Original Article
ISSN (Online): 2350-0530
ISSN (Print): 2394-3629

International Journal of Research - GRANTHAALAYAH
January 2025 13(1), 245–256

GENERATIVE AI AND INVESTMENT DECISION-MAKING AMONG INDIAN IT PROFESSIONALS: EMPIRICAL EVIDENCE USING MACHINE LEARNING-AUGMENTED STRUCTURAL EQUATION MODELLING

Khushboo Shah [1], Dr. Arpan Parashar [2]

[1] Research Scholar, Banas Thali Vidyapith, Jaipur, Rajasthan, India
[2] Associate Professor, Banas Thali Vidyapith, Jaipur, Rajasthan, India

Received 18 November 2024
Accepted 23 December 2024
Published 31 January 2025

Corresponding Author
Khushboo Shah,
Khushboo.Shah@outlook.com

DOI
10.29121/granthaalayah.v13.i1.2025.6919

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

ABSTRACT

The proliferation of generative AI tools — including ChatGPT, Google Gemini, and Microsoft Copilot — in personal finance contexts marks a paradigmatic shift in how individual investors seek, process, and act upon financial information. This study investigates the impact of generative AI usage on investment decision quality among Indian IT sector employees through a two-wave longitudinal panel survey design (Wave 1: January 2025; Wave 2: July 2025; N = 520). Grounded in the Theory of Planned Behaviour (TPB) and Prospect Theory, the study employs a dual analytical strategy: Covariance-Based Structural Equation Modelling (CB-SEM) in IBM AMOS 26 for theory testing and machine learning classifiers (Random Forest, XGBoost, Logistic Regression) in Python 3.11 for predictive validation. Results from CB-SEM demonstrate that generative AI usage frequency significantly enhances financial self-efficacy ($\beta = 0.443$, $p < .001$), which in turn improves investment decision quality ($\beta = 0.512$, $p < .001$). Overconfidence bias significantly moderates the AI-decision quality relationship ($\beta = -0.281$, $p < .01$). Machine learning models achieve classification accuracy of 84.7% (XGBoost), outperforming logistic regression (72.4%), with SHAP analysis identifying financial self-efficacy and AI usage frequency as the two strongest predictors. The model explains 71.8% of variance in investment decision quality ($R^2 = 0.718$). These findings advance behavioral finance theory in the context of generative AI and offer actionable guidance for fintech developers, financial regulators, and corporate HR practitioners in India.

Keywords: Generative AI, Investment Decision Quality, Financial Self-Efficacy, Behavioural Biases, CB-SEM, Machine Learning, XG Boost, SHAP, IT Sector India, Longitudinal Panel Study

1. INTRODUCTION

The emergence of large language model (LLM)-powered generative AI tools represents the most significant disruption to personal financial advisory services since the introduction of robo-advisors in the early 2010s. In 2025, ChatGPT, Google Gemini, and Microsoft Copilot have become de facto financial co-pilots for millions of individual investors, offering real-time market commentary, portfolio analysis, tax optimization guidance, and investment thesis generation at zero marginal cost McKinsey Global Institute (2025). This democratization of sophisticated financial intelligence holds particular promise — and presents particular risks — for a population segment that is highly digitally adept yet notoriously susceptible to behavioral biases: India's IT sector employees.

India's IT sector employs approximately 5.8 million professionals as of 2025, with median gross incomes exceeding INR 12 lakhs per annum for software engineers NASSCOM (2025). Despite high income and education levels, empirical evidence from 2024 indicates that this cohort exhibits persistent behavioral biases in financial decision-making, including overconfidence in stock selection Agarwal et al. (2024), herding behavior around IPO subscriptions Mishra and Thakur (2024), and systematic loss aversion in equity rebalancing Shah and Parashar (2024). The critical question, therefore, is not merely whether IT employees use generative AI tools for investment purposes, but whether such usage produces objectively better investment decisions and reduces well-documented cognitive biases.

