International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 11 | Nov - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Application of Generative AI in Financial Planning and Wealth Management

Ganesh S Belkone, Dr. Deepak Sahoo

College of Management and Computer Application

MIT Art, Design and Technology University

Abstract

Generative AI tools are being built into advisory workflows faster than the evidence on their effects. This paper asks what a GenAI-assisted planning process would change for retail investors in India, and where the risks sit. Because client-level advisory data are private, the study uses a synthetic dataset of 800 clients (400 GenAI-assisted, 400 conventionally advised), a synthetic audit of 1,500 machine-generated advice outputs, and a Monte Carlo simulation of a fifteen-year savings goal. In the simulated data, GenAI assistance cuts planning time by about 40% and cost per plan by about 44%. It also raises the portfolio–risk alignment score by 7.4 points after controlling for age, assets, literacy and risk tolerance. The gain is larger for clients with low financial literacy, and the satisfaction advantage shrinks with age. Error rates in the audit fall from 28.2% for ungrounded prompting to 4.6% when retrieval grounding and advisor review are combined. The Monte Carlo results suggest that allocation tweaks move goal-success probability by only a few percentage points. All effects are products of the simulation's assumptions, which are reported in full, so the findings show how an evaluation could be structured and should not be read as evidence about real clients.

Keywords: generative AI, wealth management, financial planning, robo-advisory, synthetic data, Monte Carlo simulation, responsible AI

1. Introduction

Financial planning has always been a mix of arithmetic and conversation. The arithmetic (projecting a corpus, checking an allocation against a risk profile) is the part software has handled for years. The conversation, in which an advisor works out what a client actually wants and explains why a plan looks the way it does, has been much harder to automate. Large language models change this because they can produce fluent, personalised explanations on demand, and wealth managers have started to use them for drafting client reports, summarising research, answering routine queries and preparing first versions of financial plans.

The attraction is easy to see. A planning conversation that takes an advisor several hours of preparation could be drafted in minutes, and clients who were never profitable enough to serve individually might get some form of personalised guidance. The worry is just as easy to see. A language model can write a confident paragraph containing a wrong tax rule, a stale fund figure or a product that does not suit the client, and a client with little financial knowledge is the person least able to notice. Questions of this kind sit close to long-standing concerns about regulation and accountability in finance (Gupta & Mishra, 2014; Sahoo, 2018).

Empirical work on these questions is hard to do properly. Real advisory records contain income, holdings and family circumstances, so firms rarely share them, and the technology is changing faster than a conventional study can be designed and published. This paper takes a different route. It builds a synthetic population of retail clients, assigns some of them to GenAI-assisted planning, and examines what a careful evaluation would measure: efficiency, advice quality, client experience, error risk and goal attainment. The aim is to set out a transparent framework that can later be repeated on real data, and to show which results are sensitive to which assumptions.

The paper makes three contributions. It gives an explicit data-generating process for a GenAI advisory experiment, so a reader can see exactly what has been assumed. It analyses heterogeneity by financial literacy and age, since these

determine who is helped and who is exposed. And it treats advice errors as a measured outcome with its own audit design rather than a footnote.

2. Literature Review

2.1 Language models and sentiment in financial forecasting

Much of the Indian work on text-driven finance models predates the current wave of generative tools but points the same way. Dash and Mishra (2024a, 2024b) combine natural language understanding based sentiment analysis with deep learning to forecast stock prices and market trends, and Mishra and Dash (2025) extend this to an LSTM-based forecasting model. A survey of machine-learning sentiment methods for Indian stock trends (Mishra & A. D., 2023) places this work in a wider body of research. Their relevance here is that financial text can be turned into inputs for decision support. A generative model for planning sits one step further along: it reads, summarises and then writes in the client's own context.

2.2 Traditional analysis of funds, banks and risk

Before any AI layer is added, a planner still has to judge funds and issuers. Studies of mutual fund schemes compare risk and return across asset managers (Mishra & Peerapur, 2016), examine the risk profile of a single scheme (Mishra & Borole, 2017) and compare tax-saver funds (Pol et al., 2021), while fundamental analysis has been applied to selected banks for investment decisions (Mishra & Raje, 2019). Macro conditions also matter: Mishra (2015) examines how interest rates relate to stock market performance. These studies show the kind of evidence a human planner draws on, and they are the kind of material a retrieval-grounded GenAI system would have to reproduce correctly. They also set a benchmark for the audit in this paper, which counts as errors the places where generated advice departs from such evidence.

