International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 • Website: www.ijfmr.com • Email: editor@ijfmr.com

From Rule-Based Trading to Adaptive Algorithms: A Review of AI in Financial Markets

Reyansh Mangal

Student, Shiv Nadar School, Noida

Abstract

Traditional financial frameworks, such as Modern Portfolio Theory (MPT), often fail to account for real-world frictions including transaction costs, information asymmetries, and gradual portfolio adjustment. This paper examines transition from these classical, rule-based systems to Artificial Intelligence (AI)-based trading models designed to navigate the complexities of modern, data-intensive markets. Current literature indicates that AI applications in trading can be categorized into three primary paradigms: Machine Learning (ML) for non-linear signal generation, Deep Learning (DL) for hierarchical feature extraction from high-dimensional and alternative data (such as sentiment analysis), and Reinforcement Learning (RL) for sequential decision-making that explicitly incorporates market constraints.

A notable recent development is the rise of hybrid architectures that modularize the trading process – separating signal prediction from execution optimization to balance predictive power with operational robustness. Central to this transition is a shift in training methodologies, moving from static historical modeling towards dynamic walk-forward validation and simulated market environments. Despite technical advancements, empirical evidence suggests that AI does not offer a "silver bullet" for universal outperformance; rather, its efficacy is found in improved adaptability and execution efficiency under realistic frictions.

However, the increased use of AI-based trading must account for critical ethical and systemic risks, including model opacity, algorithmic herding, and the concentration of technological advantages in the hands of a small pool of investors. While AI significantly enhances market responsiveness, its successful integration requires a synthesis of computational innovation with rigorous economic reasoning and transparent regulatory governance to ensure long-term market stability.

Introduction

Financial markets have traditionally relied on human judgement, economic theory, and rule-based decision making to guide trading and portfolio allocation. Classical financial frameworks, such as the Modern Portfolio Theory (MPT) assume rational investors, frictionless markets, and stable statistical relationships between risk and return (Markowitz). Under these assumptions, investors are expected to continuously rebalance portfolios to achieve optimal diversification and risk-adjusted returns.

However, decades of empirical evidence suggest that real-world financial markets deviate significantly from these theoretical ideals. Traditional trading is typically characterized by periodic rather than continuous portfolio rebalancing, reliance on historical data and discretionary judgement, and significant exposure to transaction costs and information frictions. Bacchetta demonstrates that international mutual

fund portfolios adjust only gradually to changes in expected returns, with actual adjustment speeds reaching just 30 to 50 percent of what frictionless theory predicts. This inertia is primarily driven by trading costs, taxes, regulatory constraints, and difficulties in accessing and processing reliable information across markets, particularly in international contexts (Bacchetta et al.). These frictions contribute to persistent phenomena such as home bias, under-diversification, delayed price discovery, and suboptimal portfolio outcomes.
As global financial markets have become increasingly complex, data-intensive, and interconnected, the limitations of human-centered and static rule-based trading systems have become increasingly evident. The growth of high-frequency data, alternative data sources (news and social-media sentiment), and rapid market regime shifts has outpaced the ability of traditional models to adapt effectively (Cohen). Early algorithmic trading systems improved execution efficiency by automating predefined rules, but their rigidity prevented them from responding dynamically to non-linear and evolving market conditions (Cohen).
Recent advances in artificial intelligence (AI) and machine learning offer a fundamentally different approach to trading and portfolio management. AI-based trading systems are designed to address key inefficiencies of trading by processing large volumes of structured and unstructured data, identifying complex non-linear patterns, and updating decision rules dynamically as new information arrives (Warin and Stojkov). Deep learning models improve pattern recognition in high-dimensional financial-data, while reinforcement learning frameworks enable sequential decision making that explicitly account for transaction costs, risk, and evolving market conditions (Mienye et al.; Bai et al.). By shifting from static optimization to adaptive learning, AI-based trading aims to reduce inertia, overcome information frictions, and improve responsiveness in modern financial markets.

This paper examines how AI-based trading models attempt to overcome the inefficiencies of traditional trading, the types of AI algorithms currently employed in financial markets, and the extent to which these systems improve trading efficacy across markets. It also considers the ethical and systemic risks associated with the increasing reliance on autonomous algorithmic decision marking in modern finance.