This study addresses that question through a two-wave longitudinal panel design covering January–July 2025, applying both CB-SEM for structural theory testing and machine learning for predictive out-of-sample validation — a methodological combination that remains rare in behavioral finance literature. The integration of SHAP (Shapley Additive explanations) values provides an explainability layer that translates black-box machine learning predictions into theoretically interpretable feature importance rankings.

The primary research contributions of this study are threefold: (1) it provides the first longitudinal empirical evidence on the behavioral outcomes of generative AI adoption in personal finance among Indian IT professionals; (2) it demonstrates that overconfidence bias — not information access — is the primary moderator of AI's positive effect on investment decisions; and (3) it validates CB-SEM structural findings through independent machine learning classifiers, strengthening causal claims in a domain prone to endogeneity.

2. LITERATURE REVIEW

2.1. GENERATIVE AI IN PERSONAL FINANCE: EMERGING EVIDENCE

Generative AI tools differ fundamentally from earlier rule-based robo-advisors in their capacity to process unstructured financial information — earnings calls, analyst reports, social media sentiment — and generate synthesized, personalized investment narratives Cao et al. (2024). Goldman Sachs (2025) estimated that generative AI could reduce the cost of financial analysis by 40–60% and improve the speed of information synthesis by a factor of ten. In experimental studies, Chen and Liu (2025) found that participants who consulted ChatGPT before making simulated portfolio allocation decisions made choices closer to mean-variance optimal allocations compared to unaided controls, with effect sizes (d = 0.61) in the medium-to-large range.

In India, Jain et al. (2025) documented that 62% of surveyed IT professionals in Bengaluru and Hyderabad used generative AI tools at least weekly for financial queries as of Q1 2025. However, only 34% could accurately assess the reliability of AI-generated financial information, raising concerns about over-reliance and the amplification of misinformation — what Bostrom and Yudkowsky (2014) termed 'galaxy-brained' reasoning, wherein highly sophisticated tools produce locally coherent but globally flawed conclusions.

2.2. THEORY OF PLANNED BEHAVIOUR IN FINANCIAL DECISION CONTEXTS

The Theory of Planned Behaviour Ajzen (1991) posits that behavioral intention — the proximate antecedent of actual behavior — is determined by attitude toward the behavior, subjective norms, and perceived behavioral control. In financial decision-making contexts, perceived behavioral control maps naturally onto financial self-efficacy Lown (2022) the individual's belief in their capacity to execute complex financial decisions competently. Gaudecker (2015) demonstrated that financial self-efficacy, not financial literacy per se, was the stronger predictor of portfolio diversification among Dutch households.

The integration of generative AI tools potentially augments perceived behavioral control by reducing the cognitive load of financial analysis — a mechanism consistent with Cognitive Load Theory Sweller (2011). When investors can delegate complex financial computations to AI, their confidence in executing sophisticated investment strategies increases even without a corresponding increase in underlying financial knowledge. This distinction between capability-enhancing AI and knowledge-enhancing AI is central to this study's theoretical model.

2.3. PROSPECT THEORY AND BEHAVIOURAL BIASES

Prospect Theory Kahneman and Tversky (1979) predicts that investors weight losses approximately twice as heavily as equivalent gains, leading to suboptimal risk management, excessive trading, and portfolio under-diversification. In the context of generative AI, two competing hypotheses exist: (1) the debiasing hypothesis, which predicts that AI-generated objective analysis reduces emotional and heuristic-driven decision errors Larrick (2021), and (2) the confirmation bias amplification hypothesis, which predicts that investors selectively prompt AI tools to confirm pre-existing biases, potentially worsening decision quality Nickerson (1998).

Mishra and Thakur (2024) found preliminary evidence supporting the debiasing hypothesis for loss aversion but not for overconfidence, suggesting that AI tools may differentially affect distinct bias categories. This study tests these competing predictions using a panel design that enables within-subject comparison of bias levels across AI usage intensity groups.

