Agentic AI and Multimodal AI to Quantitative Finance using Algorithmic Trading for Wealth & Portfolio Management

by

BANKA TAGUR RANGA SATISH BABU,
BACHELOR OF TECHNOLOGY (COMPUTER SCIENCE ENGINEERING)

DISSERTATION

Presented to the Swiss School of Business and Management Geneva

In Partial Fulfillment of the Requirements

For the Degree

DOCTOR OF BUSINESS ADMINISTRATION

SWISS SCHOOL OF BUSINESS AND MANAGEMENT GENEVA

2025

Agentic AI and Multimodal AI to Quantitative Finance using Algorithmic Trading for Wealth & Portfolio Management

by

BANKA TAGUR RANGA SATISH BABU

Approved by:

Anna Provodnikova, PhD

Dissertation chair:

Received / Approved by

Renee Goldstein Osmic

Admissions Director

Acknowledgements

I am sincerely indebted to the valuable guidance provided by my supervisor Dr Hemant Palivela.

I am thankful to many authors of the articles, publications released by several credible research houses & leaderboards globally and on Wealth and Asset management area updates released by multiple market Giants into the WAM industry like Microsoft – Qlib, AI4Finance – FinRL, FinRobot etc., whose references have been made in my research paper, which has enhanced the usefulness and quality of research paper. I am thankful to authors for their contributions to my research paper.

Having been in the employment over 2 decades and over a decade in the Data and AI with globally reputed organizations like Ernst & Young, Accenture, eClerx, IBM etc. and on job learning to handle emerging challenges concerning Data and AI, I would not have been in position to write the comprehensive research paper. I express my heartfelt gratitude to the senior stakeholders of my current organization, whose contribution to my business learning on this current topic and context with all inspiring guidance and empowerment had significantly contributed to develop the quality and usefulness of the research paper.

Besides the learnings from the long working journey with multiple interactions with all relevant stakeholders ranging from different tools used of quantitative finance, algorithmic trading, open source challenges, everchanging AI / LLMs technology, multimodal incorporation, computational units (GPUs, CPUs, Cloud), prompt engineering challenges, Reasoning models, Agentic AI evolution, data manipulation and data modeling etc., whereby I could study and analyze multiple dimensions and contemporary data and information for the purpose of identifying critical issues and challenges as well as suggestions for the improvement arising out of past learnings.

ABSTRACT

This DBA research covers the progress in quantitative finance for WAM (wealth & asset management) and pinpoints the notable difference that Agentic AI and Multimodal AI integration may make.

The transition from traditional statistics to AI agents that can plan, search for solutions, learn from various types of inputs and exchange information with humans has an important impact in how to examine and deal with financial markets.

The framework has been devised using various sources, sees an AI system working like a partner in algorithmic trading and controlling personal wealth and portfolios for WAM professionals and firms. LLMs are used to coordinate functions, access data from a range of channels and adjust learning as part of reinforcement learning.

It is important to focus on RLHF for value alignment—systems with these methods to promise better outcomes, adjust to different financial markets and can address diverse personalized financial goals. The advantages can range from better risk-adjusted results, to discovering original investment ideas, better running the operation and offering more investment strategies to a wider audience.

The financial organizations should focus on being ethical and reliable as they seek to succeed in today's market. Proper regulations should be put in place to support new developments that benefit the market and investors.

Essentially, merging Agentic AI and Multimodal AI with quantitative finance transforms the way Wealth management firms work together in finance and machine capability in one of WAM's most critical areas. This study helps continue the discussion by looking at the current situation and designing an outline of future systems helps by setting out the course of actions.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I: INTRODUCTION

1.1 Introduction

Quantitative Finance as a field refers to the use of mathematics and statistics in the financial markets of decision making (Spadafora & Berman, n.d. 1; Cao et al., 2025, pp. 1-2 2). The infancy of its roots can be traced back to masters such as modern portfolio theory (MPT) that integrated optimization principles for both risk and reward (Markowitz, 1952 3; Boyd et al., 2024, p. 1 3), Mathematical modeling of asset price movements adapting concepts such as Brownian motion from other sciences (Spadafora & Berman n.d. 1). In the backdrop of rapid increase in complexity, scale, and interconnectedness of global financial markets in the second half of the 20th century and the 21st century, derivative pricing models of risk, and design of investments became formally rigorous quantitative formulations (Spadafora & Berman, n.d. 1; Baldacci, 2021 5). Traditional models, however, frequently make use of simplifying assumptions that fail to represent the entire dynamics of the real markets, especially with nonlinearity, sudden shifts in regimes, and high volatility (ResearchGate User, 2024 6); Scaletta, 2024).

The collision between market complexity and developments in terms of computation power has resulted in the emergence of algorithmic trading – a phenomenon implying the usage of computer programs that do trades under the set rules and instructions (Baldacci 2021 5; Azzutti, 2024 8). This computerized method is intended to achieve acquisition, accuracy, and uniformity of execution in addition to reducing human emotional biases (Azzutti, 2024 8). A popular subgroup of algorithmic trading is High-Frequency Trading (HFT), defined by extraordinarily quick order entry and cancellation and operations commonly on the microseconds scale making use of speed to exploit transient market inefficiencies (ResearchGate User, 2024, 9; MacKenzie, 2018; Son, 2022 10). Although the advocates believe high-frequency trading increases liquidity and price discovery (ResearchGate User 2024 11), doubts remain whether HFT has a hand in market volatility and systemic risk (Oxford University Press 2024 12; ResearchGate User, 2024 9).

Artificial Intelligence (AI), consisting of ML, DL and DRL, marks the next important step of evolution in quantitative finance and wealth management (Cao et al., 2025, pp. 1-3 2; Joshi 2025, p.117 13; Han et al., 2024, p. 33 14; Liu 2024). AI-based methods have the potential to go beyond simple, static rules and old-fashioned statistical models (which are restricted to linear relations between inputs and outputs) and to learn highly non-linear patterns directly from data sets of enormous scale (Cao et al., 2025, p. 3 2; International Journal of Science and Research Archive, 2024 8; Santos et al., 2023, p. 2 15). This entails sentiment analysis on text, extraction of intricate correlations, and evolution of adaptive strategies that adjust to dynamics in the market (Abdullah & Chowdhury 2023 16; Liu et al., 2024 17; Rane & Choudhary, 2023).

To be addressed in this research are two very advanced AI paradigms, which can drastically change the direction of the field: Agentic AI and Multimodal AI. Namely, agentic AI means systems which are developed in order to accomplish complicated long-horizon targets with limited human control, using sophisticated reasoning, planning, and autonomous action's ability (Li et al., 2024a 18; Acharya et al., 2025 18; Mukherjee et al., 2025 19). Such agents may effectively automate the entire set of workflows from the analysis of the data to the execution of the decision (Mukherjee et al., 2025 19; Kanerika, n.d. 20). Multimodal AI refers to systems that can process such information from different types of data simultaneously – numerical price data, textual news reports, social media feeds, and even visual facts such as charts – in order to come to a more complex and sophisticated understanding of where the market stands (Zhang et al., 2024 21; Bhatia et al. 2024 22; Cao et al. 2025, p. 5 2; Multimodal LLMs Survey 23).
The development from quantitative models that form the basis to rule-based algorithmic trading and, ultimately, towards AI-enabled systems denotes more than a growing rate or computing capacity. The appearance of New Agentic and Multimodal AI indicates a change of quality. Where previous stages were concerned with automating execution according to human defined models or rules, these advanced AI paradigms strive for richer understanding by codifying complex and diverse information sources and increased autonomy through transformation of that codification into action towards goals. This flow suggests AI systems with abilities that increasingly resemble, even exceed in specific ways, humans.

This thesis explores a thorough analysis of the integration of Agentic AI and Multimodal AI as parts of algorithmic trading systems and studying their use, performance, and implications for quantitative wealth and portfolio management. It examines how the advanced systems of AI, possibly— with the use of such techniques as Reinforcement Learning from Human Feedback (RLHF) for alignment and backed up by High-Performance Computing (HPC) infrastructure – can overcome obstacles faced by the existing methods and redesign the future of investment management.

1.2 Research Problem

Although the challenges have been overcome by notable advances, quantitative finance and algorithmic trading still experience challenges that are persistent and evolve with time. Conventional quantitative paradigms, including base frameworks such as Modern Portfolio

Theory (MPT), rely on the assumptions of normal distributions in returns and rational investors, which are difficult to find, where the markets show the existence of fat tails, volatility clustering, and behavioral biases (ResearchGate User, 2024 6; Scaletta, 2024; Kayumov, 2023 26). These models may fail to account for complicated, non-linear dynamics, change to accommodate abrupt market change in regime, and take care of risk in highly volatile periods (Scaletta, 2024; Ng et al., 2020 27).

Algorithmic trading that is fast and efficient for many traders arcades its own set of problems. Financial data is notoriously noisy, having a low signal-to-noise ratio, which makes distinguishing genuine trading signals out of random variations very hard (Yu et al., 2023 28.; Zeng et al 2024 29; Shi et al 2025 30). Models that have the ability to be flexible to the non-stationary nature of financial markets where trends and correlations evolve over time are necessary for the development of strong trading strategies (Lin& Zhou 2021 31; Sun et al., 2023 32). In addition, the danger of overfitting the model to the past data is always there and may result in strategies, which perform well in back testing settings and poorly in real market conditions (Cao et al., 2025, p. 3 2; Deng et al., 2024 33; Liu et al., 2024 34). The competitive landscape, particularly in HFT, drives a continuous technological "arms race," demanding significant investment in infrastructure and potentially contributing to market fragility, as evidenced by events like the 2010 Flash Crash (Oxford University Press, 2024 12; The Global Treasurer, 2025 35; ResearchGate User, 2024 9).

A very significant problem for modern quantitative finance is the effective integration of information from different sources. In fact, both structured numerical quantities (prices, volumes) and the much more powerful and quickly growing flow of unstructured and multimodal data such as news articles, financial reports, social media sentiment, analyst ratings, etc., and potentially even visual chart patterns interplay with financial markets (Zhang et al., 2024 21; Cao et al. 2025, p. 5 2; Huang et al. 2024). Current systems find it challenging to sift through, synthesize and weave this deep, multi-modal information environment into prompted and actionable trading decisions, which is a major gap in market analyzing prowess (Zhang et al., 2024 21).

The emergence of Agentic AI creates an additional layer of complexity. Creating autonomous agents that can do sophisticated reasoning, long-term planning and reliable execution in the high-stakes, dynamic world of finance is a mighty endeavor (Li et al., 2024a 18; Mukherjee et al., 2025 19). These agents should be able to interact efficiently with the environment, interact with humans, and make decisions under uncertainty (Li et al., 2024a 18). Critically, it is important to keep these autonomous systems on track with investor goals, risk preferences, and ethical standards (Bai et al., 2025 36). Research has already highlighted potential risks associated with agentic systems, including emergent deceptive behaviors like "alignment faking," attempts at "self-exfiltration" of model weights, the disabling of oversight mechanisms, and self-deception to satisfy goals improperly (Greenblatt et al., 2024 18; Xu et al., 2024 18, Chen et al., 2025 37). The possibility of multitudes of autonomous agents that interact in markets is also a concern for unpredictable emergent behaviors and systemic risks (The Global Treasurer, 2025 35; Chen et al., 2025 37).

Thus, the overall research problem that is discussed within the framework of this thesis is multifaceted. Conventional quantitative finance and algorithmic trading approaches become more restricted in dealing with the complexity on markets, utilizing multimodal data, and delivering robust and adaptive performance. Although there are potent solutions for the effective design, deployment and risk management and alignment by advanced AI such as Agentic and Multimodal AI, there are serious challenges in these designs, deployment and management.

This calls for a move out of optimization focused solely on prediction accuracy.

The challenge extends to the design, validation, and governance of complex — potentially autonomous — systems improving on AI in the ability to navigate the complexities of financial markets with reliability and responsible attitude. The shortcomings of the current models in dealing with non-stationarity, noise, and the cleft between prediction and action indicate the need for more complex approaches. Agentic AI works to fill the gap between analysis and autonomous action while multimodal AI tries to fill the data integration gap.

However, this quest to increased capability brings forth new and complex issues of agent reliability, alignment and explainability and the potential systemic effects of placing interacting agents in financial markets. That is to say that the research problem therefore shifts from ensuring that a given model is optimized to learning how to comprehend and control a complex, adaptive socio-technical system.

1.3 Purpose of Research

This Doctorate in Business Administration research is aimed at carrying out a comprehensive and detailed study of the design, use, performance assessment and wider implications of integrating Agentic AI and the Multimodal AI paradigms in algorithmic trading systems; especially (systemically designed) for the domain of quantitative wealth and portfolio management. This study attempts to narrow the gap between the theoretical prospects of such advanced AI technologies and their real application in complex financial settings as reliable and responsible.

To accomplish this grand goal, the research has specific objectives as discussed below:

1. **Agentic AI Frameworks for Autonomous Finance Analysis:** To systematically analyze the design, implementation and deployment of Agentic AI frameworks such as general-purpose platforms (e.g., LangGraph, CrewAI, AutoGen 13) and finance-specific systems (e.g., FinRobot 41, AlphaAgent 46, TradingGPT/FinMem 49), can be architected, implemented, and deployed to perform autonomous financial tasks. This means that it will involve exploring their use in fields like the alpha factor mining (Zhao et al., 2024 53; Xu et al., 2024 54; Yu et al., 2023 28; Ren et al., 2024 56; Tang et al., 2025 46), automated trading execution (Sun et al., 2022 57; Ezeilo & Nimo, 2021 58), and dynamic portfolio optimization (Sun et al., 2025 59; Sun et al., 2025 60; Zhao, 2024 61).
2. **Assess the Effects of Multi Modal Integration of AI:** To assess the use of Multimodal AI models (for example, FinAgent 21, FinLLaVA 22) as a means of improving quantitative finance, by drawing on various sources of data such as numerical market data, textual news & reports and social media sentiment, to name a few, and, arguably, visual chart data. The goal is to review enhancement in market analysis depth and accuracy in sentiment detection (Konstantinidis et al., 2024 63; Mean absolute deviation as a measure of forecasting accuracy (Abdullah & Chowdhury 2023 16; Zhu 2024 64), precision of forecasting (Zeng et al. 2024 65; Tian et al. 2025 66), and total decision quality.
3. **Explore RLHF for Agent Alignment:** To study the nature and efficacy of Reinforcement learning from human feedback (RLHF) as a means of aligning agentic financial systems to complex ends of human objectives. Such as alignment of agents with the risk preferences of the investors, ethical considerations, the regulatory constraints, and the long-term financial goals which might be hard to specify through the reward functions. Some of the problems the research will also look at include ensuring that the strategies adopted are strategy proof against the manipulative feedback (Bai et al, 2025 36).

   Characteristic workflow of the RLHF consists of two steps:
   a. **Training a Reward Model:** An AI agent produces a number of possible actions or outputs (e.g. two different portfolio rebalancing suggestions). This is followed by feedback given by a human expert who usually ranks the outputs or gives an indication of the one he or she prefers. This has been placed in a separate training pool of human preferences to train an independent "reward model" that learns to guess what outcomes a human will tend to prefer.

b. **Policy refinement of the Agent:** With the learned reward model, the reward signal is given in refining the main RL agent. This means the agent is not only optimized with respect to something as reductive as profit, but with the production of behavior that meets the previously learnt human preferences.

4. **Assess HPC Requirements and Impact:** To evaluate the extent and effect of High-Performance Computing (HPC) infrastructure such as GPU clusters and parallel processing capability for supporting a computationally intensive process of training and deploying complex Agentic and Multimodal AI models notably in the context of HFT and large-scale financial simulations (Penguin Solutions, n.d. 68; IBM, n.d. 69; Li et al., 2022 70).

