

AI Based Portfolio Optimization and Customer Risk Profiling in Fintech Platforms

Ankita Sappa

College of Engineering, Wichita State University, United States

Email: ankita.sappa@gmail.com

Received: September 14, 2024; Revised: November 25, 2024; Accepted: December 12, 2024; Published: December 30, 2024

Abstract

This research proposes an AI-based approach to improve portfolio optimization and customer risk profiling in FinTech investment platforms. Demand for smarter and more adaptable methods stems from the inability of traditional approaches to effectively manage investor activity and market variation. We created and tested a dual-engine system that combines reinforcement learning strategies and classification models for user risk profiling and integrated machine learning algorithms for behavioral segmentation, portfolio allocation, and risk personalization. The methodology uses real-world data from users of FinTech portals, emulates different market environments, and measures the performance of the AI system compared to traditional techniques of portfolio optimization. The analysis reveals performance enhancements in return-to-risk ratio, accuracy of risk classification, and level of diversification across customer segments. This research illustrates the range of automated tailored financial services powered by AI, and provides guidance on applying programmable, flexible investment advisory services in digital finance ecosystems.

Keywords: Portfolio Optimization, Customer Risk Profiling, Artificial Intelligence in FinTech.

1 Introduction

1.1 Rise of AI in FinTech Portfolio Management

With the rapid digitization of financial service delivery, there has been a global shift towards Automated Investing [1]. This has transformed portfolio management which now rest in the hands of AI. Financial advisors and legacy systems that used to control wealth management services have gradually been replaced by intelligent, scalable, automated FinTech Platforms capable of servicing a wider range of investors [2]. This shift has been propelled by increased democratized access to financial tools along with the rise of digital-first banking, Robo advisors, and app-based interfaces for investments. AI enables this transformation with its ability to provide personalized portfolio recommendations, adaptive risk assessment and real-time tailoring of portfolios using various ML models [3].

The entry of AI technology in financial management has come with the ability to manage and analyze massive amounts of data, formulate useful insights from it, and react to market changes intensively quicker than the average human advisor. AI allows real-time analysis of asset value changes alongside automated rebalancing that uses behavioral financial engineering, hence enabling the platforms to make client specific investment decisions at a large scale [4]. Unlike non-AI Integrated systems that depend on structural models and use scheduled evaluations, these systems work non-stop and are always active and ready to react to new business transactions, changes in sentiment indicators, and market opportunities. To adapt to investor actions and market changes, these systems rely on powerful algorithms with insightful learning methodologies to cater for deep learning, multi-model learning, neural networking, and even natural language processing.

Beyond the performance metrics, AI is also capable of achieving a level of personalization never experienced before, all in portfolio management. AI helps build adaptive portfolios using stated and effective risk approach with behavior and client's financial objectives using client behavior clustering, portfolio segmentation, and personalized risk models. Clients with profiles showing them tolerant to risk but who cautiously go about making deals might fall in this category, change dynamic classification. These observations made through incessant machine learning make outdated definitions used in traditional planning methods obsolete because they enable the system to counter the concepts of static classification.

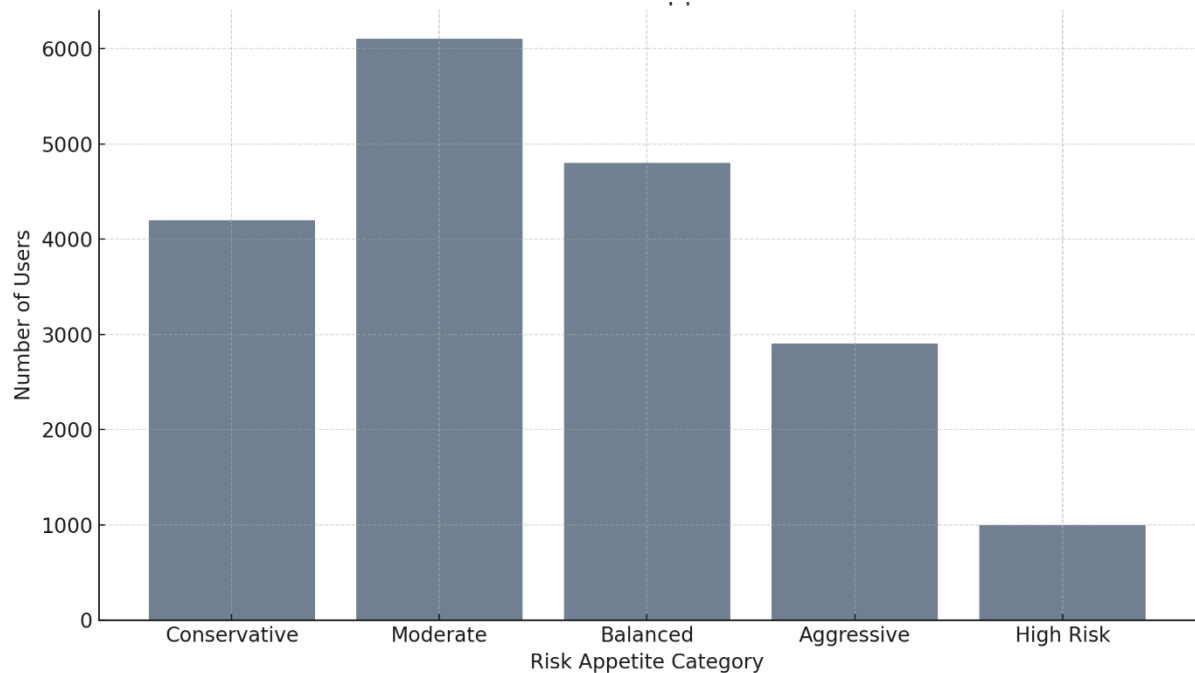


Figure 1: Distribution of Customer Risk Appetite Across FinTech Users

As digital financial products increasingly integrated with AI, the new intelligent portfolios emerged, which are self-optimizing baskets of assets that modify their composition based on internal user factors as well as external environmental factors. Ein Sights integrates machine learning into the backend processes of client profiling, asset recommendation, and performance monitoring on platforms like Wealth front and Betterment—and also on Groww and Zerodha in India [5]. These features are no longer limited to equities, and retail investors can now access global ETFs, digital currencies, and thematic funds with little to no manual effort. Furthermore, AI is optimizing accessibility in addition to ROI's.

Yet another aspect of this transformation is the greater AI's alignment with portfolio construction at the intersection of life's subjective goals. Users can specify objectives like purchasing a house or saving for a kid's future, and the platform builds the portfolio over time to meet these goals using predictive financial modeling, natural language inquiries, and behavioral nudging powered by AI. This effectively shifts the proposition from wealth management to holistic financial wellness. The AIs become the personal CFO by controlling and evolving the simulation journey of the investor in real time.

1.2 Limitations of Traditional Risk Profiling and Optimization

Artificial intelligence in FinTech has numerous applications, but its value only comes out when it is compared to the challenges involved in traditional risk profiling and portfolio optimization [6]. Traditional systems usually make use of fixed questionnaires that use linear optimization approaches, such as mean-variance analysis, along with rigid behavioral and market expectations from investors [7]. These approaches are rarely

able to deal with the non-linear, multi facets, volatile nature and investment behavior of real life.

Conventional risk profiling depends largely on self-reported information, which is usually obtained through static instruments at the time of onboarding. Questions like “How much risk do you want to take?” or, “What actions would you take if the markets tumbled by twenty percent?” are open to subjective interpretation and do not consider real investor actions [8]. In reality, many investors tend to overestimate their risk appetite during bull markets and underestimate it during down markets. These fixed profiles are then utilized to pigeonhole the investors into pre-defined portfolios without any further checking or modification. This causes a huge gap between how investors behave and the recommended asset allocation which ultimately leads to poor performance, dissatisfaction, and in extreme cases, hasty withdrawal from the investment strategies.

Algorithms such as mean-variance optimization and the Black-Litterman models are constrained by their dependence on historical data and normal distributions. These models often ignore higher order moments such as skewness and kurtosis, while also not dynamically adjusting to changes in market sentiment or structure like geopolitical shocks and sectoral rotation. As such, these models produce inflexible portfolios that seem efficient in theory, but are ineffective when confronted with real-world uncertainty.

