International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

RiskWise: An Explainable Behavioral Intelligence Framework for Detecting Risky Investor Behavior in Simulated Trading using Machine Learning

Hana Haseeb
Department of Computer Science
Atria Institute of Technology
Bengaluru, Karnataka, India

Prof. Asha P V
Department of Computer Science
Atria Institute of Technology
Bengaluru, Karnataka, India

Gulam Shabbir Khan
Department of Computer Science
Atria Institute of Technology
Bengaluru, Karnataka, India

J. Jennita Suganthi
Department of Computer Science
Atria Institute of Technology
Bengaluru, Karnataka, India

Anish Kumar
Department of Computer Science
Atria Institute of Technology
Bengaluru, Karnataka, India

***Abstract*—Individual retail investors often exhibit emotional trading patterns—fear of missing out, panic selling, revenge trading, and overconfidence—that systematically erode long-term performance. Despite the availability of sophisticated analytical tools on modern trading platforms, few systems offer direct feedback on the user's own behavioral tendencies. RiskWise is an explainable framework designed for paper trading that shifts analytical emphasis from asset forecasting to behavioral self-assessment. The system translates raw trading actions into structured behavioral events, derives a rich set of behavioral features, and employs an XGBoost classifier to label each action with a behavioral state. SHAP-based explanations render the model's decisions transparent, enabling personalized recommendations and behavioral risk scoring. Core elements include behavioral profiling, a behavioral DNA representation, simulation-driven training data, and a modular architecture that integrates behavioral finance, explainable AI, and machine learning. The platform is educational and non-investment-advisory, aiming to cultivate disciplined investing habits without compromising user autonomy. This document details the framework's design, feature engineering, learning methodology, and evaluation plan.**

***Index Terms*—Behavioral Finance; Explainable AI; XGBoost; SHAP; Paper Trading; Machine Learning; Investor Psychology; Risk Assessment; Behavioral Profiling; FOMO Detection; Panic Selling; Overconfidence**

I. INTRODUCTION

The proliferation of digital brokerage services has dramatically expanded retail investor access to global financial markets. With just a few clicks, individuals can now trade equities, options, and cryptocurrencies at near-zero marginal cost. Yet the empirical record consistently shows that the average retail investor underperforms relevant benchmarks, and a substantial portion of this underperformance can be traced to cognitive biases and emotional reactions rather than deficient market knowledge [1]–[3]. Overconfidence, loss aversion, herding, and the fear of missing out (FOMO) lead to predictable errors such as excessive turnover, chasing past returns, and panic liquidation during downturns.

Contemporary algorithmic trading assistants predominantly address market prediction. They apply deep neural networks, transformers, or reinforcement learning to forecast price movements, generate trading signals, or optimize portfolio allocations using news sentiment and technical indicators. While these models can achieve impressive predictive accuracy, they do not help investors recognize their own decision-making flaws. The underlying assumption—that superior forecasts automatically translate into better investor decisions—ignores the psychological mechanisms that often override rational analysis.

Behavioral finance research has established that emotional and cognitive factors routinely dominate investment choices, producing systematic departures from rationality [1]. These departures manifest as overreaction to recent news, reluctance to realize losses, and momentum chasing. They appear consistently across demographic groups and market regimes, implying deep-rooted psychological drivers. Hence, improving investment outcomes requires not only better market models but also a deeper understanding of the decision-maker's behavioral patterns.

RiskWise adopts a fundamentally different stance: rather than trying to predict where the market will move, it analyzes how the investor behaves. Operated within a simulated paper-trading environment, the system does not offer price forecasts, buy/sell recommendations, or real trade execution. Instead, it interprets trading activity to infer the user's behavioral state, computes a behavioral risk score, and delivers customized feedback. This approach converts raw transaction logs into structured behavioral events, extracts a comprehensive feature set, and applies a supervised learning model to classify

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

each trade as belonging to one of six behavioral categories. SHAP (SHapley Additive exPlanations) provides transparent explanations, showing which features contributed most to the classification. These explanations underpin the personalized advice delivered through the user interface.