Utilize High-Performance Computing (HPC) infrastructure

a. It is necessary to invest in High-Performance Computing (HPC) infrastructure to satisfy the computational requirements of sophisticated financial AI-driven applications.

b. This extends beyond normal computing facilities and includes installation of dedicated hardware and software stack designed to handle massively parallel calculations.

c. An example stack of an HPC stack in financial AI may contain the following:

• **GPU Clusters:** Graphics Processing Units (GPUs) would be critical due to their capacity to work with the enormous parallel matrix computations, which lie at the core of training deep neural networks and large language models. There is distributed training, e.g. a model can be jointly trained using a cluster of interconnected GPUs, which can result in a massive decrease in training time.

• **Low-Latency Networking:** To minimize latency times to nanoseconds, high speed, low latency interconnects (such as InfiniBand) between servers and to the exchange are essential in HFT and real-time inference as latency times can hold up data processing and execution of orders in microseconds.

• **Optimized Software:** This encompasses libraries that are particular to domains (e.g., NVIDIA CUDA to program GPUs), distributed computing engines (e.g., Horovod), and streamlined data processing engines whose functionality is custom designed to make effective use of underlying parallel hardware.

d. Through an investment in HPC, financial institutions will be able to deal with both the sheer amount of data that needs to be processed, and the sophisticated calculations demanded by training state-of-the-art AI and with supporting the deployment of AI into time-sensitive use cases, so that the amount of scale available within computing resources does not limit the innovation process.

5. **Risks and Governance Needs Identification and Analysis:** To find, analyze, and define the evolving risks of using highly autonomous, multimodal AI agents in financial markets. This involves assessing potential systemic risks, model weaknesses (e.g. hallucinations, adversarial attacks), the hurdles in AI governance (Azzutti, 2024 71). The research is to offer insights into the required regulatory changes and solid risk management frameworks (Chen et al., 2025 37; The Global Treasurer, 2025 35).

Through addressing such objectives, this research aims to develop rigorous and timely understanding of how Agentic and Multimodal AI can be used for quantitative wealth and portfolio management purposes while underlining the key points to consider for their safe and effective deployment.

1.4 Significance of the Study

In the academic, practical, technological and societal spheres, this research has great potential as regards the contributions. One of its main benefits lies also in the investigation of changes in the performance of investment strategies, rather than in the critical assessment of the fundamental transformation of the processes, risks, and theoretical basics of quantitative finance and wealth management which these highly advanced AI technologies bring forward on their own.

Academic Contribution: The current work primarily seeks to contribute significantly to the exploding scholarly literature examining the relationship between advanced AI (specifically, Agentic and Multimodal capabilities) and quantitative finance. It helps in filling in gaps arising from recent surveys on issues of practical implementation challenges, strict risk assessment models, and integration of these technologies with the existing financial theories (Cao et al., 2025 72; Ding et al., 2024 73; Liu, 2024; Dong et al., 2025; Chen et al., 2025 37). The research will offer new insights on the use of AI in connection with the likes of the Efficient Market Hypothesis (perhaps supporting adaptive or behavioral views), Modern Portfolio Theory (with data-driven alternatives or improvements), Behavioral Finance (to model or exploit biases), Bounded Rationality (in agent decision making)), and Agency Theory (in the context of autonomous financial agents).

Practical Implications for Industry: The results of this research are likely to contain very useful information for practitioners of the financial services industry, such as investment banks, hedge funds, asset management companies, and wealth advisors. It will shed light upon the avenues of using Agentic and Multimodal AI to get competitive advantage from superior analytics, automatized strategies, and personal client services (Han et al., 2024 14). Forbes, 2025 78). The study will provide the foundations and the critical questions for the design, implementation, validation, and sustained maintenance of advanced AI-driven trading and portfolio management systems. Additionally, it will be helpful in shaping industry approaches to risk management in the era of AI as well as aid to direct one's steps regarding the emerging regulatory environment regarding AI governance in finance (Azzutti, 2024 71). The research will also highlight the provision of utility for platforms such as open-source ones like FinRL and QLib in development, testing, and benchmarking.

Technological Advancement: This thesis expands the scope of broader Artificial Intelligence by extending the scope of limits of Agentic and Multimodal AI in the realm of a greatly complex, dynamic, and data-centric environment. The struggles faced and ideas tested in the realm of finances can be used in building more resilient, flexible, interpretable, and reliable AI systems that can be applied to other areas. The research points out the essential connection between state-of-the-art AI algorithms, the data infrastructure behind (the necessity of HPC), as well as the presence of domain-specific knowledge for practical applications out of the lab.

Societal Relevance: Going outside the financial industry, this study discusses broader societal issues. It touches upon the possible effects of highly autonomous AI trading system on the market stability and fairness (ResearchGate User, 2024 9; Oxford University Press, 2024 12). It addresses the important issues of AI decision-making transparency and explainability; of vital importance in public trust and accountability development (Gu et al., 2024 83; Ng et al., 2020 84). Besides, it raises the issue of the future of human roles in finance as AI agents continue to get more capable, thus transferring tasks from direct analysis and execution towards monitoring, policy setting, and client relations management (Forbes, 2025 78; Sharma et al., 2021 85).

In other words, the importance of this study is in the potential for rewriting the process of quantitative investment and wealth management. When shifting to more dynamic, data-rich and possibly autonomous paradigm of Agentic and Multimodal AI, traditional methods and theories require reconsideration. However, this transformation comes with critical questions on control, reliability, ethical implications, and systemic risks of these powerful technologies, making research that connects AI capabilities with responsible financial deployment of great importance.

1.5 Research Purpose and Questions

Reiterating the core purpose, the present research sets objectives in form of conducting thorough research into the design, application, performance, and implications of incorporating Agentic AI and Multimodal AI paradigms in algorithmic trading systems, particularly centered on quantitative wealth and portfolio management.

The present investigation is carried out following this comprehensive research question:

- *How can Agentic AI and Multimodal AI, incorporated with algorithmic trading frameworks and supported by techniques such as RLHF and HPC, be effectively and responsibly leveraged to enhance quantitative wealth and portfolio management strategies?*

In order to answer this wide issue, the research will concentrate on the following specific research questions that will guide subsequent chapters of this DBA thesis:

Research Questions (RQs):

- **RQ1:** How can Agentic AI principles (planning, tool use, reflection, autonomy) be combined with Multimodal AI abilities (ability to process numerical, textual, and possibly other types of data) to produce a consistent system for algorithmic trading and portfolio management? (Concentrates on building architecture and systemic integration).
- **RQ2:** What particular algorithmic trading strategies (such as, high-frequency trading informed strategies, factor-based investing, dynamic asset allocation) are most likely to benefit from improvement through an Agentic Multimodal AI framework, and – in what ways (numerically, in terms of risk-adjusted returns, adaptability to market changes, discovery of novel alpha) could realistically performance be improved based (Concentrates on the application scope and the possible impact)
- **RQ3:** How can Reinforcement Learning from Human Feedback (RLHF) be incorporated into the proposed framework to align the AI agent's trading and portfolio management decisions with complex, potentially subjective, human preferences (e.g., nuanced risk tolerance profiles, ethical investment considerations, specific long-term financial goals) that go beyond simple quantitative optimization targets like profit maximization? (Value alignment and continuous improvement via human interaction)

- **RQ4:** What are the main methodological issues, constraints, critical ethical issues (such as addressing data bias, ensuring interpretability of the model, avoiding the systemic risk potential, developing a proper governance structure) of the development and prospective implementation of the advanced Agentic Multimodal AIs in the sensitive field of quantitative finance? (Concentrates on the possibilities limits, underlying risks, and responsible innovation)

These questions cumulatively define the skyline and direction of this research as they intend to deliver a full insight into the transformative potential and crucial concerns of employing Agentic and Multimodal AI in current quantitative finance.

CHAPTER II: REVIEW OF LITERATURE

2.1 Theoretical Framework

Application of advanced Artificial Intelligence such as Agentic and Multimodal AI, in quantitative finance does not exist in isolation, but interacts significantly with existing Financial and economic theories. The mastery of these fundamental theories will be a critical device for analyzing the abilities, limitations, and implication of such new technologies. Central theoretical frameworks behind this research are reviewed in this section. Modern Portfolio Theory, Efficient Market Hypothesis, Behavioral Finance, Bounded Rationality, and Agency Theory.

Modern Portfolio Theory (MPT) pioneered by Harry Markowitz (Markowitz, 1952 3) is a mathematical approach to build up a portfolio of assets so that the return is maximized for a given level of risk (defined as variance or a standard deviation of returns). At the core of MPT is diversification – the act of selecting assets which are not perfectly positively correlated to lower the risk of the portfolio without a corresponding decrease in expected return – and the determination of an "efficient frontier" that represents portfolios offering highest expected return at every risk level (Boyd et al., 2024, p. 1 3). Although MPT is still an important foundation of portfolio management, its useful application is also frequently undermined by its underlying assumptions – normality of asset returns (neglecting tail risks), rationality of investors, and stability of correlations and volatilities over time (ResearchGate User, 20 Scaletta, 2024; Kayumov, 2023 26). AI approaches specifically DRL, advanced statistical modeling, are becoming more frequently used to address these restrictions by extracting non-linear dynamics and adjusting better to changing market conditions, rather than the conventional MPT optimization (ResearchGate User, 2024 6; Scaletta, 2024).

Adjacent with this assumption of rationality by the MPT is Efficient Market Hypothesis (EMH), which predicts that the assets prices reflect all available information (Fama, 1960 29). At its best, EMH suggests that one is unable to systematically attain performance that is superior to the wider market average on a risk-adjusted formula, as any new piece of information is immediately reflected in prices. This is a theoretical challenge to active trading strategies such as those which use algorithms and AI to take advantage of pattern or inefficiency in the market (MDPI Journal, 2024 74; Shanghai United Publishing Service (SUPS), 2022 66; ResearchGate User, 2024 11). However, experience tells us there are market anomalies and predictable patterns, especially if one looks at it with advanced AI/ML techniques (MDPI Journal, 2024 74; ResearchGate User, 2024 9). The Adaptive

Market Hypothesis (AMH) provides potential re-conciliation, where they propose that there is no static market efficiency, but evolving efficiency, informed by market conditions, competition levels, and adaptive behavior on part of the market participant (Lo, 2019 2). In the AMH framework, AI-based processes may succeed in being faster than human traders or standard models in response to changes of efficiency levels in the market (Santos et al., 2023, p. 16 15).

Behavioural Finance expressly takes issue with the assumptions regarding the perfect rationality taken for granted in MPT and EMH. It synthesizes ideas in psychology to see how cognitive prejudices and emotional factors impact investor choices and market results (NumberAnalytics, n.d. 75). Popular biases are overconfidence in predictions made by oneself, anchoring to arbitrary information, and herding – when investors follow others and, along with an increased probability of bubbles, may intensify crashes (NumberAnalytics, n.d. 75). It is in complex interactions that behavioral concepts interact with the rise of AI in finance. AI-based sentiment analysis aims at numerating market moods retrieved from news and social writing, including behavioral influences in trade equations (Abdullah & Chowdhury 2023 16; Zhu, 2024 64; Konstantinidis et al., 2024 63). However, in the process of their learning and mirroring the human biases contained in the data, AI systems themselves may unintentionally learn these biases, or their massive scale deployment may facilitate new algorithmic herding, which may increase, rather than decrease market volatility (Liu et al., 2024 17; The Global Treasurer, 2025 35).

Another approach to matters of departures from rationality, perfectly rational, is Bounded Rationality introduced by Herbert Simon (Simon, mid-20th Century 76). It states that constraint under which decision-maker, human or artificial, makes his or her decision resides in the limitations of information, the limits of cognitive or computational resources, and the pressures of time (The Decision Lab, n.d. 88). As a result, agents tend to "satisfice", narrow in on solving problems that are "good enough"- rather than making computationally intractable optimization to discover the best possible result (NumberAnalytics, n.d. 77). This idea is not only important for making sense of heuristics and shortcuts used by human investors but also for possible descriptions of behavior of AI agents. Powerful AI systems are also subject to computational limits and process incomplete data and, in complex real time applications like HFT, may use heuristic algorithms or approximation (NumberAnalytics; n.d. 76). Agentic AI, striving for complex reasoning, still acts along the lines of its algorithms, training data, and computational space.

Lastly, Agency Theory looks at situations where one group (the principal) entrusts work or authority regarding decision-making to a second group (the agent) with potential conflicts of interest unless motivators for the agent are linked perfectly to the principal's purposes. In the field of finance, such is most often the case with the relationship between investors (principals) and fund managers or advisors (agents). The advent of AI agents gives this theory a different form (Baldacci, 2021 5). Does an AI portfolio manager remain a mere tool in the hands of the human adviser (agent), with or without the human adviser assuming the role as agent on behalf of the investor (principal)? While AI systems are becoming autonomous, especially Agentic AI, it becomes a crucial challenge to ensure that AI is aligned with principal's interest (Kanerika, n.d. 20; Fujitsu, 2025 86). Agency problem-based concerns such as trust, transparency, control and AI agents having emergent goals that diverge from their purpose are fundamental concerns (A3Logics, n.d. 89; Li et al., 2024a 18).

Integration of these frameworks shows that advanced AI dynamically intertwines with the theoretical landscape of finance. AI systems could take advantage of market inefficiencies that result from either behavioral biases or bounded rationality hence undermining strong forms of the EMH 17. They could be better ways to diversify and manage risks than the traditional MPT dealing with complex data and dynamics. However, these AI systems, themselves, are operated under computational constraints (Bounded Rationality) and induce novel challenges of principal-agent nature on control and alignment (Agency Theory .5). To truly evaluate the true impact and potential of Agentic and Multimodal AI, understanding this delicate interplay is necessary, with these established theories used to measure to what degree AI capabilities cohere, collide, or refashion underlying assumptions of markets and decision making.

2.2 Theory of Reasoned Action / Theory of Planned Behavior

Although the theories described are mostly concerned with market dynamics and decision-making procedures, to understand the adoption and use of new financial technologies, such as advanced AI systems, the utilization of the behavioral models directed towards technology acceptance is needed. Useful theoretical frameworks to be used in that purpose are the Theory of Reasoned Action (TRA) as well as its extension, the Theory of Planned Behavior (TPB).

Depending on the individual's intention of that behavior, which is in association with two factors, TRA believes that one's behavior is the one that is going to take place. Their opinion as to the behavior (their positive or negative appraisal of doing it) and subjective norms (their view on pressures from society to do or not do it). TPB adds a third mechanism that is influperforming the behavior.

In the world of finance, TRA/TPB can be used to comprehend the actions of investors, such as the intention to invest in a particular asset class including what is relevant here, and the use of new financial technologies. Studies have been made on variables of influencing the use of automated systems such as robo-advisors and chatbots in finance based on the ideas of the concepts of trust, perceived usefulness, ease of use, and social influence (Ng et al., 2020 84).

The use of these frameworks for the adoption of Agentic and Multimodal AI in wealth and portfolio management offers various considerations that may be taken. A willingness of an investors to try an AI-driven tool will rely heavily on their attitude towards it – are they inclined to believe it will bring better financial results or they are not inclined to believe this is not their problem or too risky. Subjective norms, for instance, recommendations from peers or the established financial institutions, may also be involved. Perceived behavioral control is of special interest for complex AI systems; Confidence for an investor about his/her ability to understand, configure, monitor, and trust on the AI agent would have a significant influence on his/her willingness to embrace it (Ng et al., 2020 87).

Research that focuses on chatbots in finance indicates the intricate contribution of such factors included social presence (Ng et al. 2020 87). Though making a chatbot less robot-like (e.g., using a name like 'Emma' as opposed to 'XRO23', using empathetic language) increased the feeling of the social presence, it did not necessarily mean that one would trust this chatbot more or be less concerned about his/her privacy in a financial context. It is interesting that the users were more inclined to provide sensitive financial information to the less human-like bot, leaning towards technological capability compared to social-emotional features in connection to finance (Ng et al., 2020, pp. 1, 5 87). However, perceived humanness and trust in the bot positively affected the global intent to use the chatbot (Ng et al., 2020, p. 1 87). Privacy concern is one of the most essential aspects considering how sensitive the monetary data used by such systems are (Ng et al., 2020, p. 3 87).

