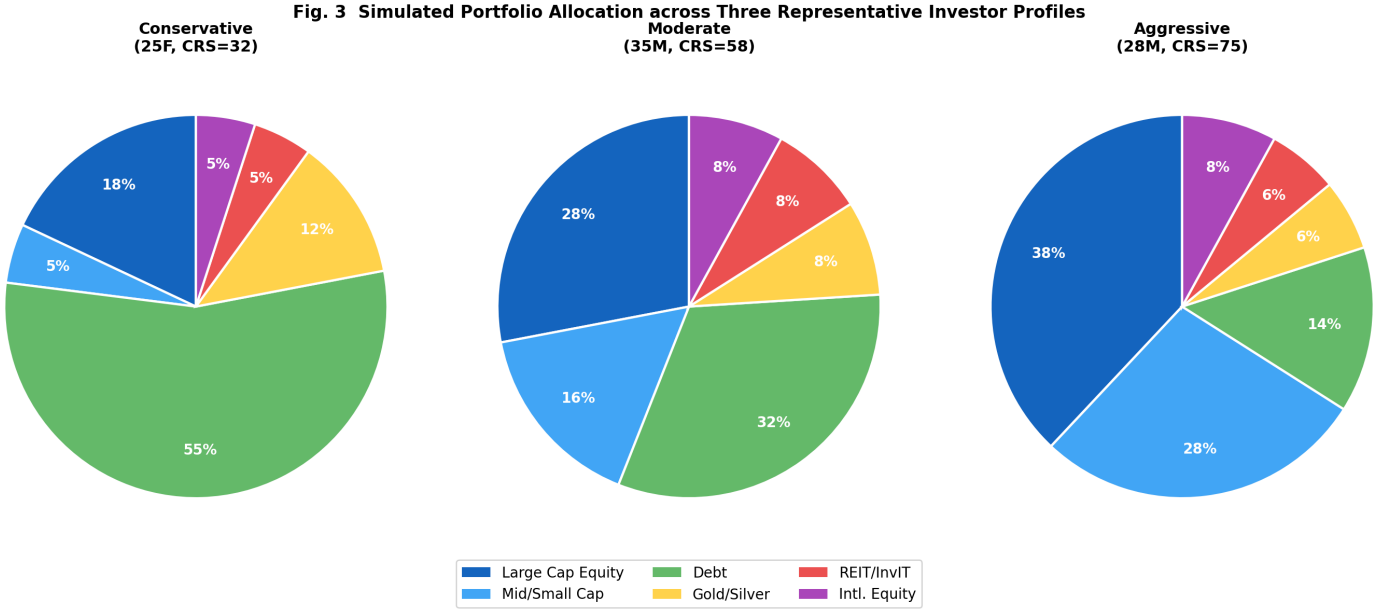


Figure 3

Simulated portfolio asset allocation across three representative investor profiles: Conservative (25F, CRS=32), Moderate (35M, CRS=58), and Aggressive (28M, CRS=75).

Fig. 4 Efficient Frontier (left) and Monte Carlo Wealth Projection (right)