AI Models for Trading

The limitations of traditional trading systems have driven adoption of artificial intelligence (AI) in financial markets. AI-based trading systems learn directly from data and adapt their behavior as market conditions evolve. The literature identifies three broad categories of AI models used in trading: machine learning models, deep learning architectures, and reinforcement learning frameworks, each addressing different aspects of the trading process (Cohen; Warin and Stojkov).

Machine Learning

The introduction to machine learning marked the transition from rule-based systems to adaptive, data-driven models. Early applications focused on supervised learning algorithms such as Support Vector Machines (SVMs), Decision Trees, Random Forests, and Gradient Boosting models. These techniques were applied to financial forecasting tasks including price direction prediction, volatility estimation, and credit risk classification (Warin and Stojkov).
These models represented a significant evolution by learning non-linear patterns directly from historical data, avoiding the strict distributional assumptions of traditional econometric models (Mienye et al.). In

practice, such models are widely used by quantitative hedge funds and asset managers as signal generation tools, particularly in equity and foreign exchange markets (Guo et al.; Cohen).

These models are typically trained using supervised learning on historical financial data. While effective, these models remain limited by their reliance on manually selected features and their difficulty in handling very high-dimensional or unstructured data (Warin and Stojkov).

Deep Learning and the Rise of High-Dimensional Financial Data

As financial markets became increasingly data-intensive, deep learning models emerged as the next stage in the evolution of trading systems. Deep neural networks differ from traditional machine learning models in their ability to automatically extract hierarchical features from raw data, reducing reliance on human-designed inputs (Mienye et al.).

Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) architectures, are widely used for modelling sequential financial data. These models are applied in equity, commodity, and cryptocurrency markets to forecast prices and volatility by capturing long-term temporal dependencies that simpler time-series models cannot represent effectively (Mienye et al.).

Convolutional Neural Networks (CNNs) have also been adapted for financial applications. By treating financial indicators or asset correlations as structured grids, CNNs identify spatial patterns across assets and time windows. These models are commonly employed in systematic trading strategies that analyze cross-asset relationships rather than single securities (Warin and Stojkov).

A major driver of deep learning adoption has been the growth of alternative data, including news articles, earnings-call transcripts, and social media sentiment. Natural language processing (NLP) models, increasingly based on transformer architectures, enable trading systems to process unstructured textual data at scale. Asset managers and proprietary trading firms use these models to react to information flows that are impractical for human traders to process in real time (Warin and Stojkov; Mienye et al.).

Reinforcement Learning and Decision-Based Trading

Beyond predictive modeling, reinforcement learning (RL) represents a fundamental shift in how trading strategies are developed. Rather than forecasting future prices, RL agents learn decision policies through repeated interaction with a market environment. At each step, the agent observes the current market state and portfolio position, selects an action such as buying, selling, holding, or rebalancing, and receives a reward based on portfolio performance, typically adjusted for transaction costs and risk exposure (Bai et al.).

This formulation allows trading to be modeled as a sequential decision-making problem, where the objective is to maximize cumulative long-term returns rather than short-term prediction accuracy. Deep reinforcement learning (DRL) algorithms extend this framework by combining RL with deep neural networks, enabling agents to operate in high-dimensional and continuous state spaces. Commonly used DRL methods in finance include Deep Q-Networks (DQN), Proximal Policy Optimization (PPO), and Soft Actor-Critic (SAC), which have been applied to portfolio allocation, dynamic rebalancing, and order execution tasks (Bai et al.; Liu et al.).

A key advantage of reinforcement learning is its ability to explicitly incorporate market frictions and constraints into the learning process. Unlike traditional optimization frameworks, RL agents can account for transaction costs, liquidity constraints, and risk penalties directly within the reward function, allowing strategies to adapt dynamically to changing market conditions. In practice, reinforcement learning systems are most commonly deployed in highly liquid markets such as equities, futures, and cryptocurrencies, where frequent interactions generate sufficient data for effective learning and strategy refinement (Liu et al.).

Hybrid AI Models in Real-World Trading Applications

In practice, most large financial institutions do not rely on a single artificial intelligence paradigm for trading decisions. Instead, hybrid models that combine machine learning, deep learning, and reinforcement learning are increasingly adopted to balance predictive accuracy, interpretability, and execution efficiency. These hybrid systems reflect the recognition that different stages of the trading process pose distinct informational and decision-making challenges (Cohen).