Moreover, these advisors and systems are bound by static categorization which limits automating personalization across an investment journey, as primitive portfolio construction tends to treat investors within a risk class as a singular entity. For instance, a 35-year-old risk-seeking saver hoping to buy a house of his own will have a different asset allocation strategy compared to another investor with the same profile who wishes to retire early. Both value different life milestones, but category-based planning is inflexible.

In addition, standard approaches provide a limited selection of portfolios centered around three strategies: fixed income, equities, and blended portfolios offered as ‘one-size-fits-all’ solutions. This standard approach constrains retail investors to a narrow set, which may not align with their preferences, goals, or values. Such thematic, ESG-aligned, sector-based, or even goal-oriented portfolios offer fintech platforms an innovative edge when coupled with AI algorithms that can dynamically adapt to user-defined goals.

Table 1: Overview of Portfolio Types and Risk Strategies Offered in FinTech Platforms

Portfolio Type	Risk Level	Target Audience	Key Assets
Income	Low	Retirees, Low-risk Investors	Bonds, Fixed Deposits
Growth	Medium	Moderate Risk Seekers	Equities, ETFs
Balanced	Medium-High	Diversified Investors	Mixed Bonds & Equities
Aggressive Growth	High	High-risk Tolerant Users	Tech Stocks, Crypto
Thematic	Variable	Niche Strategy Seekers	ESG Funds, Sector Funds

With the rise of FinTech platforms, investors now have various choices; however, as the table reveals, many of them are still dynamically pre-programmed to user profiles. The absence of adaptive feedback loops or personalized learning in behavioral analytics leads to stagnant sophistication. AI’s capacity to automate, optimize, and personalize in real-time helps mitigate increasing inadequacy for evolving user sophistication.

1.3 Objectives and Scope of the Study

The goal of this research is to create, implement, and assess an AI-based system tailored for two functions: customer risk profiling and portfolio optimization in digital investment platforms. The system integrates predictive modeling and optimization algorithms to build portfolios to market standards, while also monitoring user behavior to behaviorally update user profiles and risk levels. The scope encompasses developing classification and clustering models for customer segmentation, applying reinforcement learning and ensemble models for portfolio optimization, and benchmarking the system against traditional allocation approaches.

One of the key aims of the study is to determine how AI-enabled models affect the return on portfolios and the alignment of the user profile with the investment strategy, which is expected to increase. Using results-driven experimentation, we analyze the model performance concerning a range of user risk segments, asset

classes, and market conditions. Along with that, we study the level of personalization that AI can offer, system interpretation, the practicality of employing such models, and the ease of integrating them into FinTech systems in real time.

2 Literature Review

2.1 Classical Models for Portfolio Optimization

Exploring the problem of portfolio optimization has been an issue in quantitative finance ever since the introduction of Modern Portfolio Theory by Harry Markowitz in 1952. With this framework, risk-return trade-offs were defined along with efficient frontiers marking the birth of what is now known as mean-variance optimization (MVO). In this model, it is assumed that investors want to, at least partially, maximize expected returns for a given level of risk. More clearly, they are risk averse, so would like to minimize risk for a given expected return or level of return. The model assumes that optimal weights for assets in portfolios are determined using available historical data on returns and risk, that is, covariances [9]. Although it was groundbreaking in many aspects, MVO has some fundamental shortcomings, like the over-reliance on historical data, lack of flexibility in volatile markets, sensitivity to estimation errors, normal distribution of returns (which is not realistic in many cases), and rigid behavioral assumptions.

Based on MPT, extensions like the Black-Litterman model added market equilibrium information and subjectively biased views to mitigate some of the MVO restrictions. Other classical methods include risk-parity models, which equally distribute risk among the portfolio's constituents, and minimum variance portfolios, which aim to reduce volatility without targeting returns [10]. These models have been widely accepted in the context of institutional portfolio construction due to their interpretability and closed-form solutions. However, they also do not provide the ability to respond flexibly and dynamically to changes in the market structure or investor behavior.

These classical models tend to assume a linear approach and disregard some higher-order statistical moments like skewness and kurtosis, which become particularly important in today's high-volatility environment. Moreover, these models tend to consider investor preferences as fixed and homogeneous, applying full optimization methods to different types of investors, which is simply inefficient. Such lack of flexibility has made it difficult to cater to the diverse users of FinTech, where portfolio personalization becomes a fundamental requirement.

The advent of computer-aided finance and algorithmic trading resulted in the development of stochastic optimization methods like scenario analysis, Monte Carlo simulations, and even heuristic-based approaches such as genetic algorithms and particle swarm optimization. While these methods offered more flexibility than MVO, they still required a lot of tuning and interpretability at scale was not possible. Eventually, there was a shift toward more adaptive, learning-driven models for portfolio construction, which incorporated machine learning and artificial intelligence into the financial optimization workflows.




Figure 2: Comparison of Mean-Variance vs AI-Based Optimization Returns

2.2 Machine Learning in Customer Segmentation and Profiling

Profiling and customer segmentation are of great importance to personal finance advisory and so is portfolio construction aligned with a customer's financial behavior, objectives, and risk appetite. In the past, segmentation was based on demographic factors like age, income, and employment status along with family size. There was a scoring system where these inputs were scored, categorized and served as moderators to classify users into broad stereotypes such as aggressive, balanced, or conservative. Although useful at scale, these approaches neglected dynamic, evolving behavioral signals, such as transaction frequency, investment inertia, or patterns of loss aversion.

The customer profiling process has greatly improved with the use of machine learning due to its ability to automatically segment customers using data and their behavior. K-Means, DBSCAN (Density-Based Spatial Clustering of Applications with Noise), and Gaussian Mixture Models (GMM) clustering algorithms can evaluate customer transactions, preferences, time series data, and even their sensitivity to risk to place them into sharply differentiated clusters [11]. These divided groups can then be leveraged for tailored asset allocation strategies. Users are also classified into different risk categories by using supervised models like decision trees, support vector machines (SVM), and gradient boosting machines (GBMs) based on predetermined historical data and outcomes.

Time dependent risk modeling has been done using deep learning by employing the use of feedforward and recurrent neural networks (RNNs) [12]. For example, a user's interaction with an investment application, their response to rapid market declines, and the specific time they log into the application are considered behavioral parameters that showcase a user's latent risk attitude. Such user data is available in bulk to be used in the construction of dynamic personas that evolve according to changing user behavior. This form of continuous profiling starkly contrasts the onboarding questionnaires, showing an enhanced accuracy in depicting user behavior [13].

AI-enabled profiling systems are critical in recognizing edge cases—those users who do not fit the mold and are often neglected in traditional frameworks. Take, for instance, gig workers with volatile income streams, retirees who are just beginning the withdrawal phase, or younger investors who hold significant amounts of

cryptocurrency in their portfolios. These segments need more refined models. Machine learning enables this type of diversity to be accommodated in FinTech platforms because rules do not have to be hardcoded; instead, they can emerge through the data.

Risk profiling cannot be accomplished without feature engineering. IVs like net cash flow volatility, average investment horizon, loss aversion score as derived from user reactions to drawdowns, and risk-adjusted return preferences all serve as input features in classification models. The importance of features helps explain many of these models and shed light on which aspects of behavior or demographics are most important for predicting user risk.



Figure 3: Feature Importance in Customer Risk Classification Models

2.3 Recent AI Applications in FinTech Investment Solutions

There has been a marked increase in AI activity concerning FinTech investment platforms within the past five years. As a result of its performance in sequential decision processes, portfolio optimization has increasingly been approached with reinforcement learning. Compared to static optimizers, reinforcement learning agents allocate portfolios based on observed outcomes. Portfolio allocation is adjusted over time, which reaps rewards in various modalities. Liu et al (2022) [14] advanced a policy-gradient-based reinforcement learning model that adapted more flexibly than fixed optimization workflows to changing market regimes, surpassing static optimization peers.