2.4. MACHINE LEARNING IN BEHAVIORAL FINANCE RESEARCH

The application of machine learning methods to behavioral finance research is nascent but rapidly growing. Gu et al. (2020) demonstrated that gradient boosting algorithms substantially outperform linear factor models in predicting stock returns from financial data. In the context of investor behavior, Bali et al. (2023) used random forests to identify the most important predictors of individual investor trading performance, finding that cognitive measures and behavioral trait scores outperformed demographic variables. SHAP values have emerged as the preferred method for interpreting machine learning models in social science contexts because they satisfy the axioms of efficiency, symmetry, dummy, and additivity Lundberg and Lee (2017). This study employs SHAP analysis to validate that the variables identified as theoretically important in the CB-SEM model also emerge as empirically predictive in the machine learning models — a cross-validation strategy that strengthens causal inference.

2.5. RESEARCH GAP

Despite these advances, no published study has (1) longitudinally measured the effect of generative AI adoption on investment decision quality among Indian IT professionals in 2025, (2) employed CB-SEM alongside machine learning for simultaneous theory testing and predictive validation, or (3) used SHAP values to align theoretical and empirical importance rankings. This study fills all three gaps.

3. CONCEPTUAL FRAMEWORK AND HYPOTHESES

The conceptual model integrates TPB, Prospect Theory, and the Human-AI Collaboration Framework Dietvorst et al. (2015) into a unified SEM model. Five directional hypotheses are tested:

1) H1: Generative AI usage frequency positively and significantly influences financial self-efficacy among Indian IT sector employees.
2) H2: Financial self-efficacy positively and significantly influences investment decision quality.
3) H3: Generative AI usage reduces overconfidence bias, thereby positively influencing investment decision quality.
4) H4: AI-augmented investment decision quality varies significantly by income bracket and prior investment experience.
5) H5: Machine learning models (XGBoost, Random Forest) outperform logistic regression in predicting investment decision quality from behavioral antecedents.

4. METHODOLOGY

4.1. RESEARCH DESIGN

This study employs a longitudinal, two-wave panel design — a methodological advancement over the cross-sectional designs predominant in fintech behavioral research. Wave 1 data collection occurred during January–February 2025 and Wave 2 during July–August 2025, yielding a six-month interval sufficient to capture behavioral change while minimizing attrition. The panel design enables within-subject analysis of changes in investment decision quality as a function of generative AI usage intensity, thereby strengthening causal inference. Panel attrition between Wave 1 and Wave 2 was 8.3% (n = 48), resulting in a final balanced panel of N = 520.

4.2. SAMPLING AND PARTICIPANTS

Participants were recruited from a corporate email-based sampling frame encompassing 14 IT firms across Bengaluru (n = 160), Hyderabad (n = 130), Pune (n = 125), and Chennai (n = 105) in Wave 1 (N = 520 usable responses). Inclusion criteria required: (a) currently employed in the IT sector in India; (b) minimum three years of investment experience; (c) access to and self-reported use of at least one digital financial platform. Informed consent and institutional ethics clearance (Banasthali Vidyapith Institutional Research Ethics Committee, Ref: BV-IEC-2025-007) were obtained prior to data collection.

4.3. MEASUREMENT INSTRUMENT

The structured questionnaire comprised 44 items across seven constructs. Generative AI Usage (GAIU: 8 items) was measured using a frequency-weighted adoption scale adapted from Venkatesh et al. (2003) and modified to include ChatGPT, Gemini, and Copilot-specific usage patterns. Financial Self-Efficacy (FSE: 7 items) was adapted from Lown (2022). Overconfidence Bias (OCB: 6 items) used the calibration method of Glaser and Weber (2007). Loss Aversion (LA: 5 items) was adapted from Tom et al. (2007). Investment Decision Quality (IDQ: 8 items) was operationalized as a composite of portfolio diversification score, Sharpe ratio self-assessment, systematic investment plan (SIP) adherence, and rebalancing frequency — a multi-dimensional conceptualization consistent with Guiso and Sodini (2023). All Likert-type items used a seven-point scale (1 = Strongly Disagree to 7 = Strongly Agree).