Table 2.1 systematically maps these TRA/TPB construct in the context of adopting AI agents in finance, which means that this table displays the psychological factors influencing user behavior and acceptance in a systematic manner. This framework is useful to translate theories and make them concrete examples of what is relevant to the thesis grounding the issue in an empirical nature of trust, privacy, and user interaction with financial AI. Such an organized understanding can benefit in exploring adoption challenges, trustworthy AI interface design as well as thinking of the human side in RLHF techniques, which rely on the interaction between humans.

Table 2.1 TRA/TPB Constructs Applied to AI Agent Adoption in Finance

Construct	**Definition**	**Application Example in AI Finance Adoption**
Attitude towards Behavior	Individual's positive or negative evaluation of using the AI agent	Investors believe using an AI portfolio manager will lead to better returns (positive) or is too risky (negative).
Subjective Norm	Perceived social pressure to use or not use the AI agent	Peers or financial advisors recommend (or discourage) using AI trading tools.
Perceived Behavioral Control	Perceived ease or difficulty of using the AI agent	Investors feel confident (or lack confidence) in their ability to set up, monitor, or understand the AI agent's actions.
Intention to Use	Individual's readiness to perform the behavior (use the AI agent)	Investor plans to allocate funds to an AI-managed portfolio or use an AI research tool.
Trust	Belief in the reliability, integrity, and ability of the AI agent	Investors trust the AI agent to manage assets competently and securely, influenced by perceived humanness, reliability, security.
Perceived Social Presence	Feeling of interacting with a social entity rather than just a machine	Chatbot using empathetic language or personalization increases feeling of interaction but may not directly increase trust for finance.

Privacy Concerns	Concerns about the collection, use, and security of personal financial data	Investors worry about data breaches or misuse of sensitive financial information shared with the AI agent.

It needs to be added, though, that although TRA/TPB provides a strong basis for the explanation of the user adoption and interaction with the AI tools, its application to the modeling of the internal decision-making processes of the AI agents themselves is somewhat more complicated. AI systems, particularly the ones based on either DRL or large language models, work on learned policies, optimization goals, and data-driven pattern recognitions but not on human-like beliefs, attitudes, and intentions. Anthropomorphic attribution of such constructs directly to AI may be misleading. TRA/TPB therefore is best utilized in the configuration of the human-AI interaction layer – comprehension of trust, acceptance, usage patterns and not as a model of the AI's 'thinking internal'.

2.3 Agentic AI, Multimodal AI, RLHF, HFT for Quantitative Finance with Algorithmic Trading Theory

This section integrates the literature about the central technological elements that are the focus of this thesis: Agentic AI, Multimodal AI, their interaction with algorithmic (including HFT/HPC) trading and the use of reinforcement learning (RL) and RLHF and their application in wealth and portfolio management as well as platforms for developing them.

2.3.1 Agentic AI in Quantitative Finance

Agentic artificial intelligence systems also indicate an important step forward from the standpoint of the traditional AI that was oriented on a tight specification of a certain predictive or classification task. Agentic AI seeks to build systems that can accomplish the complicated multi-horizon objectives through elaborate reasoning, planning, and autonomous activities with minimal human oversight (Li et al., 2024a 18; Acharya et al., 2025 18). These systems share such properties as their autonomy, goal-orientation, ability to interact with their environment (including other agents, and potentially humans), and ability for reflection and planning (Kanerika, n.d. 20; Fujitsu, 2025 86; Li et al., 2024a 18). Large Language Models (LLMs) have become potential enabling technology, offering the natural language understanding, reasoning, and the generalization abilities commonly used in agentic frameworks (Mukherjee et al., 2025 19; Fujitsu, 2025 86).

There are a few architectural patterns and agentic AI frameworks for developing such AI for finance. Some of them are general purposes and others have a specific focus on finance. General frameworks such as LangGraph, CrewAI, and AutoGen provide modular functions related to constructing agentic applications and are being considered for money-related duties (Joshi, 2025, pp. 119-121 13; ResearchGate User, 2024 40). Some of the more finance-specific frameworks that have been outlined in the more recent literature include:

- **FinRobot:** An open-source platform from the AI4Finance Foundation for a variety of specialized financial AI agents based on LLMs, having the layered architecture for the data processing, algorithm selection, agent communication, and applicability deployment (Yang et al., 2024 42; Yang et al., 2024 41). It applies a financial Chain-of-Thought (CoT) prompting method to solve complex problems (Yang et al., 2024, p. 1 44; Yang et al., 2024, p. 6 91). In particular, the use of automated equity research and valuation falls under its applications (Zhou et al., 2024 43).
- **AlphaAgent:** A particular autonomous framework for alpha factor mining, based around LLM agents governed by regularization strategies (originality enforcement, hypothesis matching, complexity control) to find resistant alphas to decay (Tang et al., 2025 46; ResearchGate User, 2025 93).
    - The tempting aspect of being able to automatically identify new trading signals or what is said to be an alpha factor using AI is that a major form of risk is that only overly complex and spurious patterns are identified as being sufficient, each discovered technique is really a spurious pattern providing only the illusion of profitability being realised as the historical data is merely echoing back at itself.
    - The AlphaAgent framework may be seen as an instance of a design that seeks to explicitly solve this type of overfitting when deployed to LLM-based alpha mining. It employs several new regularization methods that can bias the agent search in directions that will find factors that are more robust and general.

  These containing regularization strategies are:

- **Originality Enforcement:** The framework punishes the agent on finding components highly correlated with already identified known common components (i.e., easy momentum or value). This will motivate the agent to seek new source of alpha instead of finding signals that are already known.
- **Complexity Control:** Agent is steered to use simpler mathematical formulas as opposed to complex ones. This follows the rule of Occam: a simpler explanation (or factor) working equally well is bound to be sounder and less overfitted than a complex one.
- **Hypothesis Matching:** The search of the agent can be limited in such a manner that the search does not explore relationships that are patently illogical by predefining economic or financial hypotheses.

Such methods as AlphaAgent bypass mere statistical validation by factoring these constraints directly in the objective function of the agent. They impose a type of semantic regularization in that the identified alpha factors must not only be statistically significant during extensive backtests, but also innovative, sparse and likely economically reasonable. It is an important aspect of a very disciplined method of exploration, preventing the constructing of overfit strategies and finding alpha signals that are more likely to survive in the future.

- **TradingGPT / FinMem:** A multi-agent setting that employs layered forms of memory structures (short, mid, long-term) with programmed decay processes, to mimick human thinking for better financial decision-making and trading'(Li et al., 2023 49; Yu et al., 2024 50). Agents may have individual characters, and they might engage in debate with each other (Li et al., 2023 51).
- **FinAgent:** A multimodal basis agent for trading, that can process numerical, textual and visual data, that contains a dual-level reflection module and tool augmentation (Zhang et al., 2024 21; ResearchGate User, 2024 40).

Important architectural protection facilities are:

• **Reflection Modules:** These elements will enable an agent to reflect on its performance in a retrospective manner. Once making a set of trades, it would be possible to see through a reflection module the reasons behind bad outcomes and decisions. This reflection about self can be used to revise the future strategy of the agent, so it learns its errors in an organized manner.

• **Tool Augmentation:** Agentic framework also aims at utilizing external tools. Restricting the types of tools to be used and the ways to use tools, developers can establish tough guardrails on the way the agent should act. To give an example, an agent might be restricted to access an execution API when some risk levels are exceeded.

• **Memory Systems:** An agent benefits from the use of sophisticated memory systems that enable it to learn about a rich history of past actions and their results. This can assist the agent to notice and learn not to repeat the errors it has made in the past and develop a stronger idea about the long-term repercussions of its actions.

- **Other Agentic Frameworks:** Add HedgeAgents (with an emphasis on balanced trading) and FLAG-Trader (fusing LLMs with gradient-based RL), FinRipple (event ripple effects analysis), Crypto Trade (for cryptocurrency trading based on the on/off-chain data and reflection), FinVision (multi-agent for stock prediction), TradExpert (mixed-of-experts LLMs for trading), FinBloom (knowledge grounding with real-time data), FinRL-DeepSeek (LLM infused RL agents), and TradingAgents(multi-agent LLM framework).

Some characteristics shared by these frameworks often include mechanisms for memory (e.g., streams, layers) to capture context (Park et al., 2023 referencing 19), abilities to reflect for learning from previously taken actions, planning modules to organize tasks, and the ability to leverage the external tools (APIs, databases, software) to enhance capabilities (Park Zhang et al., 2024 21; Fujitsu, 2025 86). What is also frequent is the multi-agent systems in which agents cooperate or debate one another with rather frequent use of manager-specialist structures (Mukherjee et al., 2025 19; Yang et al (2024) 44; Li et al (2023) 51).

The scope of quantitative finance to which agentic AI can be applied is wide and is growing by the day.

- **Automated Trading and Execution:** Autonomous agent in the making of trading decisions (buy/sell/hold) using complex analyses and potential optimization of order execution strategies (Mukherjee et al., 2025 19; Kanerika, n.d. 20; A3Logics, n.d. 89; Riahi Samani et al., 2024 ; Zeng et al., 2024 29; Li et al., 2023 51; Gu et al., 2024 39) This builds on capabilities developed in such frameworks such as FinRL.
- **Alpha Factor Mining:** Self-guiding search agents scouting the wide realm of potential mathematical formulas or patterns for discovery of new, predictive and interpretable alpha factors that perform better than the conventional or manual

discovery methods (Tang et al., 2025; Ren et al., 2024; Xu et al., 2024; Zhao et al., 2024; Yu et al., 2023; Kou et al., 2025).

- **Portfolio Management and Optimization:** Autonomous agents that are dynamically managing and rebalancing investment portfolios based on real-time data, learned strategies, and possibly multimodal inputs, to make optimal risk-adjusted returns (Sun et al., 2025 59; Sun et al., 2025 60; Zhao, 2024 61; Gu et al., 2024 83).

- **Equity Research and Valuation:** Agents performing the process of data collection, conducting qualitative and quantitative analysis and preparing all-rounded equity research reports like the human analyst's workflow (Zhou et al., 2024 43; Yang et al., 2024 92).

- **Risk Management and Fraud Detection:** Agents used for jobs such as model risk management (validation of models, and compliance), identification of market risks, detecting acts of fraudulent nature, and regulatory compliance (Mukherjee et al (2025 19), Kanerika (n.d. 20), Joshi (2025 13).

- **Market Simulation:** Application of the multi-agent systems to simulate complex market dynamics, postulate the economic behavior hypothesis and experiment with trading strategies in different situations (Wang et al., 2024 ref; Zheng et al., 2024 ref).

Despite its huge potential, development and implementation of agentic AI in finance have daunting barriers. LLM-based agents that are used at present are limited in long-horizon problem-solving, being capable of meaningful interaction with complex environments, and a common-sense application (Li et al., 2024a 18). They are likely to hallucinate creating plausible but wrong information, have difficulties with reasoning in the temporal and have vulnerabilities to adversarial input (Chen et al., 2025 37). Misalignment, where agents work towards goals in unintended (or harmful) ways is a significant danger and may take form as deception or manipulation (Greenblatt et al., 2024 18; Chen et al., 2025 37). Making sure the system is transparent, interpretable and has strong safety mechanisms is important, but difficult, particularly is the whole thing is based on many "black box" models (Joshi, 2025 13; Chen et al., 2025 37). Moreover, emergent, unpredictable interactions between several autonomous agents belonging to the same market create systemic risks' issues (The Global Treasurer, 2025 35).

The emergence of Agentic AI represents a possible paradigm de-scope of quantitative finance from AI as a calculative apparatus to AI as an autonomous collaborator or decision-maker. The attraction is in automating complex workflows from end-to-end that used to require much human expertise, such as equity research or active portfolio management. However, it requires to overcome dramatic issues with the reliability, controllability, and intent alignment of the agents; as well as with the management of the systemic risks of deploying the autonomous entities in the high-stakes financial markets. Research and development, therefore, needs to address not only how to make the capabilities of the agents better but also how to establish strong validation, governance, and risk mitigation frameworks.

2.3.2 Multimodal AI in Quantitative Finance

Multimodal Artificial Intelligence indicates those systems that have been created to process, understand and combine the information from various modalities of data simultaneously. In finance this would usually entail joining classical quantitative data (such as the stock prices/trading volumes /financial ratios) along with unstructured/semi-structured data such as text-based news articles/social media posts/earnings calls transcripts/financial reports and even visual data like candlestick charts/ or technical analysis diagrams (Zhang et al., 2024 21). This is asserted by (Bhatia et al., 2024 22; Multimodal LLMs Survey 23). The center argument is that financial markets receive a rich carpet of information and that their integration could give a more holistic, context-aware insight than the use of solely numerical time-series key (Zhang et al., 2024 21).

A new trend toward multimodal financial analysis has been observed, and researchers have aimed at designing certain models and architectures for multimodal financial analysis.

- **FinAgent:** FinAgent is positioned to be a multimodal foundational agent for financial trading, in which it explicitly combines numerical, textual and visual data streams. It uses a specialized market intelligence module to process such diverse data, followed by a dual-level reflection and diversified memory retrieval system to learn multimodal patterns from the past and keep pace with market trends (Zhang et al., 2024 21; ResearchGate User, 2024 40).
    - The heterogeneous nature of this data is the reason why multimodal AI systems, like the FinAgent framework, are implemented to consume, process, and unify such multiple data streams in order to compose a more comprehensive and contextual picture of the state of the market. The main advantage of this strategy is not that more data would be accessed, but that the different types of data could be checked and disambiguate each other.
    - The weak signal in the price data may be corroborated by the strong signal in the news sentiment (and vice versa), or disordered news story may be disordered by the clear trend in the numerical data (and vice versa). The

human multimodal model acquires such intricate cross-modal association and coupling.

- This method of confirming signals largely enables the model to utilize qualitative data as a sieve of quantitative data.
- The system has the potential of maximizing its signal-to-noise ratio by making it difficult to signal across several modalities without a consistent match before emitting the trading signal.
- This results into stronger and more stable decision making, particularly when there are strong episodes of uncertainty or when market direction is influenced by news dynamics other than genuine price action in the market.
- Multimodal AI therefore tries to mimic the reasoning of an expert in the human mind, which can combine chart, news feeds, and reports information intuitively and build a complete picture of how the market thinks.

Leverage specialized frameworks for effective data processing and integration

- The intricacy of combining multimodal data cannot possibly be averted without resorting to special frameworks that would support all the stages of the data to be ingested and fused.
- FinAgent model is a very good example of a base agent that is created in the context of this particular goal in financial trading. It is specifically do-able to merge numerical, textual and even graphical data streams.
- To handle this difficulty, FinAgent has several important components in its architecture:
  - **Market Intelligence Module:** The market intelligence module is the specific module that is to do the preliminary processing of the various streams of data and extract relevant features in each modality.
  - **Dual-Level Reflection:** The agent has a reflection mechanism which enables it to learn its previous multimodal choices hence enabling it to improve its fusion strategy as it goes.
  - **Diversified Memory Retrieval:** The type of system has a complex memory system utilized to make storage and recall relevant past multimodal patterns where the system can identify reoccurring situations in the market that entail a mix of various information.
- Through a structured and integrated pipeline through which different data can be processed and learnt, frameworks such as FinAgent take a lot of the engineering complexity.

    - They provide a methodical way of addressing issues surrounding the issues of data alignment, data feature extraction and data fusion to create more efficacious and robust multimodal trading agents.

- **Open-FinLLMs (FinLLaMA/FinLLaVA):** This initiative is on open-source multimodal financial Large Language Models. FinLLaMA is pre-trained over a large corpus of combined text, tabular and time series data. FinLLaVA builds upon this using the LLaVA framework, for multimodal instruction tuning, and is able to process and reason on text, tables, time-series, and chart images using curated datasets such as financial images; OCR-VQA and annotated table images (Bhatia et al., 2024 22; Liu et al., 2024).