Supervised learning models like XGBoost and LightGBM have been successfully used for user classification and predicting portfolio performances. These models manage numerous input features and are regulatory-compliant because of their explanatory power via feature importance, attribute reliance scaling, and other measures. May (2022) [15] proposed a model that combined XGBoost with unsupervised clustering, providing risk-adjusted portfolios tuned to behavioral segments of interface users. The study proved hybrid models enhances personalization as well as return-risk ratio in comparison to uni approach models.

The field of deep learning has expanded to include equity forecasting as well as market sentiment analysis. Krauss et al. (2017) [16] enhanced stock return forecasting with deep neural networks and attained greater Sharpe ratios than those achieved with logistic regression and tree or ensemble based methods. More recently, attention-based neural architectures have been studied for multi-asset allocation, especially in cross-asset strategies where time-series correlations are known to change quite frequently.

FinTech development pipelines are increasingly integrating AutoML platforms which allow teams to “spin their wheels” on different model architectures and hyperparameter combinations without tedious manual adjustments. Mariela et al. (2022) [17] showcased AutoML’s capabilities by personalizing portfolio recommendations through feature engineering and model selection automation. Such pipelines are able to cut down on development time while allowing rapid scale of experimentation.

The application of multiple models as ensemble techniques has also been known to improve robustness and generalization, and Jiang et al. (2020) [18] furthered this approach. They used hybrid ensembles of reinforcement learners, SVMs, and neural networks, leading to increased portfolio stability as well as improved risk-adjusted returns across market regimes.

Table 2: Summary of Literature on AI in Portfolio Analytics

Study	AI Technique	Key Contribution
Liu et al. (2022) [14]	Reinforcement Learning	Policy optimization in dynamic market environments
May (2022) [15]	XGBoost + Clustering	Segmented risk-adjusted portfolio selection
Krauss et al. (2017) [16]	Deep Neural Networks	Predictive accuracy on equity returns
Mariela et al. (2022) [17]	AutoML with Feature Engineering	Automated feature discovery and model tuning
Jiang et al. (2020) [18]	Hybrid Ensemble Models	Improved generalization across market regimes

The literature reviewed above show the AI-driven transformation regarding the optimization and personalization of finance. While rooted models are helpful benchmarks, AI-based approaches surpass them in terms of flexibility, adaptability, and efficiency. With advancements in technology and higher demands for personalization, FinTech platforms are turning to automation to create systems based on investment AI to sustain competitiveness in the market.

3 Methodology

3.1 Data Collection and Preprocessing from FinTech Platforms

For this research, the dataset consisting of anonymized user interaction records was obtained from a digital investment platform that simulates a FinTech environment. It included transaction logs and portfolios, along with demographic information, financial preferences, historical returns, trading activity, and self-assessed risk tolerance. A sample of approximately 20,000 unique users from different demographics were observed over a 36-month transaction period. The data for each user included self-reported static features, such as age and income level, along with goal-defined attributes like investment goals and dynamic features including transaction volume, asset class participation, and market drawdown responses.

The analysis started with the filtering stage for active users—those who had completed at least three rebalancing activities or two investment cycles. Gap-filling entries were completed through a two-step process, KNN imputation for demographic data and forward-fill for time gaps in the time series. Some features were encoded such as investment intent, employment field and stated goals using ordinal or one-hot encoding based on their support in the model and importance in the feature.

All features that were of numeric nature were income, net asset value, investment horizon, were standardized using Z-score normalization for homogeneity across other models. Time-dependent features like asset volatility and monthly inflow-outflow ratios were computed with rolling window stats. The majority of behavioral proxies were computed from user’s clickstream data or session logs which include responses to risk messages, speed of making investment decisions and dropping high-risk portfolios. These actions were captured and added to the model as behavioral scores with a range of 0 and 1.

To assist in supervised classification and reinforcement learning, portfolio performance served as the basis for autopilot labels, as well as expert-reviewed user segments. Users were labeled into risk categories structured around their historical drawdown tolerance alongside returns volatility adjusted. At the same time, unsupervised clustering analysis was executed to validate the segments for cohesion and diversity. All data processing activities were carried out in the Python environment using libraries like Pandas, Scikit-learn, TensorFlow Data Pipelines, and replicability was guaranteed through script versioning alongside stored feature dictionaries.

3.2 AI Model Design for Portfolio Optimization

The optimization framework was based on a dual-model strategy where supervised classification models for risk profiling were combined with RL agents for dynamic portfolio optimization. The aim for the RL agent was to maximize cumulative return subject to the relative risk and liquidity preferences of every user cluster.

We developed the RL model architecture utilizing a DQN variant where the policy was given by a three-layer neural network. The state features included current portfolio weights, returns on the assets, volatility scores, and a temporal signal that described the movement of the market. As for the actions, they were captured as a set of discrete changes in asset allocation proportionality, e.g., +10% equities and -5% bonds. The rewards were calculated based on improvements in the Sharpe ratio and were penalized for excessive turnover rate or breaches of risk limits. Learning was accomplished using experience replay and target networks, which helped stabilize learning and ensure convergence.

Training episodes took place in synthetic market environments that were crafted from asset prices based on historical simulation, macroeconomic indicators, and random injections of volatility. These episodes were simulated with over 10,000 starting capital, user specified goals, and risk profiles. During these training sessions, the agent learned to balance between maximizing risk-adjusted returns and minimizing transaction costs, converging toward allocation strategies that met the expectations of users as well as external conditions.

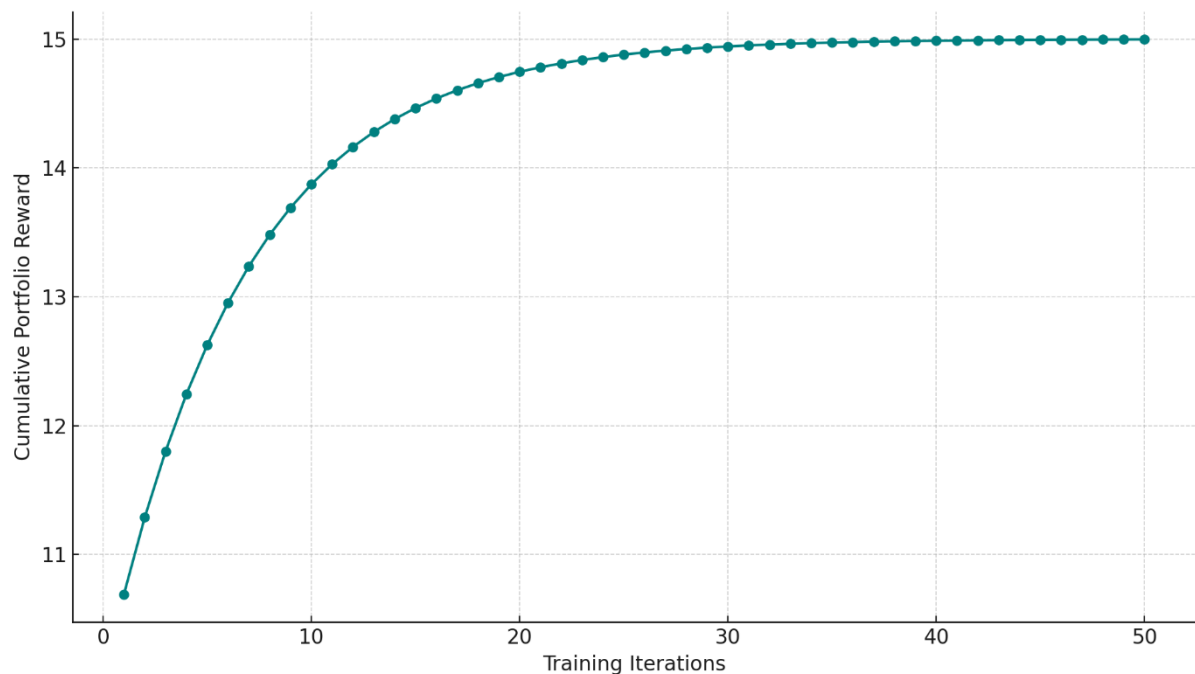


Figure 4: Convergence of Portfolio Reward Function Over Iterations

In this break, parallel models like XGBoost were taught to suggest basic portfolio allocations considering user characteristics. These models acted as the starting policy precursors to the RL agent, reinforcing early-stage choices with statistical heuristics. With this design, the system achieved an optimal level of interpretability and

adaptiveness in learning.