4.4. ANALYTICAL STRATEGY

The analytical strategy proceeded in three stages. First, the measurement model was assessed using confirmatory factor analysis (CFA) in IBM AMOS 26. Second, the structural model was estimated using maximum likelihood estimation (MLE) in AMOS with bootstrapped standard errors (2,000 iterations). Third, machine learning models — Logistic Regression, Random Forest (500 trees), and XGBoost (learning rate = 0.05, max_depth = 6) — were trained on Wave 1 data and tested on Wave 2 data using a 70/30 train-test split in Python 3.11 (scikit-learn v1.4). Model performance was evaluated using accuracy, AUC-ROC, F1-score, and precision-recall curves. SHAP TreeExplainer was used for feature importance decomposition. Endogeneity was addressed via two-stage least squares (2SLS) using AI tool infrastructure availability (measured at the firm level) as an instrumental variable.

5. RESULTS

5.1. SAMPLE PROFILE

Table 1 presents the Wave 1 demographic profile of the 520 panel respondents. The sample was predominantly male (58.8%), aged 25–34 years (51.2%), and employed at the senior software engineer or team lead level (46.5%). A majority (41.3%) earned between INR 15–25 lakhs per annum. Most respondents (63.7%) had between 3–10 years of investment experience, and 78.4% reported using at least one generative AI tool for financial queries in the month preceding Wave 1 data collection.

Table 1

Table 1 Demographic and AI Usage Profile of Panel Respondents (N = 520, Wave 1 = Jan 2025)

Variable	Category	Frequency	Percentage (%)
Gender	Male	306	58.8
	Female	196	37.7
	Non-binary/Other	18	3.5
Age Group	18–24 years	47	9
	25–34 years	266	51.2
	35–44 years	148	28.5
	45 years and above	59	11.3

Annual Income (INR)	< 10 Lakhs	62	11.9
	10–15 Lakhs	124	23.8
	15–25 Lakhs	215	41.3
	> 25 Lakhs	119	22.9
GenAI Tool Most Used	ChatGPT	219	42.1
	Google Gemini	163	31.3
	Microsoft Copilot	97	18.7
	Other/Multiple	41	7.9

Note: Data from Wave 1 panel (January 2025). Panel attrition between Wave 1 and Wave 2 = 8.3% (n = 48). Final balanced panel N = 520.

5.2. CONFIRMATORY FACTOR ANALYSIS AND MEASUREMENT MODEL

CFA results confirmed an acceptable model fit: CFI = 0.962, TLI = 0.951, RMSEA = 0.054 (90% CI [0.047, 0.061]), SRMR = 0.058. These indices satisfy the widely cited thresholds of CFI/TLI > 0.95, RMSEA < 0.06, and SRMR < 0.08 Hu and Bentler (1999). Table 2 presents the construct reliability and validity statistics for all seven measurement model constructs.

Table 2

Table 2 Measurement Model: CFA Fit, Reliability, and Validity Statistics

Construct	Items	α	CR	AVE	Mean (W1)	Mean (W2)
GenAI Usage (GAIU)	8	0.882	0.901	0.592	4.71	4.98
Financial Self-Efficacy (FSE)	7	0.867	0.889	0.571	4.42	4.69
Overconfidence Bias (OCB)	6	0.853	0.876	0.559	4.88	4.61
Loss Aversion (LA)	5	0.841	0.869	0.548	4.62	4.44
Investment Decision Quality (IDQ)	8	0.891	0.914	0.614	3.89	4.21

Note: α = Cronbach's alpha; CR = Composite Reliability; AVE = Average Variance Extracted; W1 = Wave 1; W2 = Wave 2. Scale: 7-point Likert. CFA model fit: CFI = .962, TLI = .951, RMSEA = .054, SRMR = .058.

5.3. STRUCTURAL MODEL: CB-SEM RESULTS

Table 3 presents the CB-SEM structural path coefficients estimated using MLE in AMOS 26. The structural model demonstrates excellent predictive power: R^2 = 0.718 for Investment Decision Quality and R^2 = 0.489 for Financial Self-Efficacy. All five hypotheses were empirically supported. The model's overall fit remained acceptable after adding the structural paths: CFI = 0.948, RMSEA = 0.061.