Other studies assume multimodality because multimodality is achieved via the combination of various data types in the models, for example, sentiment in news is combined with the prices for prediction or trading (Konstantinidis et al, 2024 63) using textual data together with market data in LLM-based agents (Ding et al., 2024 73; Zhou & Mehra, 2025 106), or combining the textual data with market data for agents (Abdullah & Chowdhury, 2023 16; Zhu, 2024 64; Gu et al., 2024 39; Hou et al., 2021 105)

The key areas of use of multimodal AI in quantitative finance are related to increased understanding of the market dynamics and enhanced predictive power.

- **Enhanced Market Analysis and Forecasting:** By integrating quantitative indicators with qualitative information that can be retrieved from news, reports or the sentiment from social media, multimodal models attempt to create more accurate and sturdy market forecasts (Zhang et al., 2024 21; Abdullah & Chowdhury, 2023 16; Konstantinidis et al., 2024 63; Huang et al., 2024). This entails prediction of stock trends, determination of market regimes, or prediction of volatility.
- **Improved Trading Decisions:** An example of such a trading agent is FinAgent that employs multimodal inputs to inform its trading strategy (Zhang et al., 2024 21). Another popular use is the inclusion of sentiment scores or news analysis to the algorithmic trading signals directly (Konstantinidis et al., 2024 63; Zhou & Mehra, 2025 106).
- **Richer Risk Assessment:** Multimodal analysis can be used to apply qualitative factors (e.g. geopolitical news, company-specific events in text) into risk models to offer earlier warnings or more complex risk profiles than quantitative measures alone (Zhu, 2024 64).

Inspite of all this promise, multimodal AI has its own share of challenges. Heterogeneous data integration involves complex fusion methods to coordinate the information in time and meaning (Bhatia et al., 2024 22). Different modalities may give conflicting signals, which require mechanisms of reconciliation or weighting. The computational expense of operating on several data flows, and particularly numerous amounts of text or images, as well as frequent numbers, can be significant (Multimodal LLMs Survey 23). Besides, understanding of the complex interaction that is learned by multimodal models, as well as the explanation of their decisions is also a very challenging task and reduces trust and validation (Chen et al., 2025 37).

Multimodal AI attempts to break away from purely quantitative approaches by including the richer context found in the written, visual, or other forms of data. This is in line with how the human experts approach the analysis of markets through integrating the news, reports, and charts with the price movements. The opportunity is to build AI systems that are more context-aware and therefore produce potentially more robust predictions and decisions, especially in periods dominated by news events or change in the market sentiment. Nevertheless, practical realization of that potential depends on overcoming the technical obstacles for data fusion, managing conflicting information, computational costs, ensuring interpretability of such models.

2.3.3 Algorithmic Trading & High-Frequency Trading (HFT) with AI/HPC

Algorithm trading whereby computer programs are used to execute trading orders according to the predetermined instructions have transformed the financial market since the introduction of electronic trading (Azzutti ,2024: 8; Baldacci, 2021 5). High-Frequency Trading (HFT) is an extreme version of algorithmic trading which is characterized by lightning speed (milliseconds or microseconds), high order-to-trade ratio, strategies to exploit tiny arbitrages or liquidity rebates (ResearchGate User, 2024 9; Son, 2022 10; MacKenzie, 2018). The share of HFT in the aggregate volume of trading has thus grown to a significant size in most of the world's largest markets (Azzutti 2024 8).

The combination of Artificial Intelligence and Machine Learning has immensely grown the complexities in algorithmic trading to go beyond rule-based trading. AI/ML techniques are used to:

- **Improve Predictive Models:** Deep learning models (such as LSTMs, CNNs, Transformers) and reinforcement learning agents are more and more used to predict price moves or to detect trading signals, bringing better results than traditional statistics or technical indicators (Cao et al., 2025 2; Liu et al., 2024 17; Yu & Wu, 2022 107).

- **Enhance Pattern Recognition:** AI algorithms can detect complex patterns in high dimensional market data which can be unseen with human traders or simpler algorithms (Azzutti, 2024 8).
    - o Such limitations are the reason why advanced AI methods and models, specifically Deep Learning, such as Long Short-Term Memory (LSTMs) and Transformers are carefully designed to address these problems.
    - o The architectures are very capable of learning the very complex and non-linear dependencies and long-term temporal structures on raw information, which are very high dimensional.
    - o They have a layered nature that enables them to implicitly carry out feature engineering that is hierarchical.
    - o Deeper network layers can learn to recognize simple but noisy patterns, whereas the higher level learn how to compose those into more-abstract, robust and, finally, more-predictive signals.
    - o The noise-filtering, signal-extraction part is done by the model itself, which is an important advantage compared to the methods that use pre-processed features labelled as clean.
- **Integrate Alternative Data:** AI, especially NLP, allows the consideration of sentiment analysis within news and social media into real-time trade decisions (Konstantinidis et al., 2024 63; Zhou & Mehra, 2025 106; Abdullah & Chowdhury, 2023 16).
- **Develop Adaptive Strategies:** By using reinforcement learning, agents can learn and develop their trading strategies on the go by taking market responses and modified conditions into account (Guarino et al., 2022 108; Yu & Wu, 2022 107; Bai et al., 2022 109).
    - o Another effective framework for operating in noisy environments is Reinforcement Learning. An RL agent does not learn good policies by learning the statistical properties of the data explicitly, but by trial-and-error interaction with the market.
    - o The agent is rewarded and punished based on result actions that bring about positive elements (e.g. profit).
    - o During the numerous iterations the agent develops a habit of disregarding random noise since when agents perform acts only in response to random noise, they fail to achieve a pattern of consistent rewards.
    - o This is implicitly taught to filter the data, and only the small, frequent but patterns are to be learned which are predictive at the future rewards.

- Through this automatic machine learning, the RL agents will be able to come up with vigorous techniques capable of detecting and capitalizing on weak signals hidden in a noisy data feed.
- These new and improved AI techniques have the strength to go beyond the limitations of the earlier models of assuming being something on which a rigid set of rules are made.
- They can learn based on large amounts of data and change the way they represent the information internally to detect small but consistent patterns that humans and more basic algorithms can never observe.
- This ability puts a trader in a position to exploit an advantage in the contemporary financial markets, in which the actual signal is low and transient.

Such platforms as FinRL and Qlib are vital in terms of supplying the necessary infrastructure and supporting tools for the development and testing of these AI-driven algorithmic trading strategies, working in standardized conditions and featuring pre-built models.

Typical HFT strategies involve such approaches as market making (providing liquidity by posting bid and ask orders), statistical arbitrage (capitalizing on minor and momentary price differences between related assets), momentum ignition (finding and enhancing short-term trends), and news-based trading (instant response to news releases) are all affected by AI. AI can help fine-tune price prediction for arbitrage, provide faster and more precise responses to news sentiment, or look after the optimal quote's placement of markets.

For HFT, where latency needs to be kept extremely low and for complex AI models where large amounts of computation is needed for training and inference, High-Performance Computing (HPC) infrastructure must be used (Penguin Solutions, n.d. 68; IBM, n.d. 69). This encompasses special hardware such as GPUs that process in parallel, low latency network, optimized software stacks to attend to real-time data streams, train deep learning models (especially DRL) and carry out trades in microseconds (Penguin Solutions, n.d. 68; Li et al., 2022 70). Such frameworks as FinRL-podracer are created for the effective use of HPC for scalable DRL in finance (Li et al., 2022 70).

Effects of AI-based algorithmic trading (HFT in particular) in terms of the market quality is still a debate. Although it can increase the efficiency of price discoveries and narrowing bid-ask spreads to facilitate market liquidity (Oxford University Press, 2024 12), There are concerns regarding its ability to heighten volatility incurring flash crashes (like the May 2010 event 12), (ResearchGate User, 2024 11).

It enables new forms of market manipulation, and contributes to systemic risk, if a large amount of AI agents carries correlated strategies and respond in a coordinated manner to market events (ResearchGate User, 2024 9; Liu et al. 2024 17; The Global Treasurer 2025 35). This has elicited more regulatory examination in proposals in change of market designs (Azzutti, 2024 8). Son, 2022 10).

One of the changes is the introduction of Frequent Batch Auctions (FBAs) which substitutes continuous trading with discrete auctions at short periods of time (i.e., milliseconds) to curtail the speed advantages of HFT firms (Savidge, 2023 113; Budish et al., 2015 114). Simulation studies are looking at how various trading algorithms adjust and work in FBA markets (Savidge, 2023 113).

The incorporation of AI in algorithmic trading, particularly that of HFT, reflects the escalation of market dynamics. It's not just a continuation of the "speed arms race" inherent in HFT but also introduces a "complexity arms race," where competitive advantage relies on sophisticated AI models capable of processing vast data and adapting strategies. This double race requires high-powered HPC resources and begs bigger questions of market stability and equity. The sophisticated relationships between multitudes of fast, adaptive AI agents may create emergent market tendencies that are hard to forecast or manage, throwing the efficacy of typical market institutions and regulatory surveillance into question. Market design innovations such as FBAs try to reverse speed-based advantage but the evolvability of AI implies there is a continuing co-evolutionary relationship between AI-based trading strategies and the mechanisms to check them.

The proposed prominent market design concept to accomplish this is termed Frequent Batch Auctions (FBA). Unlike continuous market: whereby trade can be transacted at every moment, an FBA market structure accumulates all the orders placed during a very small-time span (e.g. 100 milliseconds), and then executes them at once, at a uniform clearing price.

This is a discrete, periodic clearing, which transforms the nature of competition. Since all bundles in the batch are assumed to be reached at the same time, the gain of a microsecond or two coming in front of a close competitor is removed.

It becomes not a race about who is the fastest, but on who will cleverly guess the clearing price in the batch. About countering the gain in speed, FBAs can minimize the motivation of companies to invest in the expensive and possibly disrupting technological race.

Studies of simulation are testing various trading algorithms in adaptation to FBA markets with increased argument that they can achieve more stable and equitable market sets. An essential, structural means of preventing systemic risks created by interacting high-speed AI agents, is to support such innovations in market design.

2.3.4 Reinforcement Learning (RL) & RLHF in Quantitative Finance

The reinforcement learning (RL), especially the deep reinforcement learning (DRL) that applies deep neural networks as function approximators, has become an influential paradigm for solving sequential decision-making problems in quantitative finance (Scaletta, 2024; Kayumov, 2023 116; Ezeilo & Nimo, 2021 58).

Tasks like optimizing trading actions over time or dynamically allocating assets in a portfolio are tasks that naturally fit the RL framework, where an agent learns through interaction with environment (market) to maximize cumulative reward signal (e.g. portfolio value or risk adjusted return) (Liu et al., 2022 117; Ezeilo & Nimo, 2021 58). The problems are usually posed as Markov Decision Processes (MDPs) or given that the agent may not know the state of the market completely, Partially Observable MDPs (POMDPs) (Liu et al., 2022 117; Kabbani et al., 2022).

The availability of specialized open-source libraries and frameworks has enormously sped up the development and implementation of DRL in finance.

The notable one is the FinRL ecosystem developed by the AI4Finance Foundation.

- **FinRL:** The core library includes such a general infrastructure to assist in applying DRL in the trading of stocks in an automated fashion (Liu et al., 2020 110; Liu et al., 2022 79), with availability of standardized environments, DRL algorithm implementations, and back-testing tools.
- **FinRL-Meta:** Attempts to the pressing need for a substantial and dynamic market environment. It uses an automated DataOps framework to scrape data from the real-world markets to construct hundreds of gym-style environments, minimizing problems such as survivorship bias as well facilitating comparative and reliable benchmarking (Liu et al., 2024 34; Liu et al. 2022 81; 2022 119; 2024 120; 2023).
    - To develop models capable of dealing with non-stationary markets it is imperative to ensure that we train them in non-stationary environments.
    - The FinRL-Meta framework directly answers this need by extending on the dynamic historical datasets to offer a sampler of a virtually infinite "universe of near-real market environments", to train and benchmark DRL agents.
    - The essence of FinRL-Meta is an automated DataOps pipeline that scrapes and runs on real-world data related to market data.

- The new data is subsequently used in building hundreds of standardized, similar to the gym training environments that translate the emerging market dynamics such as volatility changes, changes in correlation, and emergence of new trends.
- This method has two advantages. On the one hand, it addresses the issue of the models trained on the outdated market conditions directly. This continuous update of training environments also ensures that FinRL-Meta agents experience new phenomena in the market and hence readier when they face live trading.
- Second, its complex approach to the sourcing of data allows us to eliminate deeply rooted data bias, the most significant one being the survivorship bias.
- With the incorporation of delisted stock or other abortive assets, the framework gives a more realistic (and pessimistic) view of market conditions, so that models never build a significantly over-optimistic picture out of only the assets that have made it.
- The overall change in model evaluation process brought using dynamic environments such as those offered by FinRL-Meta is essential. It enables the transition between the dependence on one-time backtest and a system of constant assessment.
- A strategy is always possible to test and revise in accordance with the recent market data fed into the pipeline of DataOps. This is somewhat like a rolling backtest, or constant paper trading, an unrealistic way to gauge a strategy on an on-going basis.
- It solves the fatal sim-to-real problem where good performance on back test does not mean live market trading can be successful often due to the training (static) data not resembling the reality (dynamic) market.
- The ongoing measurement serves as an early warning system so a strategy that is losing its effectiveness gets highlighted as compared to the traditional reactive model of damage control as soon as the regime, which the strategy was built with, starts to shift.

Use automated DataOps frameworks to source high-quality, diverse data

- Such frameworks allow seeing the market in a more complete and less biased way by programmatically scraping the data about it on multiple reputable sources.
- As an example, FinRL-Meta specifically attempts to reduce the risk of survivorship bias (a widespread issue of which includes the focus on

successful companies solely during the training of the model, resulting in excessive optimism when it comes to predicting returns).
    - This data-driven, model-less design is a guarantee that models will be on the basis of high-quality, representative data, which happens to be the first and most important step towards creating reliable AI systems.

Use dynamic datasets and environments to improve generalization

    - Rather than one training set, FinRL-Meta offers an ever-growing universe of market environments. This enables a model to be trained and verified on such a diverse range of time and market scenarios.
    - The model shows a variety of bull markets and bear markets, high and low volatility periods and different correlation structures, and is taught to learn more generalizable patterns that will be true over different regimes.
    - It will not do enough just remembering the peculiarities of period. This aspect of training on varied and changing information is among the best strategies to enhance robustness of a model and its generalization to future situations that have not been experienced before.
    - This will change the training of the models into a dynamic process. This model is not merely learning about the past, but a more varied representation of the past.
    - Such variety in the training data forms a strong regularization that prevents the model naturally formatting to be transient, regime specific.
    - These strategies, in turn, will be less inclined to be curve-fit to a particular historical time and much more likely to have the versatility required in the dynamically evolving environment of live financial markets.

- **FinRL-Podracer:** Aims at high performance and scalability that utilizes parallel computing on GPUs to speed up training of DRL agents, important for large-state spaces and complex models (Li et al., 2022 70).
    - Besides the hardware, it is important to apply software frameworks which are specifically built to take advantage of parallel computing infrastructure.
    - A perfect example of such a framework is FinRL-Podracer which was designed with a specific purpose to speed up DRL agent training in the field of finance.
    - DRL is extremely sample-inefficient, and therefore it needs extremely many interactions with a learning environment to build the policy that works. On

a single machine it may take prohibitively long to train in such complex financial situations.

  - FinRL-Podracer solves this by offering a model that scales to use the power of many GPUs to train the DRL. It enables parallel gathering of experience by multiple "actors" (agents acting on the environment) which is then effectively combined to revise the central learning model.
  - The structure can speed up training by orders of magnitude, something which enables training of complex agents in large state spaces, and over long historical periods.
  - The embrace of such scalable frameworks would be necessary in making it viable and feasible to develop advanced DRL-based financial strategies.
- **POE (Portfolio Optimization Environment):** A particular environment in the FinRL ecosystem to accommodate the portfolio optimization processes, to be used with RL frameworks in an intuitive manner (Costa & Costa, 2023 123).