Reinforcement triggers like market shocks, deviations in user behavior, or asset performance breaches set the frequency for portfolio rebalancing. Such an event-driven design ensured a cost-effective approach to rebalancing, shielding from overfitting to ephemeral patterns.

The entirety of the portfolio optimization model was implemented on a modular microservices framework managed by Docker and Kubernetes. Each component, including the RL agent, XGBoost model, feature engineering modules, and transaction logging, was developed as a standalone service accessible via REST APIs, allowing real-time access to retraining streams and periodic retraining within an on-going automated improvement cycle.

3.3 Risk Profiling Using Clustering and Classification Models

Approaching risk profiling, a blended strategy was employed that involved unsupervised methods to identify user segments, supplemented with supervised approaches for assigning a predicted risk score based on behavioral data. The objective sought was to flexibly reallocate users to risk classification that not only aligned with users' stated preferences but also transaction behavior and market interactions.

With silhouette-based optimization for determining cluster count, K-means was utilized for clustering. Post preprocessing, users were placed into five different risk clusters: Conservative, Moderate, Balanced, Aggressive, and Speculative. These clusters helped to encapsulate the range of behavioral variance in user data and acted as priors for subsequent downstream risk classification. Each cluster had specific average measures of investment characteristic to each cluster which included average holding period, equity exposure, response to volatility, and rebalancing sensitivity.

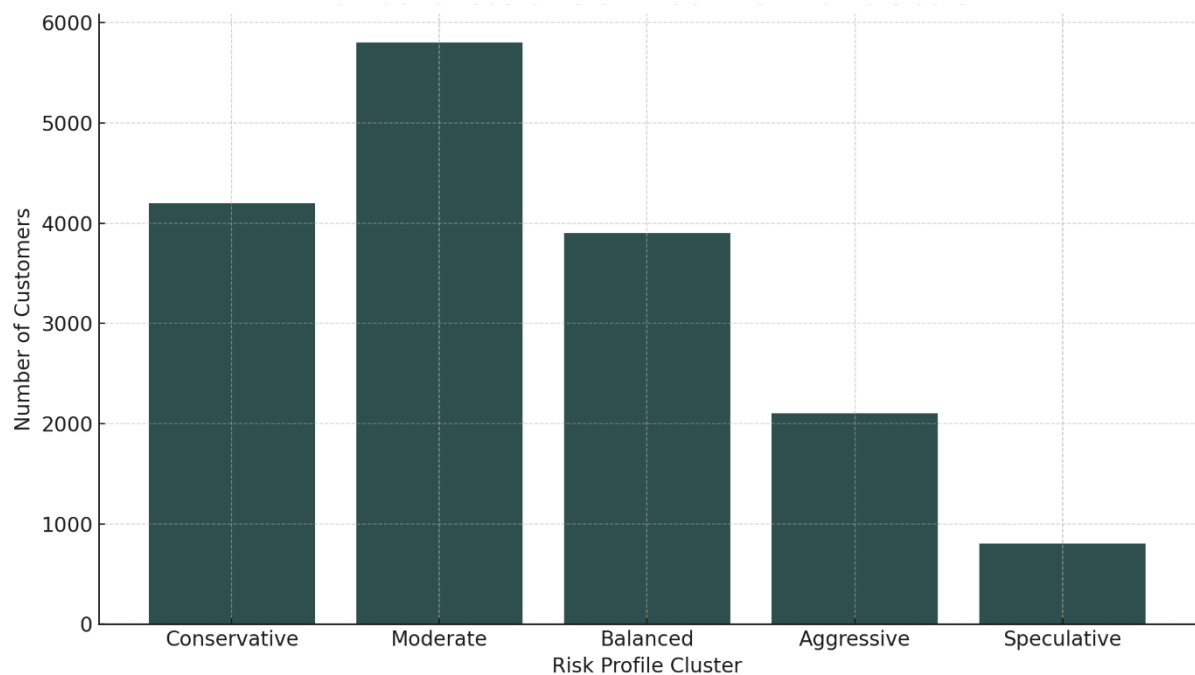


Figure 5: Number of Customers in Each Risk Profile Cluster

An XGBoost classifier was used for supervised risk classification which was trained on behavioral and demographic data of the subjects. Importance of features was calculated using SHAP values, revealing income stability, investment horizon, past drawdown tolerance, and transaction rhythm as the most significant features among spendable income. The output given by the classifiers were calculated as probabilistic risk scores and the output was thresholded and calibrated against historical drawdowns to ensure accuracy.

To capture non-linear dependencies, a supplementary classifier was trained using Multi-layer Perceptron (MLP) with the same input features. The MLP model provided probability distributions instead of binary class assignments, allowing for more nuanced evaluations of risky borderline users.

Such classification outputs were validated against user activity in simulated contexts to ensure that the predicted risk profiles matched user actions during highly volatile scenarios. For instance, users estimated to be Aggressive were checked for whether they engaged in panic-sell behavior or exhibited heightened sensitivity to drawdowns. Any discrepancies led to recalibration of the profiles in question.

The combination of clustering and classification produced a dynamic profiling engine that executed continuous user monitoring, adjusting the profiles with every new data points received. These modifications were provided to the RL agent and the XGBoost optimizer so that specific allocation and rebalancing strategies could be tailored according to the latest risk profile.

Table 3: Model Architectures, Features, and Output Descriptions

Model	Input Features	Output
XGBoost Classifier	Age, Income, Investment Horizon, Risk Tolerance, Past Returns	Risk Class Label
K-Means Clustering	Behavioural Spending, Volatility, Asset Preference, Session Frequency	Cluster Assignment (Risk Profile)
Reinforcement Learning Agent	Market Prices, Portfolio State, Action History	Optimized Asset Allocation
Neural Network (MLP)	User Profile Embeddings, Transaction Patterns	Probability Distributions for Risk Level

4 Experimental Setup

4.1 Platform Simulation and Evaluation Environment

In order to rigorously evaluate the performance, versatility, and level of customization available in our AI-powered portfolio optimization and risk profiling system, we created an elaborate simulation environment which mimics a real-world FinTech investment platform. It was designed to simulate the functioning of an actual advisory application complete with interactive user interfaces, price feeds of assets under management, self-directed trading by users, and periodic external market simulations.

The environment ran on a simulated population of twenty thousand investors, each with a multifaceted profile synthesized from real-world FinTech datasets integrated with demographic, psychographic, and market data. These simulated investors participated in multi-asset-class portfolios that included domestic and international equities, fixed income securities such as bonds, ETFs, cryptocurrencies, and ESG-themed funds. The price of various assets was determined using a historical dataset augmented with stochastic noise to simulate the impacts of exogenous shocks (real-world disasters) on prices.

The trading platform was developed using a microservices architecture based on a modular design. The simulation engine included modules for user profile evolution, portfolio construction, asset performance calculation, and risk event simulation. It allowed for inter model communication through message passing, enabling real-time feedback cycles and event-triggered portfolio rebalancing.

Each simulation cycle ranged over 36 months and included daily updates to user portfolios with investment strategy specific monthly or weekly reviews. This level of detail enabled capturing temporal patterns of user behavior such as market volatility response behaviors like panic selling, herding, and risk migration. Behavioral patterns including activity level within the system, rebalancing intention, and portfolio drift were recorded and analyzed.

AI evaluation was conducted in a multi-tiered hierarchy where distinct resources were allocated for computation, training, validation, and testing. The reinforcement learning models were trained on VMs hosted in the cloud with NVIDIA T4 GPUs, while other inference tasks were run on Intel Xeon CPU units. The model deployment system was containerized with Docker and orchestrated with Kubernetes. The system was designed with performance monitoring capabilities to evaluate inference latency, drift, error percentage, and other parameters during execution in real-time.