2.3 Digital assets, derivatives and newer products

Wealth management increasingly has to speak about products beyond equity and debt funds. Sahoo and Sahoo (2026a) analyse crypto derivatives markets in terms of price discovery, volatility and DeFi risk, and Sahoo and Sahoo (2025) review the prospects for blockchain in derivative markets. Mishra and Phansalkar (2025) survey blockchain security threats and vulnerability detection. For a GenAI planner this widens the surface for mistakes. A model that sounds fluent about a derivative or a token can still misstate its risk, and the suitability checks that apply to a plain index fund are less settled for such instruments.

2.4 Behaviour, sustainability and ethics

Personal finance is also about behaviour. Sahoo and Sahoo (2026b, 2026c) treat saving and wealth-building for younger investors as questions of habit and psychology as well as product choice, which suggests that explanation and tone could matter as much as the allocation itself. On sustainability, Sahoo and Deepak (2024a, 2024b) discuss ethical frameworks and ESG practice in sustainable finance, and Sugiawan and Managi (2026) show, in a cross-country setting, that progress on development goals does not guarantee that natural capital is preserved. The implication for planners is that a client's values and a plan's wider footprint are legitimate inputs that a generic model may overlook. Mishra et al. (2024) raise ethical dimensions of affective computing, including its societal consequences, a reminder that systems built to respond to human emotion carry their own accountability questions. Applied AI is also reaching transactional decisions elsewhere; Chaurasiya et al. (2025) describe an AI-supported bidding system for e-commerce resale. Taken together, this literature gives reasons to expect gains in process and personalisation, and reasons for care around accuracy, suitability and ethics.

2.5 Research gap and hypotheses

The works above cover forecasting, fund analysis, new asset classes and ethics, but none of them evaluates a GenAI-assisted planning workflow as a whole. This study fills that gap with a simulation and tests five hypotheses.

H1: GenAI assistance reduces planning time and cost per plan.

H2: GenAI assistance improves the alignment between a client's portfolio and stated risk profile.

H3: The alignment gain is larger for clients with lower financial literacy.

H4: The satisfaction advantage of GenAI assistance declines with client age.

H5: Retrieval grounding and advisor review reduce the share of generated outputs that contain errors.

3. Methodology

3.1 Why synthetic data

Synthetic data were used for three reasons. Client-level advisory records are confidential, GenAI planning has little track record to observe, and a simulation lets every assumption be stated and changed. The cost is that the results are only as informative as the assumptions. Where an effect appears in the results below, it appears because it was built in or because it follows from other built-in features; the analysis shows whether the proposed measures and models can recover those features from noisy data, and how large the effects look once non-random assignment is handled. It does not establish that they exist in practice. The dataset and code were generated in Python with a fixed random seed (2026) so the exercise can be reproduced.

3.2 Simulated client population

The population has 800 retail clients. Age was drawn from a normal distribution (mean 41, SD 11, truncated at 24 and 70). Investable assets were log-normal with a median of about ₹18 lakh. Financial literacy and risk tolerance were each scored from 1 to 5. Assignment to GenAI-assisted planning was not random: the probability of assignment fell with age and rose with literacy, which mimics early adopters being younger and more comfortable with finance. This produces a realistic confounding problem, and it is why the analysis relies on adjusted estimates instead of raw group comparisons. In the generated sample 400 clients ended up in the GenAI group and 400 in the conventional group.

Table 1. Assumed data-generating process for the client experiment

Outcome	Conventional	GenAI-assisted	Other terms
Planning time (hours)	N(6.5, 1.7)	N(3.9, 1.3)	Slight rise with assets
Cost per plan (₹)	N(9,500, 2,100)	N(5,200, 1,400)	Slight rise with assets
Alignment score (0–100)	68 + 3.0 × (literacy − 3.1)	+ 7.5 − 2.6 × (literacy − 3.1)	Noise SD 8.5
Risk-adjusted return	0.58 + 0.03 × (risk − 3)	+ 0.05	Noise SD 0.24
Satisfaction (1–5)	3.55 + 0.08 × (literacy − 3.1)	+ 0.38 − 0.012 × (age − 41)	Noise SD 0.62

Note. Values are the assumptions fed to the simulator. Literacy is centred at 3.1 and age at 41. Alignment and satisfaction were truncated to their valid ranges. Risk-adjusted return is a Sharpe-type ratio.