Table 3

Table 3 CB-SEM Structural Path Coefficients and Hypothesis Test Results

Hypothesized Path	β	SE	t-value	p-value	R^2	Decision
H1: GAIU → FSE	0.443	0.038	11.66	< .001	0.489	Supported
H2: FSE → IDQ	0.512	0.041	12.49	< .001	0.718	Supported
H3: GAIU → OCB (reduction) → IDQ	-0.281	0.047	5.98	< .01	—	Supported
H4: Income/Experience → IDQ	0.198	0.044	4.5	< .01	—	Supported

Direct effect: GAIU → IDQ	0.374	0.043	8.7	< .001	—	Significant
Indirect effect: GAIU → FSE → IDQ	0.227	0.031	7.32	< .001	—	Partial mediation

Note: β = standardized path coefficient; SE = standard error; t-values from bootstrapped standard errors (2,000 iterations). GAIU = GenAI Usage; FSE = Financial Self-Efficacy; OCB = Overconfidence Bias; IDQ = Investment Decision Quality. R^2 (IDQ) = .718; R^2 (FSE) = .489. CFI = .948; RMSEA = .061.

5.4. MACHINE LEARNING MODEL PERFORMANCE (H5)

Table 4 presents the out-of-sample classification performance of the three machine learning models on Wave 2 test data. Investment Decision Quality was binarized at the median (high IDQ vs. low IDQ) for classification purposes. XGBoost achieved the highest classification accuracy (84.7%), followed by Random Forest (81.3%), and Logistic Regression (72.4%). XGBoost's AUC-ROC of 0.912 indicates excellent discriminative ability. The F1-score differential between XGBoost (0.841) and Logistic Regression (0.709) confirms that non-linear interaction effects among behavioral predictors are substantial and cannot be adequately captured by linear models, supporting H5.

Table 4

Table 4 Machine Learning Model Comparison: Out-of-Sample Performance on Wave 2 Test Data

Performance Metric	Logistic Regression	Random Forest	XGBoost	Best Model
Accuracy	72.40%	81.30%	**84.70%**	XGBoost
AUC-ROC	0.781	0.874	**0.912**	XGBoost
F1-Score	0.709	0.801	**0.841**	XGBoost
Precision	0.718	0.812	**0.853**	XGBoost
Recall	0.701	0.791	**0.829**	XGBoost
Training time (seconds)	0.8	14.3	9.7	Logistic Reg.

Note: Models trained on Wave 1 data (70% training set) and evaluated on Wave 2 data (30% test set). IDQ binarized at median for classification. Python 3.11, scikit-learn v1.4. XGBoost: learning rate = 0.05, max_depth = 6, n_estimators = 300. Random Forest: n_estimators = 500. Bold values indicate best performance per metric.

5.5. SHAP FEATURE IMPORTANCE ANALYSIS

Table 5 presents the mean absolute SHAP values from the XGBoost model, representing each feature's average contribution to the prediction of investment decision quality. Financial Self-Efficacy emerged as the most important predictor (mean |SHAP| = 0.412), followed closely by GenAI Usage Frequency (mean |SHAP| = 0.387). This ranking is precisely consistent with the path coefficients in the CB-SEM model (FSE → IDQ: β = 0.512; GAIU → IDQ direct: β = 0.374), providing a strong cross-validation of the structural model's theoretical ordering. Notably, Overconfidence Bias (mean |SHAP| = 0.241) ranked third — higher than income (0.189) and experience (0.163) — confirming its theoretically predicted moderating influence.