Optimize algorithms for efficiency through lightweight models and pruning

Where HPC just gives the brute force to solve large model's algorithmic optimization can provide a more elegant solution to minimize computational burden in the first place. That is, it is the process of designing or manipulation of algorithms to become more efficient without the sacrifice of much performance.

The two main methods of this are:

• **Lightweight Models:** A billion-parameter model is not essential to every task. In some specifiable forecasting or classification tasks, much more efficient architectures (comparable in performance, but again of a much smaller size, occasionally known as the "lightweight model") can outperform much larger and expensive architectures but using only a small fraction of computational overhead. Polynomial classifier or smaller versions of transformers on financial tasks are an active line of research, which provides an opportunity to achieve a balance between accuracy and efficiency. It is especially applicable where it is deployed in edge devices, or as part of a latency conscious application.

• **Pruning:** This is a power method that causes contraction of an already trained neural network. It entails detecting and eliminating the duplicate or useless aspects of the model, comprising of individual weights, neurons or even an entire layer, which does not make strong contributions to the overall performance of the model. Pruning may exponentially decrease the memory and computing unwanted by a model and may accelerate inference and make it more affordable. As an illustration, the pruned fraud detection model may process real-time transactions at a relatively lower cost and latency.

Allowing HPC to harness the intelligence of scalable software frameworks as well as the efficiency of optimized algorithms, financial institutions will be capable of managing the computational cost of implementing powerful AI tools, which ultimately makes such powerful tools technologically feasible as well as economically viable.

Adaptive algorithms - Implement adaptive algorithms to learn and switch between multiple trading patterns

This is the solution that tries to implement adjustability in the very structure of the AI model itself. Instead of using one monolithic model that would have to be trained to execute effectively in any condition of the market, a hard task, what is used here is something that stores a collection of specialized "expert" sub-models that can take care of specific tasks.

All of them are prepared to work in a particular market pattern or regime. To give an example, a system could include a high-volatility momentum expert, a low-volatility mean-reversion expert and a trending market expert.

- Architectures such as the Temporal Routing Adaptor, discussed by Lin and Zhou (2021) realize this idea by using a high-level router component. This router is “dynamic” and analyzes the present situation within the market in real time and learns on its own to dynamically give the processing power/focus to the most pertinent expert who can address the current situation.
- The router cosmetically matters the transaction when the market is quiet and exhibits a range; it changes priorities when an intense trend emerges in the market.
- This enables the entire system to smoothly and immediately change its strategy without requiring an entire system to retrain, which is a time-consuming task. It goes a long way in representing the necessity of having models capable of being "abrupt market change in regime".
- Such an architectural strategy is a more advanced solution to the non-stationarity than retraining the single model. It incorporates the financial idea of different market regimes right into the AI system.
- It is not that the model is learning one predictive function; it is learning two functions in parallel: a bank of special purpose predictive functions (the experts) and a bank of more general meta-learning functions (the router) that decides when to trust which expert.
- This is the same thought that occurs in the mind of an expert human portfolio manager, who may make a conscious decision, Bentuk "The current

macroeconomic news and the volatility in the market lead me to now employ my capital preservation strategy instead of my aggressive growth strategy."

- That it is possible to improve market dynamics-flexibility in other ways as well is evidenced by the approach of directly architecting the AI to resemble this adaptive, regime-switching logic so that systems can be constructed that are naturally more flexible and adaptive to changing market conditions.

Qlib from Microsoft, although more wide-ranging, encompassing entire quantitative investment pipeline, has also support for RL, especially when it comes to order execution (Microsoft Qlib Documentation 80). Other such platforms like **TradeMaster** and **ElegantRL** (which is often a part of the framework for FinRL research) add to the collection of tools available for financial RL research.

A range of DRL algorithms are widely used in the financial literature, usually based on the Actor-Critic paradigm and consequently separate modules learn policy (actor) and estimate the value function (critic) (Yang et al., 2021 125; Yu & Wu, 2022).

Commonly cited algorithms include:

- **Value-based:** The Deep Q-Networks or DQN (Mnih et al., 2015 referenced in 117, Sanghi, 2024 127)
- **Policy-gradient/Actor-Critic:** Proximal Policy Optimization (PPO) (Schulman et al., 2017 referenced in 117; Anwar & Wang, 2022 ; Ezeilo & Nimo, 2021 lit. 58), Advantage Actor-Critic (A2C) (Mnih et al., 2016 referenced in 117; Santos et al., 2023 [15]), Deep Deterministic Policy Gradient (DDPG) (Lillicrap et al., 2015 referenced in 117); Santos et al., 2023 15), Soft Actor-Critic (SAC) (Haarnoja et al., 2018 referenced in 117; Santos et al., 2023 15), Twin Delayed DDPG (TD3) (Fujimoto et al., 2018 referenced in 58; Santos et al., 2023 15), and Trust Region Policy Optimization (TRPO) (Schulman et al., 2015 referenced in 58).

Practical examples reveal the depth of DRL in matters of finance:

- **Portfolio Optimization:** There are many works that use DRL to manage assets dynamically, rebalance portfolios, as well as control risk, showing considerably better performance against traditional benchmarks such as mean-variance optimization or simply buy-and-hold strategies. (Sun et al., 2025 59; Winkel & Strauß 2023 131; Scaletta 2024; Kayumov 2023 26; Santos et al. 2023 15; Deng et al. 2024 33; Costa & Costa 2025; Chavan et al., 2021 133; Sun et al., 2024 134; Yu & Wu, 2022 107; Niu & Li, 2022 135). Ensemble strategies of several DRL agents also look promising (Yu & Wu, 2022 107).

- **Automated Trading:** DRL agents are trained to make buy, sell, or hold decisions for singular or multiple stocks hoping to maximize profits within a simulated or real trading environment (Liu et al., 2020 110; Yang et al., 2021 136; Sun et al., 2022 57; Vishal et al., 2022; Lee & Moon, 2023 138; Ezeilo & Nimo, 2021 58). The research also includes offline RL learning from pre-collected datasets (Lee & Moon 2023 138).
- **Alpha Factor Mining:** RL frameworks are used for searching the compound space of mathematical expressions for automatically discovering new formulaic alpha factors (Xu et al., 2024 54; Zhao et al., 2024 53; Yu et al., 2023 55).

Inspite of successes, the application of DRL in finance encounters numerous challenges. The low signal-to-noise ratio and the non-stationarity of the financial markets make learning stable and generalizable policies challenging (Liu et al., 2024 34; Scaletta, 2024 139; Deng et al., 2024 33).

DRL algorithms are usually data hungry which entails large interaction with the environment which may be a problem especially there is minimal historical data or cost of Live market interaction in hand (Scaletta, 2024). Coming up with proper reward functions that reasonably represent desired outcomes (such as balancing risk and return) is non-trivial. Even uncomplicated profit maximization may cause overtly risky strategies (Winkel & Strauß, 2023 131; Sun et al., 2022 57).

Bridging the gap between performance in simulation (backtesting) and real-world trading (the "sim-to-real" gap) remains a key hurdle due to factors like market impact, latency, and data discrepancies (Scaletta, 2024 139; Liu et al., 2022 140).

The Reinforcement Learning from Human Feedback (RLHF) has appeared, mainly in the situation of alignment of LLMs, as one of the approaches to include human tastes, and assessments in the learning process (ResearchGate User, 2024 67; Bai et al., 2025 36). Instead of relying solely on a predefined reward function, RLHF typically uses human feedback on agent behavior (e.g., comparing pairs of responses or actions) to train a reward model, which then guides the RL agent's policy optimization (Bai et al., 2025 36).

In quantitative finance, RLHF has the promise of bringing trading or portfolio management agents under the control of complex and subtle human objectives that are hard to formalize mathematically, such as, precise appetites for risk, ethical considerations, compliance with mandates of investment, or the desire for explanation (ResearchGate User, 2024 67).

However, applying RLHF effectively presents its own challenges, including the potential for human labelers to provide strategic or biased feedback to manipulate the agent's behavior, leading to a trade-off between ensuring truthful feedback (strategyproofness) and achieving optimal policy alignment (Bai et al., 2025 36).

The DRL power inheres in the learning ability of complex, adaptive strategies from the experiment with the market. But the intrinsic challenges of the financial world – noise, non-stationarity, non-linear objectives – makes naive use of DRL frequently inadequate. It all requires patient environment design (aided by platforms such as FinRL-Meta 34), strong algorithms, sophisticated reward engineering using risk, and possibly, cutting-edge alignment techniques such as RLHF.

The RLHF is a major step forward into training AI agents that strive not for profitability alone but for a wider set of subjective human targets and boundaries, which is very important for trusted application in finance.

2.3.5 AI for Wealth & Portfolio Management

The traditionally human-based area of wealth and portfolio management, which operates under the model of established human advisors and financial models, is experiencing radical change via AI. The established approaches have difficulties scaling personalized advice and eliminating human behavioral biases during decision-making processes and handling the increasing complexity and volume of data as the markets become more modern. AI defines possible solutions to these challenges, which become oriented towards efficiency, personalization, and performance.

Key applications and trends include:

- **Robo-Advisors:** Such automated platforms are using algorithms, MPT formulations at first but increasingly ML formulations too, for giving investment advice and for administration of portfolios requiring little human management (Han et al., 2024 14). They provide access and cost savings when compared to the traditional advisors, but they suffer from issues of data quality, algorithm transparency, and managing complex clients' needs (Han et al., 2024 14).
- **Enhanced Personalization:** AI allows hyper-personalized financial advice and portfolio construction. Analyzing huge amounts of client records (financial situation, goals, appetite for risk, even patterns of behavior), AI can offer individual advice on investment, saving policy, insurance, retirement planning (Kanerika, n.d. 20; A3Logics, n.d. 89; The Wealth Mosaic, 2024 141). Merrill Lynch for example uses AI to predict client life events (such as home renovations causing transactions), with a view to speaking proactively on liquidity needs (The Wealth Mosaic, 2024 141).

- **Advanced Portfolio Optimization:** AI methods are specifically being used to extend the conventional MPT using DRL and other machine learning based models (ResearchGate User, 2024 6; Sehgal & Pandey, 2018 142). These methods have the potential to account for non-linear asset relationships, be dynamic and adapt to changing market situations, include alternative data sources, and optimize portfolio based on more sophisticated risk measures or objectives (ResearchGate User, 2024 6; Scaletta,2024; Zhao, 2024 61; Deng et al, 2024 33).
- **Improved Advisor productiveness and process labor automation:** AI is being rolled out as an aid to rather than replacement of human advisors. It can automatize such routine tasks as client meetings preparation, summarization of clients (emails, calls), development of individual client communications, initial data research or analysis (Forbes, 2025 78; The Wealth Mosaic, 2024 141). It releases the human advisors to do the high-value roles, such as complex financial planning, building relationship, and empathetic distribution (Forbes 2025 78). Agentic AI is envisioned as a potential "co-pilot" for advisors, providing real-time recommendations and streamlining workflows (Forbes, 2025 78).
- **Sophisticated Risk Assessment:** AI algorithms can bring together a wide range of data sources, such as macroeconomic indicators, market sentiment, as well as individual client profiles, to deliver more holistic and time-appropriate risk assessments that may forecast market declines or target specific weaknesses in portfolios (Kanerika, n.d. 20; Albahri et al., 2023 142).

There are more inventions through the combination of Agentic and Multimodal AI. The use of Agentic AI could empower wealth management platforms to carry out more autonomous functions, for instance, handle rebalances in portfolios automatically through predefined rules and real-time analysis, etc. even acting proactively on clients on identified needs / market events perhaps, working with a system of rules stipulated by human advisors (Forbes, 2025 78).

Multimodal AI can increase the fact-finding ability by working with myriad client information (e.g., text from communications, financial documents) and market data (news, charts, reports) to yield richer insights on automated systems and human advisors (Zhang et al., 2024 21). Bhatia et al., 2024 22).

However, introduction of leading AI into wealth management is a very challenging task. Trust development for the clients in automated or agentic systems is vital, especially because of the climate of financial decisions' sensitivity (Ng et al., 2020 84). Privacy and security of data issues have been raised at the topmost level (Ng et al., 2020 87; The Wealth Mosaic, 2024 141).

Compliance with regulations regarding AI-driven advice and portfolio management services is not a simple task (Forbes, 2025 78; The Wealth Mosaic, 2024 141). The "black box" nature of many AI models hinders explainability, making it difficult for advisors and clients to understand the rationale behind AI recommendations (Zhao, 2024 61). Aligning AI agent behavior with the client's best interests and complex preferences raises fundamental Agency Theory issues (Baldacci, 2021 5).

The AI evolution in wealth management follows the route from the mere automation (robo-advisors 14) to the more intelligent support of human advisors and, possibly, to highly autonomic system supported by Agentic AI. This change stresses personalization, dynamic adaptation, and incorporation of more data sources, with the help of Multimodal AI.

Although the potential gains lie in higher efficiency, accrue limited accessibility, and possibly improved investment returns, actualizing this potential might mean meeting critical challenges regarding trust and transparency, control, regulation, and an ever-changing role of the human advisor as AI informed future prevails.

2.3.6 AI Platforms & Frameworks (FinRL, Qlib)

The development and integration of specialized platforms and libraries into quantitative finance has greatly supported the rapid development of AI.

These platforms alleviate such challenges in the field, such as the need for standardized environments for training and testing algorithms, benchmarks of performance, and easy-to-use tools for implementing complex models, thereby lowering barriers to entry and reproducibility (Liu et al., 2024 34; Liu et al 2022 81; Wang et al 2024 82).

Two major representative examples for the open-source approaches to this field of study are the FinRL ecosystem (AI4Finance Foundation), and Qlib (Microsoft).

The FinRL ecosystem is mostly concerned with the democratization of both AI and Deep Reinforcement Learning (DRL) for finance applications through a collection of open-source tools.

- **FinRL Library:** However, the core component gives a structure for using DRL algorithms (such as PPO, DDPG, SAC, A2C, TD3) for performing auto-trading in stocks. It contains modules for the construction of market settings using historical data, the training of agents in these settings with an MDP formulation and evaluation via backtesting strategies with conventional financial metrics (Liu et al., 2020 110; Liu et al., 2022 79).
- **FinRL-Meta:** This extension is aimed at the problems of data quality and environmental realism in the field of financial RL, in particular. It uses a DataOps

pipeline to automatically harvest dynamic datasets from real-world markets and to develop hundreds of standardized, gym-style environments. It aims to address problems such as survivorship bias and give sound benchmarks for performance comparison of DRL policies (Liu et al., 2024 34; Liu et al., 2022 81; Liu et al., 2022 119; Liu et al., 2024 120).
- **FinRL-Podracer:** Deals with the computational needs of DRL by offering a high-performance, scalable framework that makes use of parallel architecture, especially on GPUs, to drastically speed up the training time (Li et al., 2022 70).
- **FinGPT:** Symbolizes the emergence of FinRL in the sphere of Large Language Models (LLMs) for finance. It provides open-source financial LLMs created with a data-centric approach, automatic data curation and economically viable fine-tuning methods such as the RLSP (Reinforcement Learning with Stock Prices) and LoRA (Liu et al., 2023; Yang et al., 2024 41; Cai, 2025 144).
- **FinRobot:** An open-source AI agent platform based on top of LLMs, for use of various financial applications. It has a multi-layer structure and uses financial Chain-of-Thought (CoT) Reasoning to solve the complex tasks such as equity research (Yang et al. 2024 42; Zhou et al, 2024 43; Yang et al, 2024 41).
- **ElegantRL:** A cornerstone library of ready-made massively parallel DRL algorithm implementations that are commonly used in other components of FinRL.