A common architecture separates prediction from execution. Supervised machine learning or deep learning models are used to generate return forecasts, volatility estimates, or market state classifications, while reinforcement learning agents are deployed downstream to optimize execution and portfolio rebalancing decisions. For example, deep neural networks may process price data and alternative data to produce signals, which are then fed into reinforcement learning agents that determine trade timing, order size, and execution strategy under transaction cost and liquidity constraints (Bai et al.; Liu et al.).

Asset managers and quantitative hedge funds frequently combine interpretable machine learning models with deep learning components to mitigate opacity risks. Tree-based models such as Random Forests or Gradient Boosting are often used alongside neural networks to validate signals and improve robustness across regimes, particularly in equity and foreign exchange markets (Warin and Stojkov; Guo et al.). This ensemble-style approach reduces reliance on any single model class and limits performance degradation during structural market shifts.

Hybridization is also evident in the use of alternative data. Natural language processing models based on transformer architectures extract sentiment or event signals from news and earnings transcripts, while traditional machine learning models integrate these outputs with numerical financial indicators. This layered structure allows firms to exploit unstructured data while maintaining control over portfolio-level risk exposures (Warin and Stojkov; Mienye et al.).

Overall, hybrid AI systems represent a pragmatic response to real-world constraints. Rather than pursuing fully autonomous trading, firms deploy modular AI architectures that combine prediction, risk management, and execution optimization, prioritizing robustness and cost-adjusted performance over theoretical optimality (Cohen).

How AI Trading Models are Trained

The training methodologies of AI-based trading models have evolved alongside increasing model complexity and data availability, reflecting persistent challenges posed by noisy and non-stationary financial markets. Across the literature, a central concern is the extraction of economically meaningful signals while mitigating overfitting and ensuring robustness across changing market regimes (Warin and Stojkov; Mienye et al.).

Training of Traditional Machine Learning Models

Early machine learning applications in trading predominantly employ supervised learning frameworks trained on labeled historical data. Inputs typically consist of engineered financial features, while targets include future returns, price direction, or volatility measures (Warin and Stojkov). During training, models minimize prediction error through loss functions that treat historical outcomes as imperfect proxies for future market behavior.

Feature engineering plays a central role in this process. Common inputs include technical indicators, macroeconomic variables, firm-level fundamentals, and cross-asset relationships (Warin and Stojkov; Cohen). To respect the temporal structure of financial data and mitigate look-ahead bias, training is commonly conducted using rolling or expanding windows, where models are repeatedly retrained on recent observations and evaluated on subsequent periods (Cohen).

While these approaches represent a significant improvement over rule-based systems, their effectiveness remains constrained by static feature definitions and limited adaptability. Empirical studies consistently note that model performance deteriorates when market conditions shift or when learned relationships become unstable, highlighting the sensitivity of supervised learning models to regime changes (Warin and Stojkov; Mienye et al.).

Training of Deep Learning Models

Deep learning models introduce a more flexible training paradigm by learning hierarchical representations directly from raw or minimally processed data. Unlike traditional machine learning approaches, deep neural networks reduce reliance on manually engineered features by automatically extracting complex patterns from high-dimensional inputs (Mienye et al.).

These models are typically trained using supervised objectives such as return prediction, volatility forecasting, or market state classification. Given the non-stationary nature of financial markets, studies emphasize rolling-window or walk-forward training schemes, in which models are periodically retrained to adapt to evolving data distributions (Cohen; Mienye et al.).

The expansion of alternative data sources, including news articles, earnings-call transcripts, and social media content, has further shaped training practices. Textual data are transformed into numerical representations and used to train models that associate information flows with subsequent market movements, enabling large-scale processing of unstructured information (Warin and Stojkov).

Despite their expressive power, deep learning models are particularly vulnerable to overfitting due to the low signal-to-noise ratio in financial data. As a result, regularization techniques such as dropout, early stopping, and parameter constraints are widely employed to improve generalization and stability (Mienye et al.).