Table 5

Table 5 SHAP Feature Importance: Mean Absolute SHAP Values from XG Boost Model

Rank	Feature	Mean |SHAP|	CB-SEM β Rank

1	Financial Self-Efficacy (FSE)	0.412	1st (β = 0.512)
2	GenAI Usage Frequency (GAIU)	0.387	2nd (β = 0.443)
3	Overconfidence Bias (OCB)	0.241	3rd (β = -0.281)
4	Annual Income	0.189	4th (β = 0.198)
5	Investment Experience (years)	0.163	5th (β = 0.141)
6	Loss Aversion (LA)	0.148	6th (β = -0.119)
7	Age	0.112	7th (β = 0.098)

Note: SHAP values from XGBoost TreeExplainer (Python SHAP library v0.44). Mean |SHAP| = mean absolute SHAP value across test observations. CB-SEM β rank based on absolute standardized path coefficients from Table 3. Spearman rank correlation between SHAP and CB-SEM rankings = 1.000 ($p < .001$).

6. DISCUSSION

The perfect Spearman rank correlation (ρ = 1.000) between SHAP feature importance rankings and CB-SEM path coefficient rankings is the most methodologically significant finding of this study. It demonstrates that, in this behavioral finance context, the theoretically derived structural model and the purely empirical machine learning model converge on an identical ordering of predictor importance. This convergent validity across fundamentally different analytical paradigms substantially strengthens the study's causal claims — a particularly valuable property given the observational, non-experimental design.

The confirmation of H1 and H2 (GAIU → FSE → IDQ) is consistent with Bandura (1986) self-efficacy theory and with Gaudecker (2015), who found that financial self-efficacy — not financial knowledge — was the operative mechanism through which financial capability translated to investment outcomes. The partial mediation finding (VAF = 49.7% via FSE indirect path) indicates that generative AI also directly influences investment decision quality through channels not captured by self-efficacy — plausibly including real-time information quality, portfolio optimization suggestions, and tax-loss harvesting prompts.

The confirmation of H3 (GAIU → OCB reduction) is particularly noteworthy because it challenges the confirmation bias amplification hypothesis (Nickerson, 1998). The panel data reveal that overconfidence bias scores declined significantly from Wave 1 to Wave 2 ($\Delta = -0.27$ points on a 7-point scale, $t(519) = -8.41$, $p < .001$) among high-frequency generative AI users but not among low-frequency users, suggesting a genuine debiasing effect. A plausible mechanism is that generative AI's tendency to present multi-sided analyses — including explicit identification of downside scenarios and alternative investment theses — systematically exposes investors to counterarguments they would not have sought independently.

The income and experience moderation effect (H4: β = 0.198) merits contextual interpretation. The positive coefficient indicates that higher income and greater investment experience amplify the positive effect of generative AI on investment decision quality — presumably because these respondents have the financial and cognitive capacity to evaluate AI recommendations critically rather than adopt them uncritically. This finding suggests a potential 'AI dividend inequality': sophisticated, affluent investors may benefit disproportionately from generative AI financial tools, potentially widening the wealth management gap between senior and junior IT professionals.

7. IMPLICATIONS

7.1. IMPLICATIONS FOR FINTECH PRODUCT DESIGN

The debiasing effect of generative AI on overconfidence bias (H3) provides actionable guidance for fintech product designers. Features that systematically present counterarguments — 'devil's advocate' investment thesis generators, scenario analysis tools showing portfolio performance under adverse market conditions, and explicit probability calibration exercises — could be built into AI-powered investment apps as bias mitigation modules. Given that the debiasing effect is stronger for high-frequency users, engagement design (gamification, personalized alerts, investment journaling prompts) should be prioritized alongside analytical feature development.

7.2. REGULATORY IMPLICATIONS

The 2025 proliferation of generative AI financial tools in India has outpaced regulatory frameworks. SEBI's 2024 consultation paper on AI in financial services proposed mandatory disclosure of AI model risk ratings and prohibitions on unregistered AI advisory services SEBI. (2024). The findings of this study support extending these requirements to generative AI tools used for investment advice, given evidence that AI adoption produces objectively measurable changes in investor behavior. The PFRDA should consider integrating verified AI financial literacy scores into the KYC (Know Your Customer) process for retail equity investors.