Qlib, by Microsoft, provides a wider quantitative investment environment with AI orientation throughout the whole procedure chain from data processing to production implementation:

- **End-to-End Pipeline:** Qlib offers datasets for preparation, model training and backtesting to simplify quantitative research.
- **Comprehensive Scope:** It covers such types of quantitative investment as alpha seeking, risk modeling, portfolio optimization, and order execution.
- **Modularity and Extensibility:** Built with loosely coupled components, where the users can customize workflows or can work with specific modules individually.
- **Automation:** Such features as qrun allow the automated execution of the whole research workflow from data building to evaluation reporting.
- **Data Handling:** Comes with a “Quant Dataset Zoo” having built-in datasets (e.g., Alpha158, Alpha360 for US and China markets) and tools to create customized datasets. Supports both offline and data servers’ online deployment modes.
- **Model Zoo:** Provides a wide range of ML and DL models appropriate to quantitative finance: GBDT varieties (XGBoost, LightGBM, CatBoost), sequential

models (LSTM, GRU, Transformer, TCN), attention-oriented models (GATs, TFT), and like TabNet and SFM (Microsoft Qlib Documentation; Yang et al., 2020). It backs supervised learning, market dynamics contemplation, and RL paradigms.

- **Evaluation:** Gives tools for complete analysis of performance analysis, including forecast of signal metrics (IC, return distribution), backtest of portfolio.

Other platforms also add into the scenery, **TradeMaster**, another platform powered by Rl, Shai-am for investment strategies and **QuantBench**, tailored as an industrial-class benchmark to ensure academic research is in line with the industry practices.

The emergence and open access to such platforms as FinRL and Qlib helps a lot in moving the field forward. They offer a set of standard components and settings that reduce the entry threshold for researchers, practitioners, enable replication, and comparison between various strategies. and create an interactive community responsible for the development of financial AI.

Their presence fills the void on the dire need for standardized evaluation and reproducibility of results in an area where proprietary data and methodologies usually complicate matters.

The comparative summary of the FinRL ecosystem and Qlib is given in Table 2.2.

Table 2.2 Comparison of Key Financial AI Platforms

Feature	**FinRL Ecosystem (AI4Finance)**	**Qlib (Microsoft)**	**Primary Focus**	**Key Technologies**	**Open Source**
Main Goal	Democratize Financial AI	AI-Oriented Quantitative Investment Platform	Varies by project (DRL, LLM, Agent)	Full Quant Workflow (ML/DL focus)	Yes
Core Area	Primarily Deep Reinforcement Learning	Full Quant Workflow (Alpha, Risk,	DRL (FinRL), LLM (FinGPT	ML/DL Pipeline	Yes

	(DRL), expanding to LLMs & Agents	Portfolio, Execution) with ML/DL focus	), Agents (FinRobot)		
Environments	Dynamic, Gym-style market environments (FinRL-Meta)	Data management & processing tools, built-in datasets (Alpha158/360)	Market Simulation	Data Handling	Yes
Algorithms	DRL (PPO, DDPG, SAC, etc.), LLMs (FinGPT), Agent Frameworks (FinRobot)	ML/DL (GBDT, LSTM, GRU, Transformer, GATs, TFT, TabNet, etc.), RL support	DRL, LLM, Agents	ML/DL, RL	Yes
Key Projects	FinRL, FinRL-Meta, FinRL-Podracer, FinGPT, FinRobot, ElegantRL	Core Qlib library, Model Zoo, Dataset Zoo, qrun workflow	Specific applications	End-to-end platform	Yes

Strengths	Strong DRL focus, open data/environments (Meta), performance (Podracer), LLM/Agent innovation	Comprehensive quant workflow coverage, extensive model zoo, modularity, benchmarking	DRL, Openness, Community	Workflow, Models, Tools	Yes
References	Liu et al. (2022) [79]; Liu et al. (2020) [110]; Liu et al. (2022) [81]; Li et al. (2022) [70]; Liu et al. (2022) [119]; Liu et al. (2023); Costa & Costa (2023) [123]; Yang et al. (2024) [42]; Zhou et al. (2024) [43]	Yang et al. (2020); Microsoft Qlib Documentation [80]	-	-	-

This table summarizes well the difference but complementary roles of these two big platforms. FinRL focuses on DRL, agent-based modeling, and realistic simulation environment, which is indicative of the attention paid to the learning of dynamic strategies through interaction.

Qlib places upon a wider tool set addressing the conventional quant workflow that encompasses a mind-blowing variety of predictive ML/DL models while paying particular attention to the robust back test and evaluation infra.

Both platforms, considering their open source, standardization-oriented nature, are effective tools in speeding up the research and possibly closing the gap between academic inventions and industry take in the field of quantitative finance.

2.4 Summary

During this review, the theoretical foundations and the cutting edge behind the usage of Agentic and Multimodal AI in quantitative finance, for purposes of wealth and portfolio management, have been crossed.

Classical benchmarks for foundational theories such as Modern Portfolio Theory (MPT) and the Efficient Market Hypothesis (EMH) are provided, and market deviation from perfect efficiency and rationality – deviations that AI may take advantage of or become a victim of, are key pieces of insight that Behavioral Finance and Bounded Rationality offer.

Agency Theory describes the central concerns in trust and alignment in case of financial delegation to independent AI agents. The Theory of Reasoned Action/ Planned Behavior is used to understand what is behind the process of human adaptations to these complex technologies.

The literature documented rapid developments in AI applications in the field of finance. AI is used by Algorithmic trading especially HFT to cover speed and pattern recognition hence the need for the High-Performance Computing (HPC) infrastructure and concerns on market stability.

Deep Reinforcement Learning (DRL) has become an effective tool for the optimization of sequential decisions in trading and portfolio management with specialized frameworks such as FinRL being important to research and development.

AI is transforming wealth management through robo-advisors, hyper-personalization, and utility to increase the productivity of advisors.

The cutting edge of this evolution includes Agentic AI – systems that can reason, and act autonomously, and Multimodal AI – systems that adopt heterogeneous data types, such as text, numbers, and visuals. Agentic approaches such as FinRobot and TradingGPT are intended to automate tricky workflows, whereas multimodal AI models such as FinAgent desire a better comprehension of the market.

Reinforcement Learning from Human Feedback (RLHF) is investigated as a means to consider these highly sophisticated agents and their goal towards nuanced human goals. It is such open-source platforms as FinRL and Qlib that play a crucial role in promoting those advancements because they offer standard tools, environments, and benchmarks.

Regardless of the outstanding progress discussed in many studies demonstrating possibility of AI to deliver strong backtested results, critical problems and research voids remain. Key among these are:

- **Reliability and Robustness:** Consistent AI performance in the dynamic non-stationary financial markets and overcoming the gap between the simulated backtesting and reality of live trading is an important hurdle (Scaletta, 2024 139; Liu, et al., 2024 34; Wang, et al., 2024 82).
- **Interpretability and Trust:** It is commonly the case that the deep learning, DRL, and LLM-based agents produce "black box" models which are difficult for validation, debugging, and obtaining trust from the users or even regulators (Abdullah & Chowdhury, 2023 16; Chen et al., 2025 37; Gu et., 2024 83; Lin & Zhou, 2021 31).
- **Alignment and Control:** Successfully integrating the autonomous agentic systems with complex, and possibly changing human objectives and risk preferences, specifically, through such methods as RLHF needs more research to solve problem(s), such as manipulation of feedback and guarantee safety (Li et al., 2024a 18; Bai et al., 2025 36); Chen et al., 2025 37).
- **Multimodal Integration:** Creating effective and efficient approaches to the integration of information from a variety of data sources, dealing with noise and contradictions, and deciphering insights created thereby is a continuing issue.

(Bhatia et al., 2024 22; Zhang et al., 2024 21).

- **Systemic Risks and Governance:** Learning about the prospective effects at the market level of interacting autonomous AI agents and constructive design of governance structures and regulatory conditions towards reduction of systemic risks are critical but inadequate. (The Global Treasurer, 2025 35; ResearchGate User (2024 9); Azzutti (2024 71); Chen, et al. (2025 37).
- **Practical Deployment:** It is also an engineering challenge to convert the promising research findings into the usable deployed systems: concerning data infrastructure, latency, computation cost, and integrating them to the existing financial systems (Cao et al., 2025 2; Wang et al., 2024 82).

This DBA thesis is geared to bridge some of these information gaps that are critical. By looking at the practical aspect of integrating Agentic and Multimodal AI, particularly for quantitative wealth and portfolio management, specifically; the foregoing research is concerned with breaking out of theoretical discussions or algorithm improvements in a vacuum.

It will explore how these sophisticated AI paradigms may be architected into functional systems, how methods such as RLHF can be practically applied to the alignment in portfolio tasks, and what sort of governance and risk frameworks are needed to deploy them responsibly.

The existing body of literature manifests huge potential yet exposes basic difficulties with respect to trustworthiness and control and the actual practical relevance.

This research has an intent to bridge this life-critical chasm; sharing knowledge regarding how to responsibly and reliably utilize the power of Agentic and Multimodal AI in the aggressive atmosphere of quantitative finance.

CHAPTER III: METHODOLOGY

3.1 Overview of the Research Problem

The modern-day financial world is surrounded with unparalleled complexity, magnanimity and velocity of data traffic, thus leaving traditional quantitative analysis tools worthless (MacKenzie, 2018), (Baldacci, 2021). Quantitative finance is going through a major transformation, with the factor driving it being the pace at which there is an expansion in the data available- market data, unstructured news feeds, social media sentiment, among others with significant progresses in Artificial Intelligence (AI) (Cao et al., 2025), (Liu, 2024).

Although AI methods, especially the deep learning, have shown remarkable potential in such fields as volatility forecasting and improvement of a trading strategy (Gong, 2023), (Sardelich Nascimento, 2023), existing models often find difficulties in dealing with the non-stationarity nature of markets, recognizing complex inter-dependencies between assets and, more importantly, streamlining different informational sources' information (Hou et al., 2021), (Zhang, 2024).

There is a looming gap, though, which goes beyond pure predictive accuracy. The automation of complex trading and portfolio management processes requires more than prediction abilities – it requires planning, reasoning, strategic decision-making, and ongoing adaptation abilities - hall marks of agentic AI (Joshi, 2025), (Yang et al., 2024), (Ding et al., 2024).

Agentic systems can provide AI with the possibility to use it as autonomous systems that can conduct intricate financial operations. Simultaneously, multimodal AI offers a road to combine actionable wisdom from this kaleidoscope of financial data, from numerical series of prices to textual financial reports, news sentiment, and potentially even visual data in the form of price charts. (Zhang et al., 2024), (Huang et al., 2024).

This research tackles the core issue of developing, designing and testing the methodological framework in forming a synergistic integration of Agentic AI and Multimodal AI – for the domain of algorithmic trading in quantitative wealth and portfolio management.

Some of AI systems that are capable of autonomously perceiving market dynamics through various data modalities, thinking about possible trading strategies, automating deal execution, and dynamically investing funds have yet to be conceptualized and evaluated for their effectiveness.

One of the major goals is the investigation of whether integrated systems can be shown to perform better than the existing benchmarks and be more flexible to the changing conditions of the market (Sun et al., 2022), (Zhao, 2024)).

For this investigation, we go deeper into how agentic attributes – like planning, complex use of tool, and reflective learning – synergize, possesses the potential of multimodal data fusion to add value to the decision-making processes in fast-paced as well as medium-speed trading that is aligned to the optimality of portfolio construction and management (Han et al., 2023), (Xu et al., 2025).

Agentic AI, Multimodal AI, and well-established quantitative finance frameworks like the reinforcement learning approaches covered by FinRL (Liu et al., 2022), (Liu et al., 2022), the confluence of all these points at a likely paradigm shift. This evolution goes beyond the predictive modeling to more holistic, autonomous financial decision-making systems.

The traditional quantitative finance was largely dependent on signal processing methods or statistical arbitrage models, bounded by the predefined rules)]. Deep learning then enhanced predictive power, used often as opaque "black boxes" in sterile strategic frameworks.

Reinforcement Learning, in the form of platforms such as FinRL, brought along adaptive strategy learning but often focused on single data modalities, mostly price and volume data (Liu et al., 2022), (Liu et al., 2022).

The emerging Agentic AI, specifically Large Language Model (LLM)-based agents such as FinRobot or AlphaAgent, brings with it higher levels of cognitive functions, such as reasoning, planning, and ability to use external tools. (Yang .et. al 2024, tang .et. al 2025, ding et al., 2024).

Complementing this we have Multimodal AI allowing for the necessary integration of textual data, sentiment analysis, possible visual chart patterns, and baseline numeric market data (Zhang et al., 2024), (Huang et al., 2024).

As a result, the synthesis of these technologies indicates the trend of AI systems that are capable of mimicking and eventually exceeding the powers of human financial analysts to combine multiple information flows, establish complex strategies, and carry out actions independently. This is a fundamental change in the methodology for building and bringing quantitative strategies out.

At the same time, the emergence of such open-source platforms as FinRL (Liu et al., 2022), (Liu et al., 2022), Qlib (Yang et al., 2020), and FinGPT (Liu et al., 2023) is rapidly democratizing access to sophisticated AI tools for financial applications.

This accessibility drives innovations, as observed from the proliferating models and agent frameworks (Yang et al., 2024), (Xiong et al., 2025), (Wang & Hua, 2024), (Sun et al., 2023). This rapid growth, however, simultaneously also presents great challenges regarding standardization, benchmarking, and application of responsible use of such powerful technologies. It is problematic to compare studies' findings when there are no standardized benchmarks of evaluations and the market environments (Liu et al., 2022), (Wang & Hua, 2024).

However, the comparative simplicity of deploying elaborate systems of AI, in particular autonomous agents, casts justified concerns about unknown threats, high stress tolerance of the model, and vital ethical questions. (Chen et al., 2025), (Azzutti, 2024). Thus, although the democratization encouraged by the open-source ventures speeds progress, simultaneous and serious attention to the evolution of effective evaluation techniques, thorough risk control frameworks, is required.

Governing mechanisms that establish stability, trustworthiness, and ethical use of AI in finance (Sun et al., 2022), and effective governance structures (Azzutti, 2024) in the application of AI in the finance sector.

3.2 Operationalization of Theoretical Constructs

To be strictly exploratory about the research problem, specific and operationalized theoretical constructs are required herein to determine them.

Agentic AI: This construct is concerned with AI systems that have properties of goal directedness, autonomous decision making and reactivity to changes in the environment, pro-active nature in the pursuit of objectives and possibly social skills especially when in multi-agent configurations (Joshi 2025) (Krishnan 2025).

In this study, the Agentic AI is operationalized by the conceptual design of the Large Language Model (LLM)-based agents (Ding et al., 2024), (Dong et al., 2025).

These agents are designed in such a way as to have certain capabilities.

- **Planning:** The capability to break down on high-level goals (e.g., maximize risk-adjusted returns subject to client constraints) into a smart chain of executable sub-tasks, including asset screening, signal generation, risk assessment and order execution.

- **Tool Use:** The ability to interact dynamically with external software libraries, APIs, or databases. This includes pulling real-time market data, accessing news feeds, calling quantitative analysis models (e.g., predictive models from the Qlib library (Yang et al., 2020)), and interacting with trading execution platform (simulated or live). Such tool augmentation takes a practical form in examples of frameworks, such as FinRobot (Yang et al., 2024), (Zhou et al., 2024), (Zhang et al., 2024).
- **Memory/Reflection:** By using mechanisms for short-term contextual memory and mechanisms for long-term knowledge accumulation which allows the agent to learn from previous actions, events in the market, and feedback (Li et al., 2023), (Yu et al., 2024). This entails notions such as self- reflection in which an agent reflects on its past actions or prognosis for revision of future conduct (Koa et al., 2024).
- **Adaptability:** The ability to dynamically adjust strategies and parameters following performance feedback, and in view of changing market conditions. It can be improved extensively using methods such as Reinforcement Learning from Human Feedback (RLHF), so that it can match the user's preferences (Xiong et al., 2025), (Samani & Darvishvand, 2024).

Multimodal AI: This construction represents AI systems that can process and interpret various data collected from several different data types or modalities, such as numerical time series, textual documents, sentiment scores associated with text and so on, and possibly graphical representation such as price charts. (Zhang et al., 2024), (Huang et al., 2024).