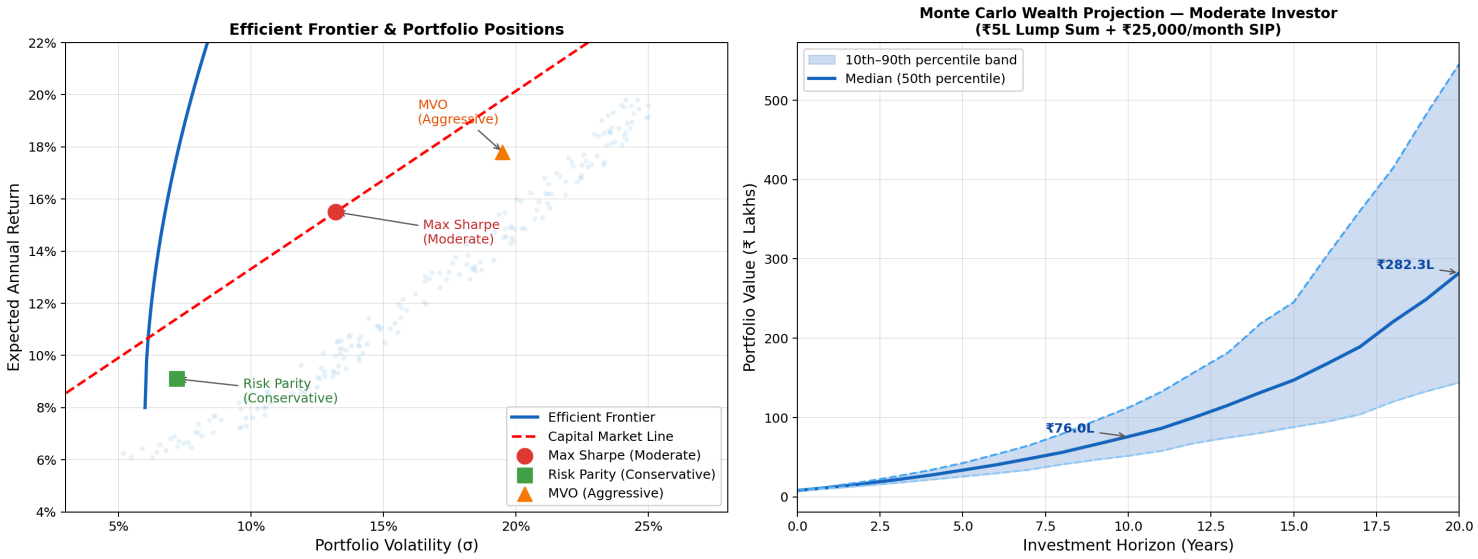


Figure 4

Left: Efficient frontier with portfolio positions for three investor profiles and Capital Market Line. Right: Monte Carlo wealth projection for Moderate investor (5,000 paths; ₹5L lump sum + ₹25,000/month SIP).

Fig. 5 Fund Scoring Framework (left) and Rebalancing Tax Impact Analysis (right)

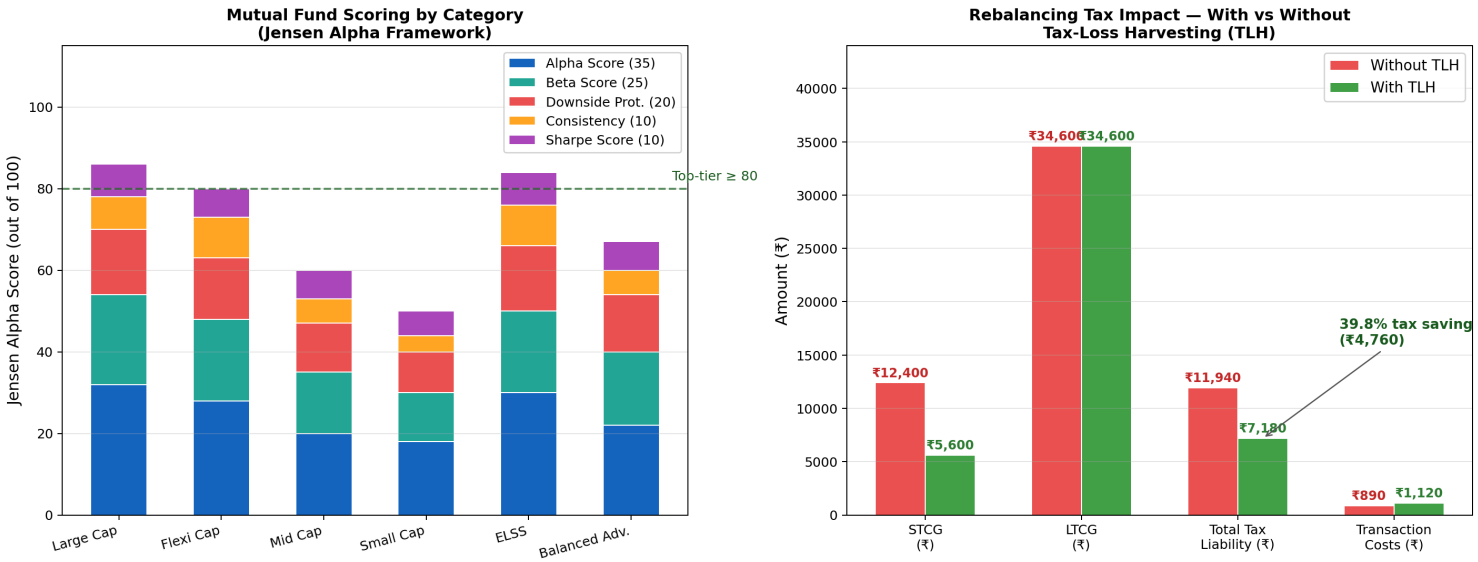


Figure 5

Left: Jensen Alpha-based mutual fund scoring framework across six fund categories. Right: Tax impact comparison of rebalancing with and without tax-loss harvesting (TLH).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- AppendixAThird.docx