3.3 Outcome measures

Five client-level outcomes were simulated over a notional twelve-month planning cycle: planning time in hours, cost per plan in rupees, an alignment score from 0 to 100 that rates how closely the recommended portfolio matches the client's risk profile, a Sharpe-type risk-adjusted return, and a satisfaction rating on a five-point scale.

3.4 Advice-quality audit

A second synthetic exercise audits 1,500 generated advice outputs, 500 in each of three configurations: a baseline prompt with no external information, a retrieval-grounded configuration that draws on a curated knowledge base, and the grounded configuration followed by review from a human advisor. Each output was flagged independently for four error types: factual or numerical error, suitability mismatch, regulatory omission and outdated information. The per-type error probabilities were set at 14.2%, 7.1%, 5.8% and 4.9% for the baseline; 5.2%, 4.8%, 3.1% and 2.2% for the grounded version; and 1.4%, 1.6%, 0.8% and 1.0% with advisor review. An output counted as erroneous if it carried at least one flag.

3.5 Goal-attainment simulation

To see how much allocation choices matter, a representative client was simulated over fifteen years with an initial corpus of ₹5 lakh, a monthly contribution of ₹25,000 and a goal of ₹1.1 crore. Annual returns were drawn from normal distributions with 12% mean and 18% standard deviation for equity and 7% mean and 4% standard deviation for debt, with no correlation between them. Three approaches were compared across 10,000 paths each: a static 60/40 equity–debt split, a 75/25 split matched to a higher risk profile, and a glide path moving from 75% to 45% equity. These parameters are illustrative.

3.6 Statistical analysis

Group differences were tested with Welch's t-test and summarised with Cohen's d. The main adjusted estimates come from ordinary least squares with heteroskedasticity-robust (HC3) standard errors, controlling for age, log assets, literacy and risk tolerance. Heterogeneity was examined through interaction terms between GenAI assistance and literacy (for alignment) and between GenAI assistance and age (for satisfaction). The audit was tested with a chi-square test of independence on the share of erroneous outputs.

4. Results

4.1 Sample description

Table 2. Client characteristics by planning approach

Characteristic	GenAI-assisted	Conventional	Full sample
Age (years)	38.9	44.5	41.7 (10.6)
Investable assets (₹ lakh)	24.2	25.2	24.7 (27.4)
Financial literacy (1–5)	3.22	2.88	3.05 (0.88)
Risk tolerance (1–5)	3.03	2.87	2.95 (0.99)

Note. n = 400 GenAI-assisted and 400 conventional. Full-sample cells show mean (SD).

The GenAI group is younger by about 5.6 years and has higher average financial literacy (3.22 against 2.88), as the assignment mechanism intended. Raw comparisons of outcomes will therefore mix the effect of the approach with the effect of who chose it, which is why the adjusted models matter.

4.2 Efficiency, alignment and client experience

Table 3. Outcome comparison by planning approach

Outcome	GenAI M (SD)	Conv. M (SD)	t	p	d
Planning time (hrs)	3.93 (1.33)	6.56 (1.73)	-24.1	< .001	-1.70
Cost per plan (₹)	5,295 (1,487)	9,401 (2,120)	-31.7	< .001	-2.24
Alignment score	75.4 (8.9)	67.2 (9.0)	13.0	< .001	0.92
Risk-adjusted return	0.632 (0.236)	0.588 (0.239)	2.6	.010	0.18
Satisfaction (1–5)	3.92 (0.63)	3.50 (0.59)	9.9	< .001	0.70

Note. Welch's t-test; d is Cohen's d (GenAI minus conventional). Negative values mean lower in the GenAI group.