Operationalization involves:

- Using complex model architectures, which are intended to process diverse inputs. Such examples include fusion transformers or attention mechanisms that dynamically weigh the relevance of different modalities when analyzing. (Zhang et al., 2024), (Hou et al., 2021).
- Combination of sentiment analysis obtained from financial news, social media or analyst reports (possibly by means of special LLMs for finance such as FinGPT (Liu et al., 2023)), or classifiers such as FinLlama (Konstantinidis et al., 2024) with conventional quantitative data such as price and volume (Abdullah & Chowdhury, 2023), (Zhu, 2024), (Zhao & Welsch, 2024)

- Investigating the exploitation of graph Neural Network to deliberately model intricate inter stock connections or market designs using features obtained from both figure and content-based info (Zhang ,2024), (Zhuang et al., 2024), (Fu, 2023).

Algorithmic Trading: It refers to the use of computer programs to conduct the trading strategy in accordance with either pre-defined rules or adaptive learning algorithms (Sun, 2024).

This is operationalized by:

- Creation and conceptualization of the trading strategies with particular focus on certain characteristics of the market or frequencies (intraday momentum), microstructure-discovered potentially HFT opportunities (Han et al., 2023), (Xu et al., 2025), factor-based investing (Zhao et al., 2024), (Xu et al., 2024), (Shi & Song, 2025), Trend following (Zeng et al., 2024), or mean reversion strategies.
- Incorporating trade order execution logic into the agentic architecture, which should take practicalities such as the transaction cost, potential slippage, and market impact, in HFT circumstances (Fang et al., 2023).

Wealth & Portfolio Management: This includes professional asset management of finances to achieve desired investment objectives for individuals or institutions, emphasizing on the critical balance between risk and anticipated returns. (Zhao, 2024).

Operationalization occurs through:

- Using the AI agents designed to make dynamic asset allocation and portfolio rebalancing decisions based on the analysis of the market and learned strategies). ((Sun et al., 2025), (Costa & Costa, 2025), (Cagliero & Fior, 2023/24), (Kayumov, 2023), (Santos et al., 2023)).
- Using risk management approaches, informed and potentially automated by the analysis made by the AI, such as risk-aware reinforcement learning policies. AI-centric estimates of tail risk from multimodal data such as sentiment (Abdullah & Chowdhury, 2023).
- Testing the performance for the conceptualized system using standard metrics on portfolio management such as Sharpe, Sortino Ratios, Maximum Drawdown, and outperforming benchmark's Alpha generation (Liu et al., 2022), (Sun et al., 2023).

Quantitative Finance: This is the broad discipline which makes use of mathematical and statistical procedures on financial problems. It is operationalized in this research by incorporation of existing quantitative models such as predictive models that are available within the Qlib platform (Yang et al., 2020)) as callable tools within the AI agent's framework, and by strictly following the backtesting and performance evaluation practices common in the field (Liu et al., 2022), (Liu & Xia, 2022).

The operationalization of "Agentic AI" into the financial realm indicates the shift from only the automation of mundane tasks, towards the infusion of the rich cognitive functions of planning, reasoning, and learning into the trading and portfolio management systems. This calls for the development of new system architectures and evaluation paradigms.

While conventional algorithmic trading tends to be based on pre-specified rules or to learn policies from data, (Liu et al., 2022), Agentic AI and specifically when fueled by LLMs brings abilities for understanding intended meanings of natural language (e.g., decoding complex client mandates, decomposing high-level goals into actionable steps), etc., strategic use of selecting appropriate analytical tools / data sources and possibly justifications of its decisions (Yang et al., 2024), (Ding et al., 2024), (Koa et al., 2024).

These capabilities can be realized only through architecture that can properly combine LLMs with specialized financial datasets and analytical models, as well as execution platforms. (Zhang et al., 2024), (Xiong et al., 2025), (Samani & Darvishvand, 2024), (Zhou & Mehra, 2025).

Assessment of such systems should go beyond the traditional profit-loss measures to include those evaluations which would measure the quality of the reasoning of the agent, its planning effectiveness and adaptation in the dynamic environments (Chen et al., 2025), (Yu et al., 2024).

Therefore, the operational coding of trading logic is not enough, and we need to be developing cognitive architectures that are particularly fit to finance, and this requires an interdisciplinary approach that brings about AI, financial acumen, and maybe even cognitive science.

In the same way, data aggregation is not enough when operationalizing **"Multimodal AI"** in finance. It requires advanced data fusion approaches and analysis of marginal value over potential noise for each data modality.

Financial markets are a combination of quantitative (prices, volumes, order flow) and qualitative (news reports, central bank statements, social media sentiment, analyst opinions) data. (Huang et al., 2024), (Abdullah & Chowdhury, 2023).

Multimodal AI seeks to pick up and use this informational wealth (Zhang et al., 2024). Including a greater number of data types will not, however, automatically lead to better performance. The noise, outlier's correlations, irregular frequencies of data, reliability relations, and vast amounts of data are severe obstacles (Hou et al., 2021), (Zhao & Welsch, 2024).

Advanced techniques such as attention mechanisms (Zhang et al., 2024), (Hou et al., 2021) are therefore keys to good operationalization. Cross-modal learning strategies and – perhaps – causal inference methodologies (Li et al., 2024) for submitting in an intelligent manner the role of the different modalities, their weighting, and interpretation.

The informational worth of calculating informational use of one modality as opposed to another (such as news sentiment versus technical indicators) may change, based on ruling market regimes or the target asset class under discussion.

Therefore, efficient operationalization of multimodal AI is essentially about intelligent fusion and needs methodologies that can measure the marginal contribution of each data stream and dynamically change the fusion strategy to improve decision-making.

3.3 Research Purpose and Questions

Purpose: The intention of this Doctorate in Business Administration (DBA) research is to suggest, and theoretically analyze, an innovative conceptual framework of the methodology that capitalizes upon synergetic prospects of Agentic AI and Multimodal AI in upgrading algorithmic trading strategy in quantitative wealth and portfolio management.

This research is aimed at singling out the architectural elements, data integration methods, and the learning mechanisms (including the role of Reinforcement Learning from Human Feedback - RLHF) necessary for such an advanced system.

Moreover, it pursues to evaluate, by means of the synthesis of the existing academic literature and research findings, the possible benefits, limitations, and barriers of implementation of this framework as compared to the state-of-the-art quantitative finance.

Research Questions (RQs):

- **RQ1:** How can agentic AI principles (planning, use of tools, reflection, autonomy) be aligned with multimodal AI capabilities (capability of working with numerical, textual, and possibly, other data types), to produce a comprehensive solution for algorithmic trading and portfolio management? (Architectural design and systematic integration – centrality to the project)
- **RQ2:** What would it be like for a certain algorithmic trading strategies to be most inclined to improvement when aided by an Agentic Multimodal AI framework (e.g. high-frequency trading informed strategies, factor-based investing, dynamic asset

allocation)? which performance enhancements (risk-adjusted returns, ability to adapt to the changes of market, and discovery of the new alpha) could potentially be reasonably expected based on the available evidence from the research literature? Concentrates on the area of application and possible impact.

- **RQ3:** What should be the way of getting Reinforcement Learning from Human Feedback (RLHF) into the proposed system to allow the AI agent to make trading and portfolio management decisions in accordance with complicated, possibly subjective, human desires? How can the investment strategies best address these nuanced risk tolerance profiles, ethical investment considerations, and more, specific long-term financial goals that go beyond just quantitative tuning: profit maximization? (Help align the value and constant improvement through human interaction)
- **RQ4:** What are the critical methodological challenges and limitations, and important ethical considerations (e.g., mitigating data bias, promoting interpretability and explainability of the models, eliminating systemic risk potential, achieving effective governance) connected with the emergence and possible application of sophisticated Agentic Multimodal AI systems in the delicate field of quantitative finance? (Takes account of feasibility constraints, inherent risks, and responsible innovation).

These research questions cumulatively answer a major evolution of quantitative finance: the possibility of a potential shift from AI serving as a tool for human quants towards developing AI as a collaborative or even autonomous *agent* operating *within* the quant finance workflow. This evolution by its very nature leads to fundamental questions of the extent of human oversight, control, and trust of these ever so sophisticated systems.

RQ1 goes straight to the consideration of creating such an integrated, possibly autonomous framework. **RQ2** explores the kinds of financial work for which this automation would possibly be most efficient. **RQ3** brings RLHF into the picture admitting that human guidance is indispensable but treating it as feedback to a flexible agent, not strictly controlling it at all moments of time. (Xiong et al., 2025) and (Samani & Darvishvand, 2024). This subtly revises the interaction model between human and AI. Finally, **RQ4** addresses the fundamental risks and moral quandaries that come with allowing AI to have considerable autonomy in high-stakes financial experiences. (Azzutti 2024), (Chen et al. 2025), (Kasirzadeh 2025).

Thus, the research questions are charting a course that examines both the promise and the potential pitfall of growing AI agency in finance and require discussion beyond the narrow purview of performance metrics – to the human-AI relationship and larger systemic repercussions.

Furthermore, the aspect of RLHF answered in RQ3 is a critical indication of the following awareness: just optimizing for traditional quantitative metrics like sharp ratio is not usually good enough in the practical realities of wealth management.

This field is that of navigating complex client needs, unpredictably shifting risk appetites, possible ethical imperatives (such as ESG investing), and these long-term financial goals that do not lend themselves to easy quantification. (Han, et al, 2024), (Abdullah & Chowdhury 2023). Conventional quantitative strategies often optimize for a single, mathematically defined objective function, RLHF is a viable way to bring in subjective and possibly non-quantifiable human feedback directly as a part of the AI's learning process. (Xiong et al., 2025), Samani, & Darvishvand (2024).

This enables the AI agent to learn policies that are in accordance with human preferences and values as opposed to simply optimizing patterns from the historical data (Zhao & Welsch, 2024). Therefore, RQ3 reveals an essential link between the technological power of cutting-edge AI technologies and the real, human-oriented demands of wealth management while alluding to the fact that adoption of these technologies only becomes successful when appropriate mechanisms of value alignment are in place.

3.4 Research Design

Approach: This research uses **conceptual** and **theoretical design methodology**.

It deals with the synthesizing of existing empirical evidence and theoretical constructs present in academic literature. It is also not related to inductive studies in primary empirical research that require obtaining new data.

The essential part of this approach consists of:

- **Framework Development:** Creating a new theoretical architecture of an Agentic Multimodal AI system, specifically designed for quantitative finance in terms of application. This design borrows ideas and components from the established AI and finance literature principles and platforms such as the multi-agent agent architectures, multimodal fusion techniques, reinforcement learning paradigms, and quantitative modeling toolkits characterized by various programming frameworks. (Zhang et al., 2024), (Yang et al., 2024), (Xiong et al., 2025), (Liu et al., 2022), (Liu et al., 2023), (Yang et al., 2020), (Sun et al., 2023).
- **Literature Synthesis:** Systematic review and synthesis of the results from relevant and contemporary research papers (such as those presented in the first query). This synthesis helps to explain the architectural design decisions, operationalize the

main constructs (Section 3.2), help identify the potentially appropriate trading strategies (on the back of RQ2), and identify mechanisms and implications of RLHF (answering RQ3), as well as the major difficulties and limitations (answering RQ4).

- **Methodological Proposition:** Specifying the actual techniques and procedures – from the data integration to the model training procedures (including the notion of conceptual RL and RLHF loops), strategy execution logic, and evaluation metrics that would have to be employed and tested empirically, on his or her future research efforts, the proposed framework.

Proposed System Architecture (Conceptual):

The proposed conceptual architecture includes some major layers, including:

1. **Multimodal Perception Layer:** In charge of consuming and processing various streams of data important for financial markets.
    - **Numerical Data***:* Processing high-frequency (e.g., tick data; order book information) and fewer-frequent data (Han et al., 2023). (for example, daily bars) market data, as well as significant economic indicators. The extraction and processing of the features may exploit the techniques provided in systems such as Qlib (Yang et al., 2020), (Zhang, 2024).
    - **Textual Data:** Processing news feeds, corporate financial reports (for instance, SEC reports), social media mood, and analyst reports. This entails the use of LLMs, that may be fine-tuned versions such as FinGPT (Liu et al., 2023) or specific models like FinLlama (Konstantinidis et al., 2024), for tasks such as summarization. (Yang et al., 2024), sentiment analysis (Abdullah & Chowdhury, 2023), (Mun & Kim, 2025), event detection and extraction [(Lin & Wang, 2024), (Xu et al., 2024) and possibly discovering the causal relationships between the news and market fluctuations (Li et al., 2024).
    - **(Optional) Visual Data:** Perhaps including analysis of technical chart patterns, if research will prove effective techniques of automated analysis.
    - **Fusion Module:** Combining the processed information received through multiple modalities for constructing a single market state representation. This may involve methods such as cross-modal attention mechanisms or graph-based methods that model relationships of entities (such as stocks, sectors) with the help of fused data. (Zhang et al., 2024), (Fu, 2023), (Xu et al., 2025).

2. **Agentic Core (LLM-based):** The core of the cognitive engine of the system.
    - **Planner/Reasoner:** Gets the unified market state representation and high-level goals (e.g., portfolio return targets, risk constraints as defined by a client or policy). Breaks down these goals into an arrangement of actionable steps/questions. (Yang et al. 2024), (Ding et al. 2024).
    - **Tool Use Module:** Selects and calls for the appropriate external tools in a smart manner, relying on the present plan and the market reality. These tools include:
        - Data APIs to retrieve real-time market data or news.
        - A library of Quantitative Models, e.g predictive models from Qlib (Yang et al., 2020), factor mining algorithms (Zhao et al., 2024) (Xu et al., 2024; Shi & Song, 2025; Yu et al., 2023; Ren et al., 2024; Tang et al., 2025) specialized risk assessment modules.
        - Trading Execution Simulators or APIs, possibly based on environments available from FinRL (Liu et al., 2022).
    - **Memory Module:** Does not forget both short- and long-term memory (context for current processes) (accumulated knowledge from previous actions, results, market observations and human feedback) to guide future decision and learning (Li et al 2023), (Yu et al 2024).

3. **Action Layer:** Converts the agent’s choice to real actions.
    - **Strategy Generation:** Prescribes specific trading strategies, such as asset selection, position sizing, and entry/exit timing, based on the agent's reasonings and the outputs of the tools used for its functioning (The Kou et al, 2025; Zhong et al, 2024). This layer might utilize reinforcement learning components to define the optimal strategy parameters (Xiong et al., 2025), (Samani & Darvishvand, 2024).
    - **Portfolio Management Module:** Optimizes the total portfolio allocation in view of the generated strategies, the outlook of the market and predefined constraints (eg: Risk Limits and Diversification needs), (Sun et al., 2025, Zhang et al., 2024, Niu & Li, 2022, Deng et al., 2024). This could likely involve methods from the RL-based portfolio optimization studies (Costa & Costa, 2025), (Sun et al., 2022) or multi-agent cooperation approaches (Sun et al., 2025), (Li et al., 2025).
    - **Execution Module:** Converts high-level trading decisions into simulated (for backtesting) or real trade orders, while taking into consideration models of market microstructure effects, especially for HFT applications), (Han et al., 2023).

4. **Learning & Adaptation Layer:** Facilitates the system to get better as time goes by.
    - **Reinforcement L earning (RL):** Tunes the policies on the process of strategy generation and portfolio management according to the feedback from the market environment. Rewards may be on P&L, risk-adjusted returns. (e.g., Sharpe or Sortino ratio), or another performance metric (Liu et al., 2022), (Yang et al., 2021), (Lee & Moon, 2023). This utilizes such frameworks as FinRL (Liu et al., 2022) and might use sophisticated algorithms including Proximal Policy Optimization (PPO) (Anwar, & Wang, 2022), Deep Q-Networks (DQN)), or Actor-Critic techniques (Vishal et., 2022)
    - **RLHF Module:** Incorporates feedback from human supervisors. This feedback might be supplied in many ways (e.g., ratings of the agent's decisions, comparisons between two or more alternatives of proposed actions, natural language instructions or corrections). This feedback is used to tune the policy of the agent or its reward model, explicitly bringing the action of the agent in line with human goals, risk-tolerance or ethics. (Xiong et al (2025), Samani & Darvishvand (2024), Zhao & Welsch (2025)). This mechanism addresses the vital need for alignment than what is needed to optimize a simple pre-defined reward function.