The main contributions are: (i) a complete behavioral intelligence architecture for simulated trading with explainable outputs; (ii) a systematic behavioral feature engineering method that captures multiple facets of trading behavior; (iii) an XGBoost-based classifier with SHAP explanations for behavioral risk quantification; (iv) a simulation-based data generation pipeline that creates labeled training sets from persona-driven scenarios; and (v) a recommendation engine that translates classification outcomes into actionable educational interventions.

The remainder of this paper is organized as follows. Section II reviews relevant literature, covering behavioral finance theories and machine learning techniques for financial data. Section III articulates the specific gaps we aim to address. Section IV provides an overview of our proposed RiskWise framework. Section V details the underlying system architecture. Section VI unpacks the core behavioral intelligence engine. Section VII explains our feature engineering strategy. Section VIII presents the machine learning methodology. Section IX lays out the experimental evaluation plan. Section X discusses anticipated outcomes. Sections XI and XII acknowledge limitations and outline promising future directions, and Section XIII offers our concluding remarks.

II. Related Work

Before introducing our own approach, it is useful to situate it within the broader landscape of existing research. To that end, this section surveys the foundational contributions from behavioral economics, as well as the methodological advances in machine learning that have shaped our design decisions.

A. Behavioral Finance Foundations

A cornerstone of our work is prospect theory [1], which offered the first cohesive model for grasping how people make decisions when faced with risk. Its central insight is that individuals measure outcomes against a personal reference point and exhibit marked loss aversion—meaning the pain of a loss outweighs the pleasure of an equivalent gain. This fundamental asymmetry helps to explain puzzling market behavior, such as the disposition effect, where investors tend to cash in winners prematurely while stubbornly hanging onto losing positions. Subsequent studies have pinned down several other biases that frequently affect investors: overconfidence, which tends to increase trading volume at the expense of returns [2]; fear of missing out (FOMO), which fuels trend-chasing; and panic selling, which often surfaces in periods of high volatility. Critically, these tendencies appear consistently across diverse demographic profiles and market environments [3]. Taken together, the evidence makes a compelling case for tools that proactively tackle these behavioral shortcomings rather than merely accepting them as inevitable.

B. Machine Learning in Finance

Over the past few years, machine learning has firmly established itself as a go-to solution for financial forecasting. The toolkit here is broad, encompassing everything from classical regression techniques to advanced deep learning architectures capable of ingesting unstructured data like earnings calls, social media chatter, and high-frequency tick data. For instance, we have seen transformer models applied to sentiment analysis, large language models utilized for gauging credit risk, and specialized domain-adapted networks trained to forecast equity returns. Nevertheless, despite their impressive track record in terms of raw predictive power, these systems are rarely equipped to offer behavioral feedback at the individual level. In practice, they churn out numerical predictions or directional signals, but fall short of providing users with a meaningful, interpretable critique of their own trading decisions.

C. Explainable AI

Explainability has gained traction in finance, especially for regulatory compliance and risk management. Techniques such as SHAP and LIME provide feature attribution for black-box models. However, their application to personalized behavioral coaching remains limited. Most current systems use explainability to justify model decisions to regulators or auditors, not to educate end-users about their own behavioral patterns. RiskWise bridges this gap by embedding SHAP into a closed-loop educational feedback system.

III. Research Gap

Despite the advances noted above, four major gaps persist. First, the predominant objective of machine learning in finance remains market prediction; sentiment analysis is almost always a means to that end. Second, explainability is typically deployed for audit or compliance, not for personal behavioral improvement. Third, individual-level behavioral profiling—tracking changes over time and adapting to each investor's unique baseline—is underdeveloped. Fourth, few systems attempt to operationalize behavioral finance theory through explicit feature design and theory-driven model construction. RiskWise is explicitly constructed to fill these voids by focusing on behavior, providing individualized explanations, and embedding a theory-informed feature pipeline.