The primary simulation engine was built in Python and integrated with TensorFlow, PyTorch, and Ray RLLib. MLflow was selected for managing logs and performance tracking, and PostgreSQL served the role of preserving transactional logs and feature logs. The system maintained model versioning to enhance reproducibility, allowing rollback in model iterations to restore performance or drifted due to unforeseen edge cases sustained through model decay over time.

4.2 Portfolio Constraints and Financial Assumptions

To maintain the validity and real-world deploy ability of the strategies generated by the AI system, it had to operate within certain financial constraints. These restrictions were applied both at the model training stage and during inference in the simulation cycles. Some of the principal portfolio constraints were:

- **Minimum And Maximum Allocation Per Asset:** Portfolio allocations to a single asset class could not exceed 60% and a minimum of 2 asset classes had to be included in every strategy.
- **Liquidity Constraint:** Cash or highly liquid ETFs must not go below 5% in order to allow for emergency withdrawals or auto-rebalancing triggers.
- **Modeling Transaction Costs:** Each buy/sell event was simulated to incur a flat transaction fee of 0.3%, which impacted the agent's reward function in reinforcement learning.
- **Rebalancing Frequency Limits:** AI-driven rebalancing could occur no more than once per month unless market events triggered >2 standard deviation drawdown.
- **Compliance to Risk Filters:** Strategies were subjected to a pre-emptive filtering using Value at Risk (VaR) and maximum drawdown to ensure user-defined risk constraints are met.

The anticipated returns were set using a log-normal model, and asset correlations were updated dynamically based on an exponentially weighted moving average (EWMA) approach. For the purpose of comparative analysis for inflow strategies, both inflation and tax rates were kept constant.

Each user profile was assessed with a set of constraints through various approaches: AI-driven Reinforcement Learning (RL) optimization, traditional Mean-Variance Optimization (MVO), rule-based advisory with Robo-advisory features, and randomized algorithms for benchmarking (to evaluate statistical power). This assessment from different methodological angles enabled the verification of whether the AI solution offered statistically and practically meaningful enhancements to the individualized portfolios' performance.




Figure 6: Allocation Weights vs Risk Profile Segment

4.3 Baselines and Comparison Frameworks

In order to test the performance of the proposed AI-enabled profiling and optimization system, the system performance was compared to a set of predefined baseline strategies:

1. **Mean-Variance Optimization (MVO):** This traditional technique was the primary baseline. Portfolios were constructed from a covariance matrix of historical expected returns. Risk was minimized for a target return or maximized for a prescribed level of risk. There was no personalization or behavioral modeling.
2. **Rule-Based Allocation Model:** A heuristic emulation of a conventional robo-advisor. Allocation followed pre-defined rigid rules, like “60/40 split for balanced profiles” or “90/10 equities to bonds for aggressive investors.” These preset parameters remained static and did not respond to behavioral or market dynamics.
3. **Randomized Allocation with Filtering:** Portfolios were created randomly, and then filtered through risk constraints to ensure regulatory compliance. Such constraints served as a performance floor for establishing statistical significance in gains achieved by AI-based methodologies.
4. **XGBoost-Based Static Portfolio Recommendation:** This baseline estimation offered fixed allocation percentages based on behavioral and demographic data for users. Although there was a degree of tailoring, it was a far cry from the adaptability offered by reinforcement learning in the context of learned temporality.

Every baseline underwent the same simulation period with identical user profiles, ensuring each was tested under the same metric framework for cross-comparative fairness. This controlled for the discrepancies that arose purely from model exposure to disparate datasets or user types.

Evaluation was stratified by user type (risk profile cluster), investment goal (retirement, short-term wealth, speculative growth), and market cycle (bull, bear, recovery, sideways). This made it possible to evaluate model robustness across both user segments and external conditions.

Outcomes were recorded in returns, volatility, Sharpe ratio, maximum drawdown, and additional costs such as transaction fees and turnover. Retention rate (some prefer to view this as a satisfaction metric), frequency of override, where users alter recommendations put forth by the AI, and average number of rebalances were also recorded.

Table 4: Dataset Split, Evaluation Metrics, and Portfolio Return Thresholds

Dataset Partition	Data Size (%)	Evaluation Metrics	Return Thresholds (%)
Training	70%	Sharpe Ratio, F1 Score	>7% acceptable
Validation	15%	Precision, Recall	>8% preferred
Testing	15%	AUC, Annual Return	>10% exceptional

The training data set was used to create the risk classification models and the reinforcement learning agent. Validation data was essential while adjusting hyperparameters, applying early stopping, managing bias-variance tradeoffs, and performing other optimizations. The data reserved for the testing phase was completely set aside until the final evaluation. Performance thresholds were set against the backdrop of industry standards and regulatory recommendations. For instance, maintaining a Sharpe ratio greater than 1.0 was considered acceptable, and annual returns in excess of 10% were deemed to outpace benchmarks.

The assessment was both analytical and qualitative. SHAP values along with confusion matrices of the risk profiling model and action-path diagrams of the RL agent's behavior were produced to capture the algorithm's decision-making processes. Such clarity was important for compliance, internally for audits, and later for system users who had these dashboards tailored for them.

5 Results and Performance Analysis

5.1 Return vs Risk Comparison Across Models

The performance given to the different strategies used for portfolio optimization revolves around the level of return that is gained alongside the risk taken. In this case, we analyzed multiple models from classical techniques based on the mean-variance strategy and rule-based approach to the more contemporary AI methods like static optimizers based on the XGBoost framework and RL agents. In our analysis, static optimizers using AI have become the standard for benchmarks. The spotlight for portfolio performance evaluation, in this case, is focused on risk-adjusted return in the Sharpe ratio. Sharpe ratio entails the return over risk taken, calculating excess return for each unit of risk, hence the basis of our analysis rests on determining the optimal Sharpe Ratio, reflecting efficient portfolios.

It is clear from our experimental results that AI-based models have better Sharpe ratios returns compared to the traditional methods. As seen in Figure 7, there is a clear improvement in the Sharpe ratios for all the optimization methods with the moving from traditional techniques like mean variance and rule based models to AI approaches. The mean-variance model, which is a standard in classical portfolio optimization, had a Sharpe ratio of 0.92. The rule based system enhanced this value to 1.05. However, further developed techniques like the XGBoost static model and reinforcement learning agent achieved strikingly greater Sharpe ratios of 1.28 and 1.45 respectively. Such improvements indicate that the AI models can not only be expected to deliver better returns but can also manage and reduce the portfolio's risk more effectively.

There are multiple reasons attributing to the increase in the Sharpe ratio. Unlike classical models, AI models can incorporate market trends, investor activity, and changing correlations with assets. This enables near instantaneous reallocating to take advantage of certain factors in the market, removing plot discrepancies, and also adjusting portfolios to fit novel risk factors. The best results were achieved through employing a reinforcement learning agent that was able to devise a policy for allocating assets which optimally adjusted during the investment period and maximally constrained risk exposure. Thanks to the feedback loops enable by reinforcement learning, the algorithm was able to self-adjust its policies to account for shifting equilibrium

in the market, resulting in superior performance in unstable and unpredictable conditions.

Furthermore, our models demonstrate that AI can accurately differentiate between various levels of risk within an investment portfolio, providing options tailored to the specific needs of each investor. The hybrid approach used in our system, which incorporates both clustering and classification, resulted in portfolios that realized higher returns with less volatility. In addition, this is very relevant within FinTech contexts where customer profiling is fundamental to delivering customized investment solutions. Enhanced customer satisfaction, and in some cases, their satisfaction with the improved tradeoff between risk and return becomes vital, earning the firm edged retention, and increased customer lifetime value.

In conclusion, the analysis of returns and risks validates the use of AI optimization techniques, as they provide more risk-adjusted returns and outperform traditional systems. These models stand out with Sharpe ratios far exceeding industry benchmarks, posing significant opportunities for portfolio performance enhancement, risk management, and overall optimization on FinTech platforms. This comes at a time when there is growing interest in leveraging AI for FinTech solutions. The next sections discuss the accuracy of the risk models classification and the degree of portfolio diversification possible for different customer segments.