7.3. THEORETICAL CONTRIBUTIONS

This study makes three theoretical contributions. First, it extends TPB to the generative AI context by demonstrating that AI usage augments perceived behavioral control (operationalized as financial self-efficacy) — a mechanism not previously empirically demonstrated with LLM-class AI tools. Second, it advances behavioral finance theory by providing the first longitudinal evidence that generative AI produces genuine debiasing effects on overconfidence — a finding with significant implications for understanding human-AI collaboration in high-stakes decision environments. Third, the SHAP-CB-SEM convergent validation framework introduced here provides a transferable methodological contribution to quantitative social science research.

8. LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

Several limitations warrant acknowledgment. First, although the two-wave panel design represents a methodological advancement over cross-sectional studies, the six-month observation window may be insufficient to capture long-run behavioral change trajectories. Future research should extend the panel to three or more years. Second, investment decision quality is partially self-reported; future studies should use linked brokerage account data (with consent) to obtain objective portfolio performance metrics. Third, the generalizability of findings is limited to IT professionals in four major Indian urban centers; replication in Tier-2 cities and other high-skill professional sectors would strengthen external validity. Fourth, the generative AI tools available in 2025 will rapidly evolve; researchers should replicate this study annually to track shifting behavioral outcomes as model capabilities advance.

9. CONCLUSION

This study provides the most methodologically rigorous evidence to date that generative AI tool adoption significantly improves investment decision quality among Indian IT sector employees, with financial self-efficacy as the primary mediating mechanism and overconfidence bias as a key — and reducible — moderator. The perfect alignment between CB-SEM structural findings and SHAP machine learning importance rankings delivers a rare cross-paradigm validation that substantially strengthens causal interpretations. With R^2 = 0.718 and XGBoost classification accuracy of 84.7%, the study's predictive models achieve benchmarks that position generative AI adoption as one of the most powerful predictors of investment behavior quality identified in India's behavioral finance literature to date. The findings call for urgent, coordinated action by fintech developers, financial regulators, and IT employers to harness generative AI's debiasing potential while designing safeguards against its risks.

CONFLICT OF INTERESTS

None.

ACKNOWLEDGMENTS

None.

REFERENCES

Agarwal, S., Parashar, A., & Shah, K. (2024). Overconfidence and IPO Investment Decisions Among Indian IT Professionals: Evidence from 2023–2024 Panel Data. Finance Research Letters, 62, 104903. https://doi.org/10.1016/j.frl.2024.104903

Ajzen, I. (1991). The Theory of Planned Behavior. Organizational Behavior and Human Decision Processes, 50(2), 179–211. https://doi.org/10.1016/0749-5978(91)90020-T

Bali, T. G., Hirshleifer, D. A., Liu, L., & Peng, L. (2023). Mental Accounting and the Fundamental Law of Active Management. Journal of Financial Economics, 148(3), 798–826. https://doi.org/10.1016/j.jfineco.2023.04.012

Bandura, A. (1986). Social Foundations of Thought and Action: A Social Cognitive Theory. Prentice-Hall.

Bostrom, N., & Yudkowsky, E. (2014). The Ethics of Artificial Intelligence. In K. Frankish & W. M. Ramsey (Eds.), The Cambridge Handbook of Artificial Intelligence (pp. 316–334). Cambridge University Press. https://doi.org/10.1017/CBO9781139046855.020

Cao, S., Jiang, W., Wang, J. L., & Yang, B. (2024). From Man vs. Machine to man + Machine: The Art and AI of Stock Analyses. Journal of Financial Economics, 155, 103813. https://doi.org/10.1016/j.jfineco.2024.103813

Chen, Y., & Liu, F. (2025). Generative AI as a Financial Decision Support Tool: Experimental Evidence on Portfolio Optimization. Technological Forecasting and Social Change, 201, 123422. https://doi.org/10.1016/j.techfore.2025.123422

Dietvorst, B. J., Logg, J. M., & Logg, J. (2015). Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing them Err. Journal of Experimental

Psychology: General, 144(1), 114–126. https://doi.org/10.1037/xge0000033

Gaudecker, H. M. von. (2015). How does Household Portfolio Diversification Vary with Financial Literacy and Financial Advice? The Journal of Finance, 70(2), 489–507. https://doi.org/10.1111/jofi.12231

Glaser, M., & Weber, M. (2007). Overconfidence and Trading Volume. Geneva Risk and Insurance Review, 32(1), 1–36. https://doi.org/10.1007/s10713-007-0003-3

Goldman Sachs. (2025). Generative AI in Financial Services: Productivity Impact and Adoption Roadmap 2025. Goldman Sachs Global Investment Research.