Training of Reinforcement Learning Models

Reinforcement learning (RL) models differ fundamentally from supervised approaches by learning through interaction with a simulated market environment rather than labeled outcomes. The environment specifies observable market states, available actions, and a reward function that reflects portfolio performance (Bai et al.; Mienye et al.).

Rewards are commonly adjusted for transaction costs, risk exposure, and other market frictions to ensure realistic learning dynamics. Through repeated interactions, the agent updates its policy to maximize cumulative long-term rewards rather than short-term predictive accuracy. Deep reinforcement learning extends this framework by using neural networks to approximate value functions or policies in high-dimensional state spaces. Algorithms such as Deep Q-Networks (DQN), Proximal Policy Optimization

(PPO), and Soft Actor-Critic (SAC) have been applied to portfolio allocation, dynamic rebalancing, and execution problems (Bai et al.; Mienye et al.).

The literature consistently emphasizes that RL agents trained in frictionless environments fail to generalize to real markets. Incorporating transaction costs, liquidity constraints, and execution limits during training is therefore essential to avoid unrealistic strategies and misleading performance estimates (Warin and Stojkov; Liu et al.).

Common Training Challenges and Validation Practices

Across all AI model classes, training methodologies face shared challenges related to non-stationarity, noise, and data leakage. To mitigate these risks, studies emphasize strict temporal validation schemes such as walk-forward testing, in which models are trained on one period and evaluated on the next (Mienye et al.; Cohen).

In addition, transaction costs, slippage, and market impact are increasingly incorporated during training or evaluation to ensure that reported performance reflects deployable strategies rather than backtest artifacts. This shift reflects a broader movement in the literature toward realism and robustness in the training of AI-based trading systems (Bai et al.; Cohen).

Efficacy of AI-Based Trading Models

The efficacy of AI-based trading models has evolved alongside advances in computational power, data availability, and modeling techniques. Early algorithmic systems primarily sought marginal improvements over classical benchmarks through enhanced prediction accuracy or execution efficiency. Over time, evaluation criteria shifted toward adaptive decision-making and cost-adjusted performance, producing increasingly nuanced evidence on the advantages of AI relative to traditional approaches.

Early empirical studies assessed efficacy using predictive accuracy and excess returns relative to classical benchmarks such as buy-and-hold strategies, mean–variance portfolios, and factor-based models. Supervised learning methods, including Support Vector Machines, Random Forests, and Gradient Boosting, demonstrated modest improvements in return predictability and volatility forecasting, particularly in equity and foreign exchange markets (Warin and Stojkov). However, these gains were highly sensitive to feature selection and market regimes, and improvements in risk-adjusted performance were often small or statistically unstable once transaction costs were considered (Cohen)

The introduction of deep learning marked a second phase in efficacy evaluation. Deep neural networks improved the modeling of non-linear relationships and high-dimensional dependencies, particularly when alternative data sources were incorporated. Empirical evidence suggests that deep learning models can outperform traditional machine learning and econometric models in specific tasks, such as short-horizon return prediction and volatility estimation (Mienye et al.). Nevertheless, out-of-sample performance frequently deteriorates during regime shifts, indicating that superior pattern recognition does not guarantee temporal robustness (Warin and Stojkov; Cohen).

More recent research evaluates efficacy through decision-based metrics within reinforcement learning frameworks. Rather than focusing on point forecasts, RL systems are assessed on cumulative portfolio performance adjusted for transaction costs, risk constraints, and turnover. Studies applying deep reinforcement learning algorithms such as DQN and PPO report improved performance relative to static optimization benchmarks, including mean–variance portfolios and rule-based rebalancing strategies, particularly in high-liquidity environments (Liu et al.; Bai et al.)

Despite these advances, the literature does not support claims of persistent, universal outperformance by AI-based trading systems. Performance advantages are context-dependent and diminish as strategies become crowded or market conditions change. When realistic costs and frictions are incorporated, many AI strategies perform comparably to well-designed traditional factor models rather than dominating them outright (Cohen; Liu et al.). Consequently, efficacy is increasingly judged by improvements in adaptability, risk management, and execution efficiency rather than raw excess returns.