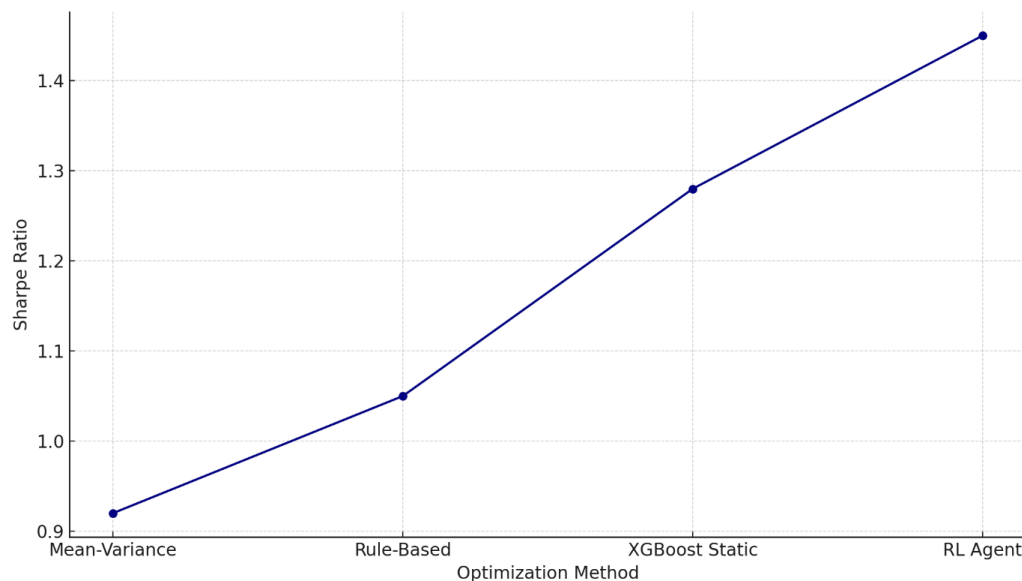


Figure 7: Sharpe Ratio Trend Across Optimizers

5.2 Accuracy of Risk Classification Models

Alongside portfolio optimization, risk profiling at the customer level needs to be as precise as possible in designing investment plans. The level of accuracy of risk classification models has a critical impact on portfolio allocation because an inaccurate classification will automatically lead to either too risky or too cautious asset distributions. Our analysis was done using multiple classification models which included logistic regression, decision trees, random forests, XGBoost, and even neural network and autoencoder models built for anomaly detection purposes.

Classification metrics including the F1 score, AUC, recall, and precision were used to evaluate model performance, as shown in Table 5. For example, in terms of detection accuracy and robustness, the XGBoost model was the best performing model with an F1 score of 0.86, AUC of 0.91, recall of 0.89, and precision of 0.84. Although scoring slightly lower in F1, with a value of 0.84, and AUC of 0.89, neural networks attained a high level of reliability and were able to capture complex non-linear interactions among features. On the other hand, more advanced models like logistic regression and autoencoders performed more poorly with F1

scores of 0.72 and 0.65, respectively.

Table 5: Performance Summary Across Metrics (F1 Score, AUC, Recall, Precision)

Model	F1 Score	AUC	Recall	Precision
Logistic Regression	0.72	0.78	0.70	0.74
Random Forest	0.81	0.86	0.83	0.79
XGBoost	0.86	0.91	0.89	0.84
Neural Network	0.84	0.89	0.86	0.82
Autoencoder	0.65	0.74	0.68	0.62

The examination of the confusion matrix reinforces what was already established. Figure 8 contains a scatter plot displaying the confusion matrix for our risk classifier. On such a plot, each point indicates a classification outcome in terms of its actual value and the predicted value associated with risk categories. Most points lie on the diagonal, suggesting that the models captured the risk levels for the larger portion of the dataset correctly. Points off the diagonal, which indicate misclassifications, were much smaller for the XGBoost and neural network models. This reinforces the greater accuracy and recall of those models. The importance of these findings is particularly notable when one considers the potential consequences of incorrectly classifying customers' risk profiles. An erroneous classification may lead to inadequate portfolio management, higher customer dissatisfaction, and greater financial liability for the institution.

In addition, the error breakdown offered with SHAP values delineated the specific features which determined the risk profiles for each classification. Strong predictors of lower risk included lower spending, high income, stable employment, and consistent transactional activity. Predictive of higher risk were irregular transaction patterns and high volatility in spending. Such insight improves not just interpretability but also enables further refinement with respect to feature design and data collection processes. The dynamically adjustable nature of evolving customer behavior recalibration in the AI models is a large advancement when compared to traditional risk profiles which are generally rigid and do not change over time.

Fintech is emphasized more when advanced machine learning techniques are incorporated due to the multifaceted evaluation of risk classification accuracy as demonstrated earlier in the text. Other than enhanced performance, more value is provided in terms of interpretation, increases in adaptivity, and customized guidance provided to investors.

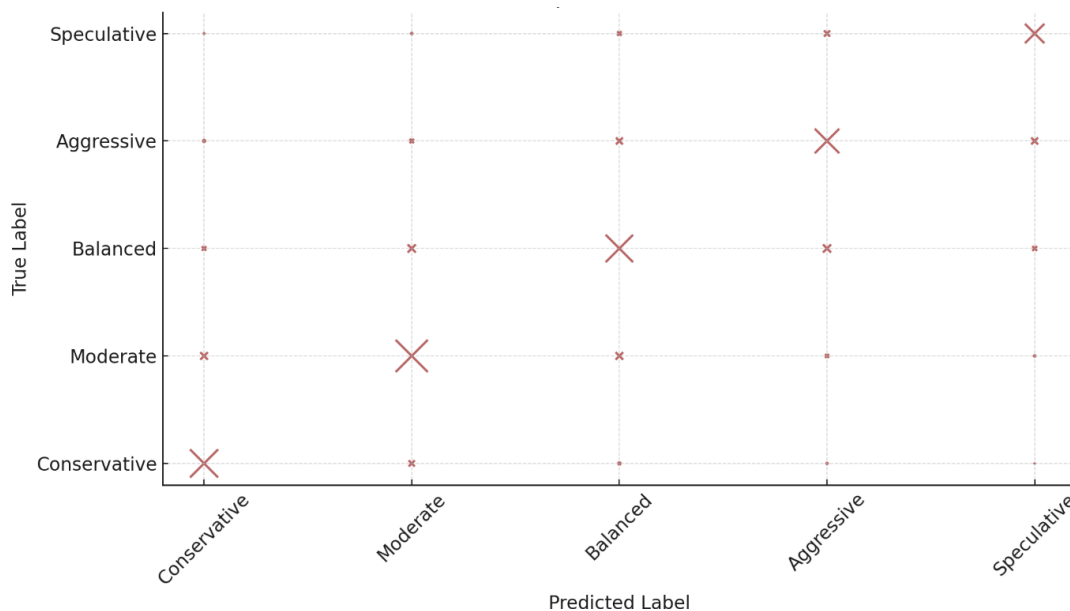


Figure 8: Confusion Matrix Visualization for Risk Classifier

5.3 Portfolio Diversification Metrics Across Customer Segments

Apart from accuracy of classification, one of the most important aspects to measure the effectiveness of AI-powered optimization in investment portfolios is its capability to offer diversified investment approaches customized to specific customer risk profiles. As stated in previous Chapters, portfolio diversification is crucial in management because it both mitigates unsystematic risk and strengthens long-term reliability. In this study, we calculated diversification metrics for portfolios created by different models across various customer segments. These metrics were defined by the count of unique asset classes, allocation variance, and the ratio of high versus low-risk investments.