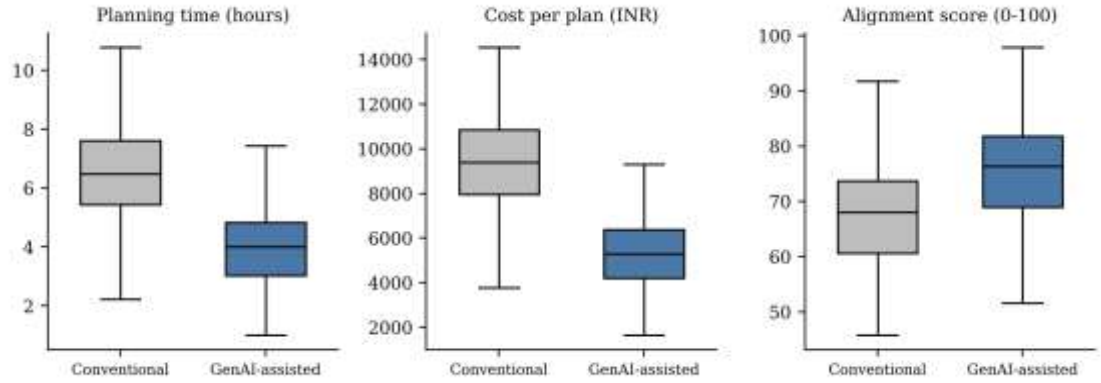


Figure 1. Planning time, cost per plan and alignment score by planning approach (outliers not shown)

Planning time fell from 6.6 to 3.9 hours and cost per plan from ₹9,401 to ₹5,295, reductions of roughly 40% and 44%. H1 is therefore supported, though it is worth remembering that the reduction was written into the assumptions. The alignment score was 75.4 in the GenAI group against 67.2 conventionally (d = 0.92). The satisfaction gap was 0.43 points (d = 0.70). The difference in risk-adjusted return was small (0.632 against 0.588, d = 0.18), which fits what one would expect: better alignment with a client's risk profile is not the same as higher returns.

4.3 Adjusted estimates

Table 4. OLS estimates of the effect of GenAI assistance (robust standard errors)

Term	Alignment (1)	Alignment (2)	Satisfaction (3)	Risk-adj. return (4)
GenAI-assisted	7.36 (0.65)***	18.65 (2.23)***	0.394 (0.046)***	0.039 (0.018)*
Financial literacy	2.24 (0.37)***	4.08 (0.46)***	0.047 (0.025)	–
GenAI × literacy	–	-3.69 (0.70)***	–	–
Risk tolerance	–	–	–	0.026 (0.009)**
R^2	0.216	0.242	0.115	0.026

Note. All models also control for age, log assets and risk tolerance (except where risk tolerance is shown). Entries are coefficients with robust standard errors in parentheses. Significance: one asterisk $p < .05$, two asterisks $p < .01$, three asterisks $p < .001$.

After adjustment the alignment effect is 7.4 points (SE 0.65), close to the raw difference of 8.2. In this simulation the confounding from age and literacy is therefore small for alignment. The adjusted satisfaction effect is 0.39 points and the adjusted risk-adjusted-return effect is 0.039 ($p = .030$), modest and explaining very little of the variance ($R^2 = 0.026$).

4.4 Who benefits? Literacy and age

The interaction of GenAI assistance with literacy is -3.69 (SE 0.70, $p < .001$), so each extra point of literacy cuts the GenAI alignment advantage by about 3.7 points. Splitting the sample into literacy terciles shows the same pattern (Figure 2). The advantage is 11.3 points in the lowest tercile, 7.0 in the middle and 4.5 in the highest. H3 is supported.

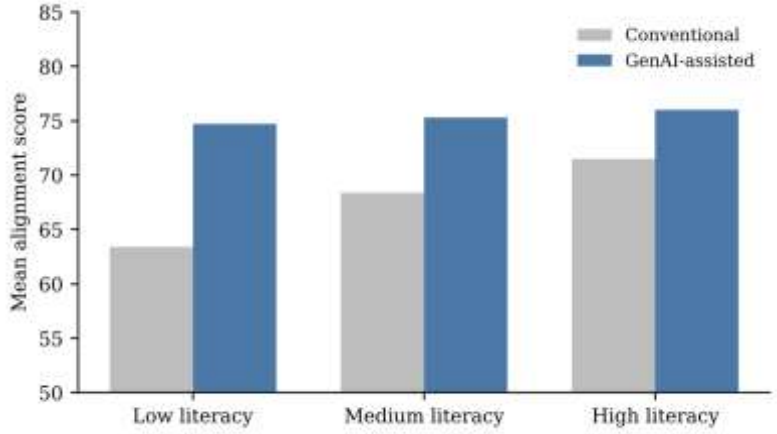


Figure 2. Mean alignment score by financial literacy tercile and planning approach

For satisfaction, the GenAI × age interaction is -0.0088 per year (SE 0.0042, p = .033). Over a twenty-year age gap, the satisfaction advantage shrinks by roughly 0.18 points. H4 is supported, though weakly compared with the other effects.