Ethical Concerns of Artificial Intelligence in Financial Trading

The growing deployment of artificial intelligence in financial trading raises ethical concerns that extend beyond firm-level performance to issues of market fairness, transparency, and systemic stability. AI-based trading systems operate at speeds and levels of complexity that often exceed effective human oversight, creating governance challenges that existing regulatory frameworks struggle to address (Cohen; OECD).

A central ethical concern is opacity and lack of explainability. Many AI trading systems, particularly those based on deep learning and reinforcement learning, function as black boxes whose internal decision processes are difficult to interpret. Unlike classical financial models with explicit assumptions and interpretable parameters, AI models learn latent representations that resist economic explanation, complicating regulatory supervision and post-event analysis during periods of market stress (European Securities and Markets Authority.).

Closely related is the issue of accountability. When AI-driven systems contribute to losses or market disruptions, responsibility becomes difficult to assign due to the distributed nature of decision-making across developers, traders, and automated execution mechanisms (Cohen). This diffusion of responsibility is ethically problematic when autonomous systems exert significant influence over market outcomes with limited direct human control (OECD).

AI-based trading also raises concerns about market fairness and unequal access. Advanced models require substantial capital, proprietary data, and computational infrastructure, concentrating technological advantages among large financial institutions (Guo et al.). This intensifies informational asymmetries and places smaller participants at a structural disadvantage, raising questions about competitive equity and market inclusiveness (Warin and Stojkov).

Another ethical risk involves market manipulation and unintended coordination. Reinforcement learning agents optimized solely for profit may discover strategies that exploit market microstructure in ways resembling manipulative behavior, even if manipulation is not explicitly programmed (Bai et al.). Additionally, the widespread use of similar data and objectives can lead to strategy convergence and herding, amplifying volatility and reducing market resilience (Kirilenko et al.).

Systemic risk represents a further concern. AI-driven trading systems can propagate shocks rapidly through feedback loops that unfold faster than human intervention is possible. The adaptive nature of these systems may exacerbate pro-cyclical behavior during stressed conditions, increasing the likelihood of cascading failures (European Securities and Markets Authority.; Cohen).

Finally, ethical questions arise from data usage and human agency. The use of alternative data raises concerns regarding consent, transparency, and bias, particularly when individuals are unaware that their digital behavior informs trading strategies (European Securities and Markets Authority.). At the same time, increased automation displaces human judgement, prioritizing algorithmic optimization over deliberation and accountability in financial decision-making (Cohen).

Overall, the ethical challenges posed by AI in trading are central to market integrity rather than peripheral design issues. The literature increasingly emphasizes the need for transparency standards, clearer accountability mechanisms, and regulatory frameworks that recognize the systemic implications of autonomous trading systems (OECD). Without such safeguards, efficiency gains risk being achieved at the expense of market stability and public trust.

Conclusion

Financial markets have long been shaped by theoretical models that assume rational behavior, frictionless trading, and stable relationships between risk and return. While these frameworks have provided foundational insights, empirical evidence consistently demonstrates that real-world trading is constrained by information frictions, transaction costs, and delayed portfolio adjustment. These limitations have created persistent inefficiencies that traditional rule-based and discretionary trading systems struggle to overcome.

Artificial intelligence has emerged as a response to these structural constraints rather than a replacement for financial theory. Machine learning, deep learning, and reinforcement learning models introduce adaptive mechanisms that allow trading systems to process large volumes of data, capture non-linear relationships, and update decision rules as market conditions evolve. Unlike classical optimization frameworks, AI-based systems are not restricted to static assumptions and can incorporate alternative data, temporal dependencies, and sequential decision-making directly into the trading process.

The literature reviewed in this paper suggests that the contribution of AI to trading efficacy is incremental rather than revolutionary. Early machine learning models improved predictive accuracy under specific conditions but remained sensitive to feature selection and regime shifts. Deep learning expanded the capacity to model high-dimensional and unstructured data, particularly through the integration of textual information, yet continued to face challenges related to overfitting and temporal instability. Reinforcement learning represents a conceptual shift by framing trading as a dynamic decision problem, enabling explicit treatment of transaction costs, liquidity constraints, and risk. However, its effectiveness depends critically on realistic training environments and careful validation.