IV. Proposed Framework

The RiskWise framework transforms trading actions into behavioral intelligence through a multi-stage pipeline that emphasizes modularity, explainability, and educational value.

A. Design Tenets

The proposed RiskWise platform has been designed around five core principles that guide both its development and operation. First, the platform is intentionally limited to behavioral analysis and does not attempt to predict stock prices, recommend investments, or execute financial transactions. This ensures that the system remains educational rather than functioning as an investment advisor. Second, every behavioral

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

prediction is accompanied by an explanation generated using SHAP (SHapley Additive Explanations). Instead of presenting users with unexplained outcomes, the platform clearly identifies the factors that influenced each prediction, making the decision-making process transparent and easier to understand.

Third, the primary objective of the platform is to improve users' awareness of their own trading behavior. Rather than encouraging aggressive trading strategies, RiskWise helps users recognize emotional patterns that may affect their financial decisions and promotes healthier trading habits. The fourth principle focuses on modularity. The system is organized into independent components, allowing additional data sources, machine learning models, or analytical features to be incorporated in the future without requiring major architectural changes. Finally, RiskWise is designed to support users rather than replace their judgment. The platform provides insights and recommendations based on observed behavior, while leaving all investment decisions entirely in the hands of the user.

B. Operational Concept

Raw trading events—orders, position adjustments, portfolio rebalancing—are continuously ingested and transformed into behavioral events. Each event is enriched with contextual information, then passed through a feature extraction pipeline that generates a high-dimensional vector covering trade characteristics, portfolio exposure, session dynamics, historical baseline, and market conditions. This vector feeds an XGBoost classifier that predicts the most likely behavioral state among FOMO, revenge trading, panic selling, overconfidence, compulsive trading, or rational trading. Simultaneously, SHAP values are computed to identify the most influential features. The user's behavioral profile is updated incrementally, and a behavioral risk score is derived from the classification probabilities and feature contributions. Finally, personalized recommendations are generated from a knowledge base of behavioral finance principles.

C. Behavioral Event Model

Every trading action gives rise to a behavioral event that records the action type (buy, sell, limit order, market order, stop-loss, rebalancing), timestamp, session ID, order size relative to portfolio, and contemporaneous market conditions. This event structure supports both micro-level (per-trade) analysis and macro-level (session or user) aggregation.

D. Behavioral Profile Architecture

The behavioral profile comprises three interrelated components: (i) a behavioral DNA that captures long-term stable traits (e.g., risk tolerance, loss aversion intensity, impulsivity) derived from the entire trading history; (ii) a dynamic layer tracking recent behavior over a rolling window; and (iii) a historical log of all classifications and risk scores. The profile is updated after each event, ensuring that the system adapts to gradual changes in the user's behavior.

E. Educational and Advisory Scope

RiskWise is explicitly non-advisory in the investment sense. It provides behavioral feedback, highlighting tendencies, suggesting cooling-off periods, recommending diversification, or proposing pre-commitment strategies, but never tells the user which stocks to buy or sell. This design maintains a clear separation from regulated financial advice and keeps the system safe for educational use.

V. System Architecture

The system follows a modular monolithic design that separates concerns while allowing independent scaling of machine learning components.

A. Frontend Layer

Built with React, the frontend delivers a dashboard that integrates portfolio visualization, trading controls, a behavioral timeline, and notification panels. Real-time updates are delivered via WebSockets, and all backend interactions occur through RESTful APIs.

B. Backend Layer

Spring Boot serves as the application core, providing authentication, paper-trading simulation, portfolio valuation, and behavioral event processing. A market-data integration service fetches and caches live prices, while a recommendation engine translates classification outputs into actionable messages.

C. Data Layer

MongoDB stores user profiles, portfolio states, behavioral events, and configuration settings. Redis caches frequently accessed data, including market prices and session state, to ensure low-latency operations.