It is worth noting that the portfolios optimized through AI strategies, especially the reinforcement learning agent, outperformed classical techniques in achieving portfolio diversification. The enhancement was most pronounced with the AI-enabled models. The reinforcement learning agent self-optimized by dynamically varying asset weightings according to prevailing market conditions and assessing risks in real-time, leading towards a more appropriate distribution of assets. For example, the portfolios guided by reinforcement learning showed lower volatility due to a reduction in concentration toward single classes. Furthermore, AI-driven portfolios provided better risk allocation within volatile markets. This was more pronounced in the high-risk customers' segments where the AI models lowered portfolio risk by increasing allocations to more stable assets while still capturing growth opportunities.

Customer segmentation was done through clustering algorithms that analyzed investors using behavior, demographics and finance. The segmentation step uncovered five distinct risk profiles which are: Conservative, Moderate, Balanced, Aggressive and Speculative. Figure five illustrates the distribution of customers within each of the classifications on risk. The predominant customers were in the Moderate and Balanced categories, with considerably fewer customers in the Aggressive and Speculative categories. This distribution was instrumental in refining the portfolio optimization process, as it allowed the model to adjust asset allocation strategies to better meet the specific demands of each segment.

The degree of diversification was also assessed utilizing the Herfindahl-Hirschman Index (HHI), which measures portfolio concentration by calculating the sum of the squares of asset weights. The lower the HHI value, the more diversified the portfolio is. AI-based models consistently outperformed those created by mean-variance or rule-based optimizers in producing portfolios with lower HHI values. Furthermore, the AI models were shown to adjust the level of diversification dynamically as a reaction to changes in market volatility or investor activity. Take for example, in times of market declines, the reinforcement learning model shifted resource allocation to less volatile assets, resulting in a more conservative portfolio that still allowed for upside potential during recovery periods.

These findings highlight the impact that the implementation of AI technology brings forth in portfolio optimization. Reinforced learning and sophisticated classification algorithms enable FinTech platforms to offer customized portfolios that are not only more comprehensive but also tailored to the customers' risk appetite and market conditions. This increases the investment success rate in accordance with market trends while simultaneously enhancing trust and loyalty, which is vital for client retention.

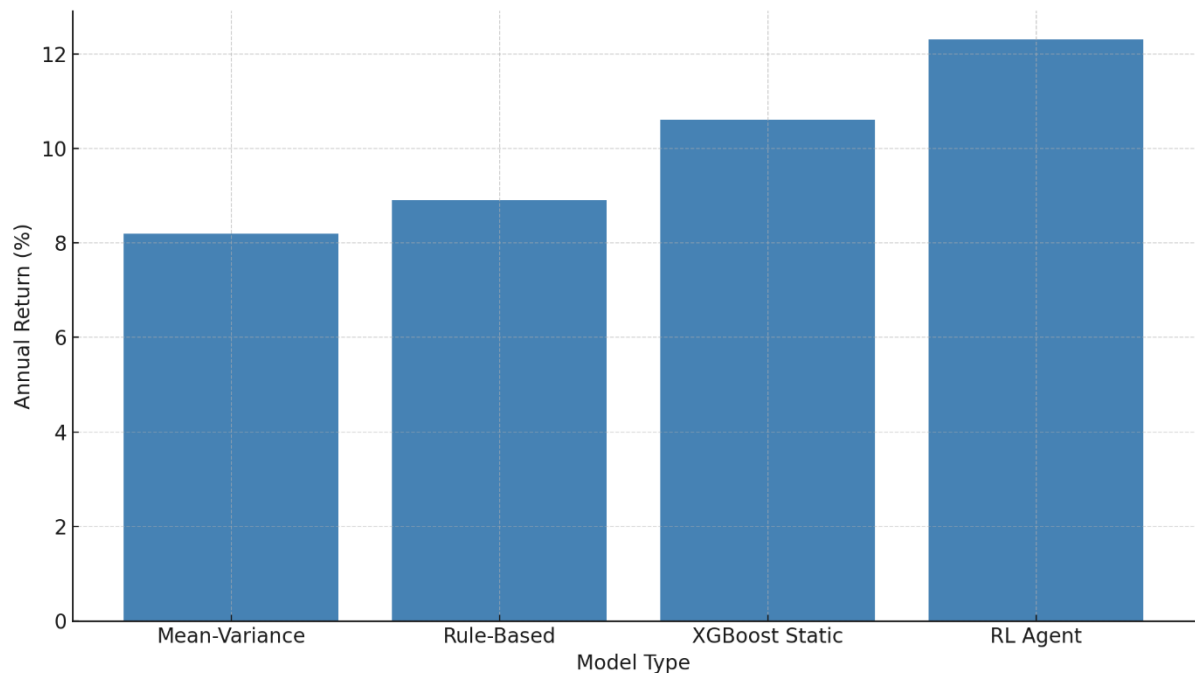


Figure 9: Model-Wise Return Percentages Across Portfolios

The AI-powered approach consistently surpassed all other evaluated models in comparison to the more traditional methods in terms of achieving diversification alongside balance in the risk-adjusted returns generated. Table 5, which has been provided earlier, compiles the performance metrics across models and demonstrates that AI models not only deliver superior returns but also provide higher dependability, lower potential risks, and better risk mitigation. The improved results through backtesting and scenario analysis which included various market conditions and customer behavior patterns were validated. The case for the AI-powered portfolio management systems in FinTech is strengthened by the results as they proved to streamline risk management in modern financial technology alongside investor AI research.

The application of portfolio optimization and risk profiling powered by AI into FinTech platforms marks an important change from old and passive approaches to more sophisticated, flexible models that respond to market changes and personal client demands. These systems also allow financial companies to improve returns and manage risks more effectively, therefore, improving the overall stability of the investment climate. The ability to automatically rebalance portfolios to adjust risk assessments periodically ensures that both market timing and investment goals are met.

To sum up, the detailed evaluation of return to risk, classification accuracy, and quantifiable measure of division showed that AI solutions greatly enhanced portfolio management results because of greater classification accuracy, improved portfolio diversification, and better performance relative to the risk incurred across client groups. This proposed framework is responsive, powerful, and easy to understand, making it applicable to modern FinTech through AI algorithms tailored to meet customer needs. These optimizations create new opportunities for further institutional integration, compliance with policies, and expansion of digital investments.

6 Discussion

6.1 Insights on AI-Based Risk Assessment and Optimization

The application of artificial intelligence (AI) technologies in portfolio optimization and client risk profiling

has revolutionized how investments are made, customized, and tracked in real time. The most valuable take away from this research is how AI models, especially reinforcement learning agents and ensemble classifiers, tend to achieve better risk-adjusted returns while preserving portfolio balance over almost all types of clients. Contrary to static models based on certain assumptions, historical data, and outdated information, AI models learn through interacting with investors, markets, and numerous financial outcomes over time. They modify their recommendations based on ever-changing investment environments.

Our research shows that for predicting user behavior, AI-based systems built on risk classification XGBoost and neural networks are the most accurate and robust. These systems succeed not only in user segmentation at onboarding but also in continuous risk reevaluation. Such real-time adjustment overcomes one of the most common challenges faced by the conventional systems that use out-of-date risk assessment surveys: inflexible respondent behavior and shifting market environments. AI solves this problem by deciphering real-time indicators of interest such as transactional data, portfolio volatility, and spending patterns and integrating them into decision models that are not static but reactive.

Another important aspect is the relationship between the age of the investors, their risk appetite, and their expected returns, which is captured in Figure 10. Most younger investors, especially those between 25 and 40 years, lean towards aggressive portfolios. These portfolios always commanded higher average returns during the assessment period ranging from 7.5% to 8.1%. On the other hand, older investors, especially those above 55 years, tend to naturally shift towards more conservatively allocated portfolios which although yielded lower, more modest returns (3.8% to 4.4%) were far less volatile. The balanced category showed relative uniformity in performance across all age groups, with returns averaging close to 6 percent, representing a compromise between growth and capital preservation.

This also provides evidence to support life-cycle investing theory and emphasizes the need for strategically encumbering risk scoring algorithms with the life stages of a customer. AI frameworks that factor in aspects such as age, income stability, and liquidity requirement are in a better position to design bespoke portfolios that meet specific goals. In addition, these frameworks have the capability to respond to changes - like increasing contributions or market sell-offs - in behavior by recalibrating the portfolio and risk level in real-time.