There is a tension here that the numbers make visible. Low-literacy clients stand to gain the most alignment from GenAI-assisted planning, and in this sample they were also the group least likely to be assigned to it (110 of 267 against 162 of 267 in the highest tercile). Who gets access matters as much as how well the tool works.

4.5 Error risk

Table 5. Audit of 1,500 generated advice outputs (500 per configuration)

Configuration	Factual	Suitability	Regulatory	Outdated	Any error
Baseline prompt only	61 (12.2%)	40 (8.0%)	31 (6.2%)	20 (4.0%)	141 (28.2%)
Retrieval-grounded	28 (5.6%)	31 (6.2%)	16 (3.2%)	9 (1.8%)	81 (16.2%)
Grounded + advisor review	8 (1.6%)	7 (1.4%)	6 (1.2%)	3 (0.6%)	23 (4.6%)

Note. Counts of flagged outputs with percentages in parentheses. An output can carry more than one flag, so type counts do not sum to the "any error" column. Chi-square test on any error: $\chi^2(2) = 101.9$, $p < .001$.

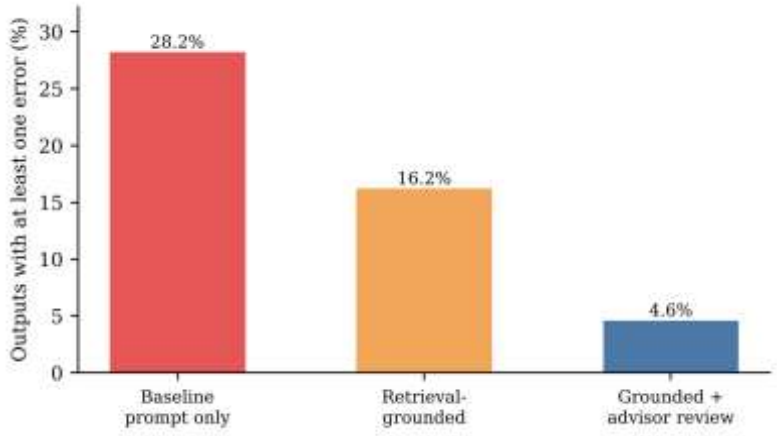


Figure 3. Share of generated outputs containing at least one error, by configuration

More than one in four ungrounded outputs (28.2%) contained at least one error. Retrieval grounding brought this to 16.2%, and advisor review to 4.6%. The difference across configurations is large ($\chi^2(2) = 101.9$, $p < .001$), which supports H5. Factual errors fell most steeply with grounding, while suitability mismatches fell more slowly, since a model that has the right facts can still recommend the wrong product for the person. Even in the best configuration about one output in twenty needed correction.

4.6 Does the allocation matter? Goal-attainment simulation

Table 6. Simulated outcomes for a fifteen-year ₹1.1 crore goal (10,000 paths)

Approach	Goal success (%)	Median corpus (₹ lakh)	10th percentile (₹ lakh)
Static 60/40 rule	63.5	121.1	84.1
Risk-profiled 75/25	65.7	127.1	80.7
Glide-path 75 to 45	61.6	118.8	84.8

Note. Equity returns N(12%, 18%) and debt returns N(7%, 4%), uncorrelated; ₹25,000 invested monthly from a ₹5 lakh start.

Success probabilities sit within about four percentage points of each other: 63.5% for the static split, 65.7% for the 75/25 split and 61.6% for the glide path. The higher equity share lifts the median outcome and also lowers the 10th percentile (₹80.7 lakh against ₹84.1 lakh), so the better odds come with a worse bad-case outcome. In this setup the contribution amount and the time horizon matter much more than the exact mix, which suggests that a planner's best use of a GenAI tool may be explanation, goal-setting and keeping the client contributing, rather than hunting for a cleverer allocation.