Importantly, empirical evidence does not support claims of consistent or universal outperformance by AI-based trading systems (Cohen; Warin and Stojkov; Liu et al.). When evaluated under realistic assumptions that include transaction costs, market impact, and non-stationarity, AI strategies often perform comparably to well-designed traditional models. As a result, efficacy is increasingly assessed in terms of adaptability, robustness, and execution efficiency rather than raw excess returns (Liu et al.; Cohen; Bai et al.). This has led to the widespread adoption of hybrid architectures, where different AI paradigms are combined to address distinct stages of the trading process while mitigating model risk.

The growing reliance on AI in financial markets also raises ethical and systemic concerns that cannot be treated as secondary considerations. Opacity, accountability, unequal access to technology, and the potential amplification of market instability pose significant challenges to market integrity. The adaptive and autonomous nature of AI-based trading systems complicates regulatory oversight and blurs responsibility when failures occur. Without appropriate governance mechanisms, efficiency gains may come at the cost of increased systemic fragility and reduced public trust in financial markets.

Overall, AI-based trading should be understood as an evolutionary development rather than a disruptive break from existing financial practice. Its value lies in improving responsiveness to information, managing complexity, and optimizing decisions under realistic constraints. At the same time, its limitations underscore the continued relevance of economic reasoning, robust validation, and regulatory oversight. Future progress in AI-driven trading will depend not only on technical innovation but also on the ability to align adaptive algorithms with market stability, transparency, and accountability.

References

1. Bacchetta, Philippe, et al. "International Portfolio Choice with Frictions: Evidence from Mutual Funds." *The Review of Financial Studies*, vol. 36, no. 10, Oct. 2023, pp. 4233–70. *Silverchair*, https://doi.org/10.1093/rfs/hhad027.
2. Bai, Yahui, et al. "A Review of Reinforcement Learning in Financial Applications." arXiv:2411.12746, arXiv, 1 Nov. 2024. *arXiv.org*, https://doi.org/10.48550/arXiv.2411.12746.
3. Cohen, Gil. "Algorithmic Trading and Financial Forecasting Using Advanced Artificial Intelligence Methodologies." *Mathematics*, vol. 10, no. 18, Jan. 2022, p. 3302. *www.mdpi.com*, https://doi.org/10.3390/math10183302.
4. European Securities and Markets Authority. *Artificial Intelligence in EU Securities Markets.* Publications Office, 2023. *DOI.org (CSL JSON)*, https://data.europa.eu/doi/10.2856/851487.
5. Guo, Jian, et al. "Quant 4.0: Engineering Quantitative Investment with Automated, Explainable and Knowledge-Driven Artificial Intelligence." arXiv:2301.04020, arXiv, 13 Dec. 2022. *arXiv.org*, https://doi.org/10.48550/arXiv.2301.04020.
6. Kirilenko, Andrei A., et al. "The Flash Crash: High-Frequency Trading in an Electronic Market." SSRN Scholarly Paper no. 1686004, Social Science Research Network, 6 Jan. 2017. *papers.ssrn.com*, https://doi.org/10.2139/ssrn.1686004.
7. Liu, Xiao-Yang, et al. "FinRL: A Deep Reinforcement Learning Library for Automated Stock Trading in Quantitative Finance." arXiv:2011.09607, arXiv, 2 Mar. 2022. *arXiv.org*, https://doi.org/10.48550/arXiv.2011.09607.
8. Markowitz, Harry. "Portfolio Selection." *The Journal of Finance*, vol. 7, no. 1, 1952, pp. 77–91. *JSTOR*, https://doi.org/10.2307/2975974.
9. Mienye, Ebikella, et al. "Deep Learning in Finance: A Survey of Applications and Techniques." *AI*, vol. 5, no. 4, Dec. 2024, pp. 2066–91. *www.mdpi.com*, https://doi.org/10.3390/ai5040101.
10. OECD. *Artificial Intelligence, Machine Learning and Big Data in Finance: Opportunities, Challenges and Implications for Policy Makers*. OECD, 2021. *DOI.org (Crossref)*, https://doi.org/10.1787/98e761e7-en.
11. Warin, Thierry, and Aleksandar Stojkov. "(PDF) Machine Learning in Finance: A Metadata-Based Systematic Review of the Literature." *ResearchGate*, Oct. 2025. *www.researchgate.net*, https://doi.org/10.3390/jrfm14070302.