D. Machine Learning Service

This standalone microservice handles feature extraction, XGBoost inference, SHAP value computation, and recommendation generation. The service loads a pre-trained model and explainer at startup and exposes a lightweight REST interface. Using a shared feature pipeline across training and deployment mitigates feature drift.

E. Modular Monolith

Although deployed as a single Spring Boot application, the codebase is organized into well-defined modules (user management, trading engine, portfolio manager, event processor, ML inference, notification). This structure preserves maintainability and testability while simplifying deployment.

F. End-to-End Flow

When a user submits a trade, the frontend sends it to the backend, which validates and executes it, updates the simulated portfolio, and forwards the action to the behavioral event service. The event service persists the event and triggers the ML service, which computes features, performs inference, and returns classifications and SHAP explanations. The results are stored in the user's profile, and the frontend is updated with the new behavioral insights.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

VI. Behavioral Intelligence Framework

This core analytical engine converts raw trading data into behavioral assessments through a clearly defined pipeline.

A. Taxonomy of Behavioral States

Each trade is assigned to one of six categories, defined as follows:

- **Fear of Missing Out (FOMO):** impulsive buying following price increases, momentum chasing, and entry without a clear rationale.
- **Revenge Trading:** escalating position sizes after losses, rapid attempts to recover losses, and departure from the user's typical strategy.
- **Panic Selling:** liquidating positions during downturns, selling at unfavorable prices, and overreacting to volatility.
- **Overconfidence:** abnormally high trading volume, outsized position sizes, and frequent rebalancing.
- **Compulsive Trading:** excessive trade frequency, very short holding periods, and patterns inconsistent with market conditions.
- **Rational Trading:** moderate frequency, position sizes aligned with risk tolerance, and evidence of systematic planning.

This categorization is grounded in behavioral finance literature and is designed to be empirically distinguishable from trading data.

B. Profile Components

The behavioral profile includes:

- **Behavioral DNA:** a compact set of stable traits (risk appetite, loss aversion, overconfidence propensity, emotional reactivity, impulsivity) derived from long-term trading patterns.
- **Behavioral Risk Score:** a continuous value from 0 to 100, where higher values indicate greater emotional influence. It is computed as a weighted sum of classification probabilities and feature-level SHAP contributions.
- **Behavioral Timeline:** a chronological record of all classifications and risk scores.
- **Intervention History:** a log of recommendations provided and their outcomes.

C. Knowledge Base

A curated knowledge base associates each behavioral state with its psychological underpinnings, common triggers, and evidence-based intervention strategies. This repository is continuously updated with new research findings.

D. Recommendation Generation

Recommendations are generated via a three-step procedure: (1) the current behavioral state is identified; (2) candidate interventions are retrieved from the knowledge base; and (3) they are personalized using the user's profile and history. Recommendations are framed constructively—e.g., "Consider a 5-minute cooling-off period after a loss" or "Review your portfolio concentration before adding more to a winning position."

VII. Behavioral Feature Engineering

Feature engineering is central to the system's predictive power and interpretability.

A. Feature Dimensions

Features are organized into six groups:

- **Trade Characteristics:** position size relative to portfolio, order type, limit price relative to current price, holding duration, realized profit/loss.
- **Portfolio Exposure:** concentration ratios, sector weights, number of positions.
- **Session Activity:** trade frequency, inter-trade intervals, session duration.
- **Behavioral History:** rolling averages of position size, holding period, and prior classifications.
- **Behavioral Deltas:** deviations from the user's own historical baselines.
- **Market Context:** volatility (VIX or realized), trend direction, sector performance.

B. Feature Vector Formulation

For an event e, the complete feature vector is:

$$\mathbf{x}_e = [\mathbf{x}_e^{\text{trade}}, \mathbf{x}_e^{\text{portfolio}}, \mathbf{x}_e^{\text{session}}, \mathbf{x}_e^{\text{history}}, \mathbf{x}_e^{\text{delta}}, \mathbf{x}_e^{\text{market}}]^\top \quad (1)$$

All features are normalized to have zero mean and unit variance to ensure consistent scaling.