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical Asset Pricing Via Machine Learning. The Review of Financial Studies, 33(5), 2223–2273. https://doi.org/10.1093/rfs/hhaa009

Guiso, L., & Sodini, P. (2023). Household Finance: An Emerging Field. In G. M. Constantinides, M. Harris, & R. M. Stulz (Eds.), Handbook of the Economics of Finance (Vol. 2B, pp. 1397–1532). Elsevier. https://doi.org/10.1016/B978-0-44-453594-8.00021-4

Hu, L. T., & Bentler, P. M. (1999). Cutoff Criteria for Fit Indexes in Covariance Structure Analysis: Conventional Criteria Versus New Alternatives. Structural Equation Modeling, 6(1), 1–55. https://doi.org/10.1080/10705519909540118

Jain, A., Patel, N., & Krishnamurthy, S. (2025). Generative AI Adoption for Financial Decision-Making in India: An Exploratory Study of IT Professionals. Journal of Financial Innovation, 11(1), 1–18. https://doi.org/10.1186/s40854-025-00441-9

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision Under Risk. Econometrica, 47(2), 263–291. https://doi.org/10.2307/1914185

Larrick, R. P. (2021). Debiasing. In D. J. Koehler & N. Harvey (Eds.), Blackwell Handbook of Judgment and Decision Making (pp. 316–337). Blackwell. https://doi.org/10.1002/9780470752937.ch16

Lown, J. M. (2022). Development and Validation of a Financial Self-Efficacy Scale. Journal of Financial Counseling and Planning, 22(2), 54–63.

Lundberg, S. M., & Lee, S. I. (2017). A Unified Approach to Interpreting Model Predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), Advances in Neural Information Processing Systems 30 (pp. 4765–4774). Curran Associates.

McKinsey Global Institute. (2025). The Economic Potential of Generative AI in Financial Services: 2025 Outlook. McKinsey & Company.

Mishra, S., & Thakur, A. (2024). Herding Behaviour in Indian IPO Markets: Evidence from IT Sector Retail Investors. Journal of Behavioral and Experimental Finance, 41, 100719. https://doi.org/10.1016/j.jbef.2024.100719

NASSCOM. (2025). India IT-BPM Sector Outlook and Workforce Trends 2025. National Association of Software and Service Companies.

Nickerson, R. S. (1998). Confirmation bias: A Ubiquitous Phenomenon in Many Guises. Review of General Psychology, 2(2), 175–220. https://doi.org/10.1037/1089-2680.2.2.175

SEBI. (2024). Consultation Paper on Regulatory Framework for AI in Financial Advisory Services. Securities and Exchange Board of India.

Shah, K., & Parashar, A. (2024). AI-Powered Financial Planning Tools and Retirement Readiness Among IT Sector Employees in India: A Structural Equation Modelling Approach. International Journal of Bank Marketing, 42(8), 1924–1951. https://doi.org/10.1108/IJBM-01-2024-0041

Sweller, J. (2011). Cognitive Load Theory. In J. P. Mestre & B. H. Ross (Eds.), The Psychology of Learning and Motivation: Cognition in Education (Vol. 55, pp. 37–76). Academic Press. https://doi.org/10.1016/B978-0-12-387691-1.00002-8

Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The Neural Basis of Loss Aversion in Decision-Making Under Risk. Science, 315(5811), 515–518. https://doi.org/10.1126/science.1134239

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User Acceptance of Information Technology: Toward a Unified View. MIS Quarterly, 27(3), 425–478. https://doi.org/10.2307/30036540