5. Discussion

Read as a demonstration, the study shows what a GenAI advisory evaluation can and cannot tell us. The efficiency result is the least surprising and the most dependent on assumptions. If a draft plan really does take four hours instead of six and a half, the economics of serving smaller clients change, but that depends on advisors actually reviewing outputs, and the review step is exactly what the audit suggests cannot be skipped.

The distributional result deserves more attention. Gains in alignment were largest for low-literacy clients, as H3 predicted, yet the same clients were less likely to receive GenAI-assisted advice, and they are the ones least equipped to spot a wrong answer. A firm that deploys these tools with no review for retail clients might therefore be putting the heaviest error exposure on those who benefit most. This echoes the behavioural emphasis in the personal finance literature (Sahoo & Sahoo, 2026b, 2026c): for many clients the obstacle is understanding and habit, and a tool that explains well could address it, provided the explanation is right.

The audit results make a practical point about architecture. Moving from a bare prompt to a grounded system removed most factual and outdated-information errors, which fits the way earlier forecasting work depends on well-prepared financial text and data (Dash & Mishra, 2024a, 2024b; Mishra & Dash, 2025). Suitability errors were more stubborn. Deciding whether a product fits a person is a judgement about that person's circumstances, which is why advisor review kept its value in the last configuration. This matters even more for newer instruments such as crypto derivatives, where risks are harder to characterise (Sahoo & Sahoo, 2026a).

Finally, the Monte Carlo exercise is a reminder to keep the claims modest. A GenAI tool does not produce better markets. Its plausible contribution is a better process: faster drafts, clearer explanations, closer matching of risk profile and portfolio. The small spread in goal-success probability, and the almost negligible difference in risk-adjusted return in the client experiment, point the same way. Sustainability and ethical preferences are another area where a generic tool could fall short (Sahoo & Deepak, 2024a, 2024b; Sugiawan & Managi, 2026), and a planning assistant would need to ask about them explicitly.

6. Implications

For wealth managers, the results argue for deploying GenAI as an assistant inside a controlled workflow rather than as an autonomous advisor. That means grounding outputs in a maintained knowledge base, keeping a named human accountable for the final advice, and logging outputs so they can be audited. For regulators, the audit design here could be adapted into a routine test of a firm's AI-generated advice, with error categories that map onto existing suitability and disclosure duties (Gupta & Mishra, 2014; Sahoo, 2018). For clients, the main message is that generated advice is a draft

of someone's opinion and should be checked, especially when it concerns complex products. For researchers, the framework can be repeated with real or partner-supplied data, with the synthetic parameters replaced by estimates.

7. Limitations and Future Research

The most important limitation is that every finding is conditional on the assumptions in Table 1 and Section 3.4. Effects were built into the generator, so statistical significance here says nothing about whether the real effects exist or how large they are. With 800 simulated clients and strong assumed effects, tiny p-values are expected. The size of the effects should not be treated as a forecast.

Other limits follow. Assignment to GenAI assistance depended on only two observed variables, whereas real adoption is shaped by unobserved factors such as trust, advisor preference and income stability, so real confounding would be harder to remove. The twelve-month window ignores market shocks, behavioural responses to losses and regulatory change. The Monte Carlo model assumes normally distributed, uncorrelated returns, which understates tail risk. The audit probabilities describe an imagined system and not any particular product, and the error taxonomy would need validation by practising advisors.

Future work could proceed in three directions: a randomised or quasi-experimental pilot with a partner advisory firm, an audit using real model outputs scored by licensed planners, and a study of how clients with different levels of literacy interpret generated advice. The simulation code can be reused to run power calculations before such a pilot.

8. Conclusion

This paper set out a simulation framework for evaluating GenAI in financial planning and wealth management and applied it to synthetic data. Under the stated assumptions, GenAI-assisted planning is faster and cheaper, aligns portfolios with risk profiles better, and is rated more highly by clients, with the largest alignment gains among less financially literate clients and a satisfaction advantage that fades with age. Unchecked generated advice carried at least one error in more than one output in four, and grounding plus advisor review brought that to under five in a hundred. The goal-attainment simulation shows that allocation choices move the odds of reaching a goal far less than saving behaviour does. The practical lesson is a conditional one: the technology can widen access to personalised planning if it is grounded, audited and supervised by accountable humans, and it can do harm to the clients least able to detect errors if it is not. These claims should be tested on real data before anyone relies on them.