C. Personalized Baselines

A key innovation is the use of personalized baselines for each user. For feature f, the baseline μ_f is updated via exponential smoothing:

$$\mu_f^{(t)} = \alpha\mu_f^{(t-1)} + (1-\alpha)f^{(t)} \quad (2)$$

where α is the smoothing factor. The deviation is then normalized by the feature's standard deviation:

$$\delta_f = \frac{f^{(t)} - \mu_f^{(t)}}{\sigma_f^{(t)}} \quad (3)$$

This personalization allows the model to flag deviations that are significant for that specific user, rather than applying global thresholds.

D. Feature Selection

Feature selection combines statistical tests (e.g., ANOVA) with domain expertise to retain only the most discriminative and stable features. SHAP importance scores from the trained model further inform feature refinement.

VIII. Machine Learning Methodology

The machine learning pipeline is designed for accurate classification with transparent explanations and real-time inference.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

A. Simulation-Based Data Generation

Because labeled behavioral data is scarce, we generate a large synthetic dataset using trader personas and scenario libraries. Personas define behavioral parameters (e.g., impulsivity, risk tolerance, reaction to gains/losses), and scenarios provide sequences of market events (bull runs, crashes, news shocks). The simulation produces millions of labeled events, covering both typical and extreme behaviors.

B. Shared Feature Pipeline

The same feature extraction and transformation code is used in training and production, ensuring consistency and minimizing feature drift.

C. XGBoost Classifier

XGBoost [4] was chosen for its excellent performance on tabular data, regularization to prevent overfitting, and fast inference. The objective is to minimize the regularized loss:

$$\mathcal{L}(\phi) = \sum_i \ell(y_i, \hat{y}_i) + \sum_k \Omega(f_k) \tag{4}$$

where ℓ is categorical cross-entropy and Ω penalizes tree complexity. Hyperparameters are tuned using randomized search with cross-validation, and early stopping prevents overtraining.

D. SHAP Explainability

SHAP values are computed for every prediction [5]. For a prediction $\hat{y}$ from feature vector $\mathbf{x}$, the contribution of feature i is:

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(|N| - |S| - 1)!}{|N|!} \left[\hat{y}(S \cup \{i\}) - \hat{y}(S)\right] \tag{5}$$

These values quantify each feature's marginal contribution, accounting for interactions. They are presented to the user as the top driving factors behind the behavioral label.

E. Behavioral Risk Score

The behavioral risk score R for a given event is:

$$R = \sum_{b \in \mathcal{B}} w_b\, p_b + \lambda \sum_{i \in \mathcal{I}} I_i\, \phi_i \tag{6}$$

where p_b is the predicted probability for state b, w_b is the severity weight, λ is a scaling factor, $\mathcal{I}$ is the set of risk-relevant features, I_i indicates elevation, and ϕ_i is the SHAP value. The score is updated with each new event and displayed to the user.

F. Evaluation and Validation

The model is evaluated using accuracy, precision, recall, F1-score, and ROC-AUC, with stratified cross-validation. Inference latency is measured to ensure sub-100 ms performance.

IX. Experimental Design

The evaluation plan encompasses classification performance, behavioral change, interpretability, and system efficiency.

A. Classification Metrics

We will report standard metrics (accuracy, precision, recall, F1, ROC-AUC) along with a confusion matrix to identify misclassification patterns.

B. Behavioral Outcomes

We will measure the change in the frequency of risky states and the behavioral risk score over repeated sessions. User engagement (dashboard visits, time spent on feedback, adherence to recommendations) will also be tracked.

C. Interpretability

We will assess SHAP explanation fidelity (how well they match model behavior), user comprehension (via quizzes), perceived usefulness (surveys), and trust (self-reported and behavioral).