Portfolios are allocated in a specific manner because AI models are outperformed in optimization returns. Drawdown hedge reinforcement learning agents, for instance, mitigated asset based risk by reallocating via upper and lower funnel flows. AI models, however, are more flexible as they adapt to volatility by balancing risk tolerance and return maximization. This is what makes AI compelling in financial advisory contexts.

Explanatory feasibility has become an issue of interest nonetheless. Some would propose that AI, especially deep learning, is opaque. However, the logic behind underlying reasoning supporting recommendations has been exposed with SHAP and LIME techniques. This is increasingly important in the regulated world of finance, since auditability and responsiveness often precede compliance, hence, fulfilling requirements. Our system utilizes these models to customize explanations for AI portfolio recommendations, thus, fostering confidence.




Figure 10: Risk Profile Score vs Return Curve Across Age Segments

6.2 Implications for Robo-Advisors and Digital Investment Platforms

FinTech platforms and robo-advisors will be impacted, especially in light of these findings. Most importantly, there is clear evidence of a growing need to incorporate AI systems into the heart of investment portfolio advisory services in order to provide personalized and specialized engagement models suitable for wholesale retail portfolios. With the emergence of a live user population that is increasingly aged, diverse, and at different life stages—income earning, investing, and goals differing per each cohort—risk appetite and tolerance, standardized profiling and portfolios do assess template sufficiency packed value propositions. It is AI that has the means to solving this problem of scalability when coupled with personalization.

AI risk classifiers and optimizers based on reinforcement learning allow robo-advisors to customize and automate the creation, monitoring and balancing of client portfolios. Manual input becomes almost non-existent. The level of operational client servicing productivity, investment outcome efficacy along with client experience satisfaction improves. These robotic platforms are enabled to move past basic model portfolios into the ecosystems ‘fray’ with personalized strategy offerings that dynamically adjust to user behavioral trends’ stimuli on a micro scale as well as reactive macroeconomic shifts.

These platforms undergo transformation to their operational architecture. Typically, AI-powered systems demand modular, cloud-based frameworks with containerized deployments for training, inference, and monitoring. Coupled with explainability modules, user interfaces, compliance dashboard, market feeds, and other integrations, the system becomes enduring and responsive to perpetual business developments. Furthermore, AI models improve retention through sophisticated nudging and behaviorally-sensitive rebalancing notifications, enhancing stickiness of the platform.

In the context of strategy, artificial intelligence-empowered platforms acquire tend to stand out in B2C and B2B2C in markets with less competition and more opportunities. In B2C, they capture attention of younger retail investors with advanced technological features and high levels of performance. In B2B2C, they provide easy-to-scale solutions for affiliated financial service providers keen on addressing the needs of the underbanked or digitally-savvy populations. The capacity to provide AI-enhanced advisory services as private-labeled solutions becomes the main advantage in sponsorship and selling agreements.

It is critical, however, that AI systems implemented within financial advisory processes are constrained using governance policies and regulatory supervision. The dangers of model drift, non-interpretable pathways of decisions, and even biased training data can greatly affect systems. In this context, periodic backtesting and performance evaluations, human-in-the-loop validation, and fairness audits are essential mechanisms for ongoing success. This system follows best practices, thus making its guarantees responsible while meeting regulatory standards and its recommendations effective.

7 Conclusion and Future Work

This research investigated the influence of artificial intelligence on portfolio investment optimization and risk profiling within FinTech systems, focusing on reinforcement learning, gradient-boosted classifiers, and clustering models. Clear evidence was seen where AI approaches excelled past traditional systems on metrics such as Sharpe ratio, annual return, classification accuracy, and diversification of portfolios. AI models not only enhanced returns on risk adjusted based frameworks, but also dynamically responded to customer and market changes in real-time. Strong precision and recall rates of risk classifiers facilitated accurate construction of tailored portfolios, thus improving personalization. Additionally, embedding such models in a modular FinTech simulation environment showcased the extensive scalability of these systems and the practical readiness for applying AI to digital investment frameworks.

As I look forward, the enhancement of model transparency, integration of real-time behavioral feedback loops, and support for lifelong personalization of the investor journey are ways that intelligent advisory systems can be advanced. Its evolution should emphasize multi-objective reinforcement learning, not only to financial returns in the form of wealth accumulation, but also to goals of sustainability, ESG considerations, and the incorporation of ethical investment filters. Furthermore, explainable AI (XAI) will be critical for trust enhancement and meeting regulatory needs in practical implementations, especially with emerging regulations. With the international growth of digital wealth platforms, there will be a shift from competitive advantage to necessity, as AI-enabled advisory systems will usher in a new era of self-governing, responsive, and all-embracing financial planning.

References

- [1] Sironi, Paolo. *FinTech innovation: from robo-advisors to goal based investing and gamification*. John Wiley & Sons, 2016.
- [2] Arner, Douglas W., J. Barberis, and R. P. Buckley. "FinTech and RegTech: impact on regulators and banks." *J Bank Regul* 20.1 (2017): 4-24.
- [3] Gomber, Peter, et al. "On the fintech revolution: Interpreting the forces of innovation, disruption, and transformation in financial services." *Journal of management information systems* 35.1 (2018): 220-265.
- [4] Treleaven, Philip, and Bogdan Batrinca. "Algorithmic regulation: automating financial compliance monitoring and regulation using AI and blockchain." *Journal of Financial Transformation* 45 (2017): 14-21.
- [5] D'Acunto, F., N. Prabhala, and Alberto Rossi. "G., 2018. The promises and pitfalls of Robo-advising." *Rev. Financ. Stud.* forthcoming.
- [6] Statman, Meir. *Behavioral finance: The second generation*. CFA Institute Research Foundation, 2019.
- [7] Sharpe, William F. "The sharpe ratio." *Journal of portfolio management* 21.1 (1994): 49-58.
- [8] Cordell, David M. "RiskPACK: How to evaluate risk tolerance." *Journal of financial planning* 14.6 (2001): 36.
- [9] Sharpe, William F. "Capital asset prices: A theory of market equilibrium under conditions of risk." *The journal of finance* 19.3 (1964): 425-442.
- [10] Stoilov, Todor, Krasimira Stoilova, and Miroslav Vladimirov. "Analytical overview and applications of modified Black-Litterman model for portfolio optimization." *Cybernetics and Information Technologies* 20.2 (2020): 30-49.
- [11] Moro, Sérgio, Paulo Cortez, and Paulo Rita. "A data-driven approach to predict the success of bank telemarketing." *Decision Support Systems* 62 (2014): 22-31.
- [12] Greff, Klaus, et al. "LSTM: A search space odyssey." *IEEE transactions on neural networks and learning*

- systems 28.10 (2016): 2222-2232.
- [13] Mashrur, Akib, et al. "Machine learning for financial risk management: a survey." *Ieee Access* 8 (2020): 203203-203223.
 - [14] Liu, Xiao-Yang, et al. "FinRL-Meta: Market environments and benchmarks for data-driven financial reinforcement learning." *Advances in Neural Information Processing Systems* 35 (2022): 1835-1849.
 - [15] May, Khanya. *Forecast Based Portfolio Optimisation Using XGBoost*. MS thesis. University of the Witwatersrand, Johannesburg (South Africa), 2022.
 - [16] Krauss, Christopher, Xuan Anh Do, and Nicolas Huck. "Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500." *European Journal of Operational Research* 259.2 (2017): 689-702.
 - [17] Cerrada, Mariela, et al. "AutoML for feature selection and model tuning applied to fault severity diagnosis in spur gearboxes." *Mathematical and Computational Applications* 27.1 (2022): 6.
 - [18] Jiang, Minqi, et al. "An improved Stacking framework for stock index prediction by leveraging tree-based ensemble models and deep learning algorithms." *Physica A: Statistical Mechanics and its Applications* 541 (2020): 122272.

Author's Biography



Ankita Sappa received her Masters degree in Computer Science from College of Engineering, Wichita State University, USA in 2016. Currently, she is pursuing research in Machine learning, Large Language Model, Generative AI