References

Chaurasiya, S., Mishra, D. U., & Shekhar, S. (2025). Smart reselling: Leveraging MERN stack and AI for intelligent bidding in e-commerce. *International Journal of Advanced Research in Science, Communication and Technology, 5*, 402–410.

Dash, A. S., & Mishra, U. (2024a). Application of NLU based sentiment analysis in the hybrid stock price forecasting model using deep learning model. *2024 IEEE Pune Section International Conference (PuneCon)*, 1–6.

Dash, A. S., & Mishra, U. (2024b). Stock market trend prediction model using deep learning based sentiment analysis of financial data. *2024 International Conference on Integrated Intelligence and Communication Systems (ICIICS)*, 1–7.

Gupta, M., & Mishra, U. (2014). *Financial regulatory framework*.

Mishra, D., Deshpande, S., Anna, M. G., & Tiwari, A. (2024). Exploring the ethical dimensions and societal consequences of affective computing. In *Affective computing for social good: Enhancing well-being, empathy, and equity* (pp. 91–105).

Mishra, D., & Phansalkar, S. (2025). Blockchain security in focus: A comprehensive investigation into threats, smart contract security, cross-chain bridges, vulnerabilities detection tools & techniques. *IEEE Access*.

Mishra, U. (2015). A study on impact of interest rates and the performance of stock market. *International Journal of Advances in Management, Technology and Engineering Sciences, 4*(12[III]), 28–33.

Mishra, U., & A. D. (2023). Sentiment analysis using machine learning for forecasting Indian stock trend: A brief survey. *Finance: Theory and Practice, 27*(6), 136–147.

Mishra, U., & Dash, A. S. (2025). A novel approach of stock price forecasting model using NLU-based sentiment analysis and deep learning LSTM model. *Neural Computing and Applications*, 1–30.

Mishra, U., & Raje, H. (2019). Fundamental analysis of selected banks for investment decisions. *International Journal of Research in Social Sciences, 9*(2).

Mishra, U. M., & Borole, M. C. (2017). A comparative risk analysis of Kotak Select Focus Fund (G) mutual fund scheme. *Imperial Journal of Interdisciplinary Research, 3*(3), 4.

Mishra, U. M., & Peerapur, V. (2016). Comparative analysis of Birla Sun Life mutual fund schemes with other asset management company's schemes. *International Journal of Engineering and Management Research, 6*(1), 241–244.

Pol, A., Raje, H., & Mishra, U. (2021). Comparative performance analysis of selected mutual fund schemes in tax saver category. *PalArch's Journal of Archaeology of Egypt/Egyptology, 18*(7), 2325–2337.

Sahoo, D. R. (2018). Corporate governance: Need of the hour (Companies Act 2013). *International Journal of Management Studies, 4*(ii).

Sahoo, D. R., & Deepak, S. V. (2024a). Exploring the ethical perspectives of sustainable finance: A research study. In *Advances in Business Strategy and Competitive Advantage* (pp. 119–142). IGI Global.

Sahoo, D. R., & Deepak, S. V. (2024b). Exploring the nexus of sustainable finance: ESG ethical framework practices. In *Ethical quandaries in business practices: Exploring morality and social responsibility* (pp. 253–280). IGI Global.

Sahoo, D. R., & Sahoo, V. D. (2026a). Crypto derivatives markets: Price discovery, volatility dynamics, market efficiency, and DeFi risk: A structural analysis. *Journal of Engineering and Technology Management, 80*(April–June), 774–799.

Sahoo, D. R., & Sahoo, V. D. (2026b). *Personal finance for millennials & Gen Z: From broke to building wealth*.

Sahoo, D. R., & Sahoo, V. D. (2026c). *Wealth begins within: The psychology of money, meaning, and the mind*.

Sahoo, V. D., & Sahoo, D. R. (2025). Prospective of blockchain in derivative markets: An empirical review. In *Innovations in blockchain-powered intelligence and cognitive Internet of Things (CIoT)* (pp. 229–252).

Sugiawan, Y., & Managi, S. (2026). SDG progress and the challenge of inclusive wealth sustainability. *iScience, 29*(6), 116058. https://doi.org/10.1016/j.isci.2026.116058