D. User Study

A between-subjects study will be conducted with retail investors. The treatment group receives full behavioral feedback; the control group receives only portfolio information. Pre- and post-measures will evaluate awareness, risk scores, and trading behavior.

E. System Performance

We will measure inference latency, throughput under load, and resource utilization to validate scalability.

X. Expected Results and Discussion

We anticipate high classification accuracy for states with clear behavioral signatures (e.g., compulsive trading, panic selling). More subtle states like FOMO and overconfidence may show moderate accuracy but remain useful. The personalized baselines should improve sensitivity to individual deviations.

We expect that users receiving feedback will exhibit a decrease in risky-state frequency and lower behavioral risk scores over time. The SHAP explanations should enhance trust and comprehension, particularly when they highlight specific, actionable features.

Compared with market-prediction systems, RiskWise offers a unique educational value and inherent interpretability. Its simulation-based training may not fully capture real-world complexity, but it provides a solid foundation for initial deployment.

XI. Limitations

Several limitations must be acknowledged. The training data is synthetic and may not reflect all nuances of human behavior. XGBoost does not model temporal dependencies, though this can be addressed in future work. The system only observes behavior, not internal psychological states. Moreover, the educational impact depends on user engagement; passive users will derive limited benefit. Finally, the simulated environment may not replicate the emotional intensity of real trading.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 08 , August - 2026

XII. Future Work

Future directions include integrating biometric signals (e.g., heart rate) and user-interface interactions to enrich behavioral features. Sequential models such as LSTMs or Transformers could capture temporal dynamics. Longitudinal field studies will evaluate long-term behavioral change. Integration with live trading platforms (as an overlay, not an execution system) would extend the system's practical utility. Social features, such as peer comparison and collaborative coaching, could enhance engagement. Finally, more advanced natural-language explanations could improve user comprehension.

XIII. Conclusion

RiskWise provides a novel approach to investor education by focusing on behavioral self-awareness rather than market prediction. The framework integrates behavioral finance principles, comprehensive feature engineering, XGBoost classification, and SHAP explainability within a modular architecture designed for simulated trading. By delivering transparent, personalized feedback, it empowers users to recognize and mitigate harmful trading patterns. The system is educational, non-advisory, and respects user autonomy. Future development will extend its data sources, modeling capabilities, and real-world applicability, with the ultimate goal of fostering healthier investment habits.

Acknowledgment

We are deeply grateful to everyone who helped shape the RiskWise framework into its current form. The software engineering team deserves special recognition for their hard work in building the modular architecture and making sure all the machine learning components worked smoothly within the trading simulation environment—a task that was often easier said than done. Our sincere thanks also go to the behavioral finance experts who generously shared their knowledge; their insights were instrumental in developing the behavioral taxonomy and feature engineering approach, grounding the system in solid theory.

We would also like to thank the beta testers and user study participants whose thoughtful feedback made a real difference. Their engagement and constructive critiques helped refine the educational recommendations and significantly improved the overall user experience, making the system more practical and user-friendly.

A big thank you to our colleagues and peers for the many stimulating conversations, critical feedback, and encouragement along the way. The collaborative spirit within our research group was a constant source of motivation and played a crucial part in seeing this work through to completion.

Finally, we are grateful to the anonymous reviewers for taking the time to carefully read our manuscript. Their insightful suggestions helped us clarify our ideas and strengthen the paper considerably.

References

[1] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica*, vol. 47, no. 2, pp. 263–291, 1979.
[2] B. M. Barber and T. Odean, "Boys will be boys: Gender, overconfidence, and common stock investment," *The Quarterly Journal of Economics*, vol. 116, no. 1, pp. 261–292, 2001.
[3] T. Odean, "Volume, volatility, price, and profit when all traders are above average," *The Journal of Finance*, vol. 53, no. 6, pp. 1887–1934, 1998.
[4] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 785–794.
[5] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proc. 31st Int. Conf. Neural Information Processing Systems*, 2017, pp. 4765–4774.