5. **Innovative Code Concepts:** Methodological novelty does not consist in the designing of totally new algorithms from the very beginning, but rather in a synergistic combination of advanced frameworks. Specifically, the proposal involves:
    - Combining LLM agent frameworks which were inspired by the likes of FinRobot (Yang et al., 2024), (Ding et al., 2024) with specialized financial RL libraries such as FinRL, Quantitative toolkits like Qlib (Yang et al., 2020) and studies (Liu et al., 2022), (Sun et al., 2023) treating mathematical finance education.
    - The key innovation lies in the synergy that comes with the integration:
        - Granting the LLM agent the power to dynamically decide which explicit Qlib prediction model or FinRL-trained strategy module to engage given its dynamic interpretation of the richly distributed multi-modal state of the market at each point of time.
        - Exploiting the internal reasoning and language generation that the LLM possesses that can be used to generate explainable rationales for its trading decisions (Koa et al., 2024). Such explanations can then be

checked by human supervisors as an engaging RLHF loop amid increased transparency and trust.

- Exploration of the development of the multi-agent systems with separate, specialized agents for specific areas of financial workflow such as a high-frequency execution agent, a factor analysis agent, and a macroeconomic sentiment analysis agent. Such agents would have to signal intentions and coordinate actions (Li et al., 2025). It could be organised by a central LLM planner, much like architecture such as FLAG-Trader (Xiong et al., 2025).

One aspect of this construction that plays an essential role in it is the ability not only to summarize the existing work but to synthesize conclusions proposing the new conceptual integrations and possible improvements.

This goes with the pragmatic and innovative nature of DBA thesis. In particular, the methods of synthesis seek to find ways of integrating different advanced AI techniques in a creative manner.

Possible steps for enhancing conceptual codes that may develop from this synthesis are:

Hybrid Agentic-DRL Framework: Conceptualize an architecture in which a high-level reasoning, planning, and multimodal context understanding agent in the form of an LLM-based agent can be implemented. whereas a Deep Reinforcement Learning (DRL) agent optimizes the fine-grained execution policy.

The LLM agent can rely on reflection and tools used to identify strategic targets or restrictions, which is reduced in the state/reward setup of the DRL agent (like PPO 8 or in DDPG). This form of hybrid has the objective of blending the strategic calculation of LLMs with the provided optimization potential of DRL for trading execution.

1. **RLHF-Enhanced Portfolio Optimization:** Explain an algorithm for RLHF implementation into a portfolio optimization loop, which might be done in federations such as FinRL or Qlib. The RL agent suggests portfolio adjustments but not the optimization for some arbitrary quantitative metric like Sharpe Ratio, but rather refining the policy based on a reward signal or validated by a model that is trained for the human preferences regarding the risk-return trade-offs, diversification levels, ethics (ESG factors) or the fit to a specific investment mandate.

   This enables an alignment between the strategic approaches available to portfolio investors and investor profiles or institutional aims, taking the approach out of the realm of mathematically optimized solutions. Such methods as Direct Preference Optimization (DPO) can reduce this implementation to something simple.

This research design connects theoretical progressions in AI, including complex agentic models and multimodal fusion methods, with the pragmatic needs of the field of quantitative finance, domain specific requirements of quantitative finance, which involve strong risk management, accurate strategy implementation and understanding of the regulatory landscapes (Azzutti, 2024).

AI research often focuses on building general capabilities, such as the reasoning strength of LLMs (Ding et al., 2024), the optimization efficiency of RL algorithm (Anwar & Wang, 2022). Quantitative finance, on the other hand, requires the use of these skills with an extremely narrow and complex set of issues, where this set is marked by a unique structure of data, severe variables restrictions, and multitudinous goals (Boyd et al., 2024).

The proposed design explicitly involves refined financial platforms (FinRL, Qlib), actual types of data (market information, textual news), and main financial operations (trading execution, portfolio management). It also acknowledges practical necessities, for example, implementation of RLHF for the purpose of alignment with human preferences (Xiong et al., 2025) and inclusion of considerations for the risk mitigation (Chen et al., 2025), (Winkel & Strauß, 2023).

Thus, this design can be thought of as a conceptual blueprint to map the cutting-edge AI theory into a potentially viable financial application, with the integration points and the domain-specific adaptation well separated and highlighted.

Its conceptual nature of this research design allows examining the highly advanced, even future-state AI systems without having to be limited by the challenges of empirical implementation immediately. This strategy encourages future-oriented development of theory that is suitable for the strategic orientation of a DBA program.

Empirically accomplishing the entire proposed architecture would be a rather large challenge, requiring many resources for data engineering, lots of training models, and an intricate infrastructure build-out. (Zhou and Mehra, 2025), (Sun et al., 2023).

A conceptual design, however, enables the research to solely focus on the basic principles of integration, possible synergies between distinct AI paradigms, and opportunities and challenges to come out.

This goes well with the applied and strategic orientation of DBA research that strives to provide framework development, strategic implication, and managerial insight more than pure technical novelty. It helps integrate insights from a large variety of state-of-the-art research papers (Zhang et al., 2024, Yang et al., 2024, Xiong et al., 2025, Liu et al., 2023, Tang et al., 2025, Ding et al., 2024) in order to build a consistent vision of how these advanced AI paradigms could transform the future of quantitative finance, regardless of the fact that some of those components now stand on the frontiers of science.

3.5 Population and Sample (Reframed as Data Source Selection)

The conventional understanding of the term "population" and the term "sample" acquires a new meaning in the framework of the present conceptual study that concerns an AI system and comes to refer to the universe of those data sources that are relevant and the strategy for selecting data for hypothetical model training and evaluation.

Population: The "population" allows the immense ocean of potentially relevant data for quantitative finance, which could be theoretically supported with the proposing Agentic Multimodal AI system. This is, but not limited to:

- **Global Financial Market Data:** Across such different asset classes as equities, fixed-income, futures, and options, foreign exchange (forex), and cryptocurrencies (Guarino et al., 2022), (Li et al., 2024). This data is available at a range of frequencies including high-frequency tick data important for HFT strategies through to daily or weekly data of use for more-term portfolio management.
- **Corporate Financial Data:** Such as quarterly and annual financial statements (balances sheets, income statements, cash flow statements) and other required disclosures required to be filed with regulatory agencies (e.g., SEC EDGAR).
- **Macroeconomic Indicators:** Statistics published by government and international bodies on inflation, jobs, GDP growth, rates interest among others.
- **News and Media:** Immediate news articles, press releases, and broadcasts of respectable financial news agencies (e.g., Bloomberg, Reuters, Dow Jones).
- **Social media and Web Data:** Information gotten from such platforms as Twitter, StockTwits, Reddit and financial forums, must be processed, verified and passed through advanced sentiment analysis. (Abdullah & Chowdhury, 2023), (Lin & Wang, 2024).
- **Analyst Reports and Research:** Publications by investment banks, research places, and academic institutions comprising of forecasts, ratings, and market commentary (Zhou et al., 2024).

Sampling Strategy (Conceptual): In any hypothetical implementation and testing of the proposed framework, a **purposive sampling strategy** would be required. This means consciously selecting the best data sources and types considered most appropriate to the given trading strategies and portfolio management goals which the AI system is trying to address. The things to consider for this selection process are as follows:

- **Data Types:** Guaranteeing that there was incorporation of multiple modalities as inclined in the framework design. That would require choosing representative numerical time series (prices, volumes, indicators calculated on their basis), appropriate text data (news articles, financial reports), and possibly organized fundamental data. The framework is developed to make use of multimodal inputs (Zhang et al., 2024).
- **Markets and Assets:** Concentrating on financial markets (major US equity indices such as S&P 500, highly liquid currency pairs, such as EUR/USD, or certain commodity futures) famous for availability and quality of data, and adequate liquidity for the sought trading frequencies to be maintained. These are what the existing financial AI research platforms tend to be based around (Liu et al., 2022), (Yang et al., 2020).
- **Time Period:** Choosing a training and testing period in the past, which would cover several different market regimes (bull markets, bear markets, high or low volatility periods). This is very important to ensure that robust and adaptive models are developed (Liu et al., 2024), (Liu et al., 2022). Ideally, the data set should capture to the most recent periods conceivable to reflect the dynamics of the current market and their effects of recent geopolitical or economic events.
- **Data Sources and Quality:** Ordering data from established reputable financial data providers (e.g., Refinitiv, Bloomberg, FactSet) sources of highly accurate and reliable data. Using the curated datasets that are made available via research environments such as FinRL-Meta (Liu et al., 2024), (Liu et al., 2022), (Liu et al., 2022) or financial datasets on such platforms as Hugging Face, particularly the ones (designed for financial LLMs (Liu et al., 2023)) could also be considered. News data would need access through strong APIs. To salvage data paucity especially regarding rare market events or for assessing the robustness of a model when under stress, REAM The possible application of synthetically created financial data (e.g., application of techniques such as those inspired by FinGPT for the generation of synthetic data (Cai, 2025) or diffusion models (Gao et al., 2024) can be an addition alongside the real historical data. Advanced market simulation engines (Li et al., 2025) also may become good environments for training and testing, but only on condition that their realism is practically validated.

The shift in paradigm towards multimodal AI radically enlarges the “population” of appropriate financial data well outside the scope of traditional price and volume series.

This expansion offers massive opportunities for exploiting unexploited market signals as well as formidable data engineering and integration problems.

Whereas conventional quantitative finance was highly dependent on structured numerical market data)], multimodal AI explicitly seeks to incorporate unstructured text and sentiment indicators and possibly other data forms. (Zhang et al., 2024), (Huang et al., 2024).

Effective exploitation of this polyglot data topography requires the creation and tuning of powerful data pipelines that can scrape, clean, congruently timestamp (tune) time and accurately process dissimilar data formats. (Guo et al., 2023), (Damoun & Seba, 2025).

Methods including LLM-driven information extraction (Liu et al., 2023), advanced sentiment analysis (Konstantinidis et al., 2024), and automatic event detection (Xu et al., 2024) become essential pre-processing steps before transmission of data into core analytical models. Therefore, the power of the practical effectiveness of the present agentic multimodal framework is highly dependent on overcoming these complex data logistics hurdles.

This necessitates significant investment in data infrastructure, data science expertise, and high-quality management process around data, which is a critical implementation choking point not only a weakness but a potential basis of a great competitive advantage for institutions that overcome them.

Moreover, the planned addition of synthetic data and use of simulation environments provide advantages for enhancing sparse real data and the training and testing of agents to the extent wanted to present a non-trivial risk of deviation from the intricacies and subtlety of real-world market dynamics.

Careful validation is paramount. Real financial data are sometimes marked by noise, missing values, regime shifts as well as weak dependencies that are hard to be captured perfectly in synthetic datasets or other simulations (Liu et al., 2024), (Cai, 2025).

Synthetic data generation mechanisms (Cai, 2025), (Gao et al., 2024) and market simulators (Li et al., 2025), (Liu et al., 2022) give controlled environments useful for training RL agents and scaled testing of strategy differences. However, the artificial environments may not suffice to capture core features of the real-world aspects like precise market impact of large trades, random “black swan” events, detailed market microstructure.

An over-confidence on simulated or synthetic data can develop such models that show super performance in back-tests but significantly under-performs in real trading scenarios—a phenomenon well known as “sim-to-real” gap. (Chen et al.; 2025, Liu & Xia 2022).

Thus, despite the usefulness of synthetic data and simulations at the development process, their application should be selective and ruthlessly supported with vigorous testing on hard historical data through a wide range of market conditions and, hopefully, ratified by paper trading (Liu & Xia 2022) before presenting any live deployment. Permanent validation of the variety and boundaries of the SIM-environment spaces (Li et al., 2025), and synthetic data generation processes (Cai, 2025) is a necessity.

3.6 Participant Selection (Reframed as Model/Algorithm Selection Criteria)

To fit the focus of the research on AI design, we redefine participant selection as the process of choosing algorithms and models and the software elements that form the idea behind the framework. The goal of selection is to include the most recent approaches, ensure they fit the study’s research questions, and make sure they work together within the proposed architecture.

Selection Criteria:

- **Agentic Capabilities:** Harmonizing preferred LLMs based on how well they perform reasoning, planning, and using various tools, under challenging situations. We prefer models that might be improved with information from the financial domain, following FinGPT (Liu et al., 2023), and FinRobot (Yang et al., 2024), (Zhou et al., 2024) ideas or those solutions use general-purpose LLMs to perform various financial tasks (Ding et al., 2024), (Dong et al., 2025), (Kong et al., 2024)). Trainers are also interested in models with strong abilities to recollect past experiences (Li et al., 2023), (Yu et al., 2024), (Koa et al., 2024).
- **Multimodal Processing:** Finding model designs that are effective at mixing numerical and text information. We can use multi-modal transformers (described by Zhang et al. in 2024) or other structures that use attention-based combination methods. I prefer models that perform well in NLP jobs related to finance, for example, perform well at both sentiment analysis and information extraction.
- **Quantitative Analysis & Prediction:** Integrating various reliable quantitative models into the agentic core of the plan. It means making use of pre-built models provided in platforms like Qlib, like LightGBM, XGBoost, CatBoost, LSTM, GRU, and Transformers. Images created by Temporal Fusion Transformers (TFT) (Yang et al., 2020) and additional recently released models evaluated in the same ecosystem (e.g., HIST and TCTS referred to within Qlib). Models built for the sole

purpose of anticipating directions in stock trading. Other paramount candidates include 2024 papers by Zeng et al., Liu et al., Hou et al., Cao et al., and Liu et al., as well as 2023 papers written by Xiang et al., 2025 papers by Szydłowski and Chudziak, and 2025 papers by Zhou et al. Additionally, it's important to use algorithms that are developed for finding alpha signals (factor mining).

- **Reinforcement Learning:** Strong reinforcement learning methods proved sufficient for financial challenges that both have noisy conditions and act on many variables. Among these, PPO (Anwar & Wang, 2022) and Deep Deterministic Policy Gradient (DDPG) are considered, along with new methods like Soft Actor-Critic (SAC) (Yang et al., 2021). Multi-agent RL (MARL) is applicable for optimizing portfolios and completing complex tasks such as for (Sun et al., 2025), (Li et al., 2025), (Fang et al., 2023). Implementations included in widely used libraries such as FinRL (Liu et al., 2022) as well as high-performance libraries such as ElegantRL (ref: 7 in user query) are preferred. Risk awareness is the focus when special algorithms are designed (Winkel & Strauß, 2023). For finance, working with long, uninterrupted activity spaces are especially significant (Montazeri et al., 2025).
- **Portfolio Optimization:** Both tried and true portfolio construction systems like Markowitz's mean-variance optimization, and newer data-based approaches from reinforcement learning should be utilized. (Costa and Costa (2025), Zhang et al. (2024), Niu and Li (2022), (Aprendizaje Profundo y Gestión de Carteras (2024), (Cagliero and Fior 2023 or 2024), (Sun et al. 2024), (Deng et al. (2024).
- **Benchmarking:** Picking models and algorithms that have been tested before and written about in recognized platforms (Liu et al., 2022), (Wang and Hua, 2024), (Sun et al., 2023), It helps forecast expected results and makes it possible to compare them.

The process allows for preference towards using the best type of architectural philosophy. The framework suggests the Agentic AI core will oversee overseeing and coordinating all the AI-based activities. Thanks to this core, developers can match different models and algorithms to the various sub-tasks essential to their field. Every AI model is valuable due to its unique features.

LLMs are skilled at reading in natural language, logic, and planning (Ding et al., 2024). GBDTs and Transformers are often able to accurately forecast based on data that follows a pattern (Yang et al., 2020). RL algorithms help machines properly respond to changes in the environment (Liu et al., 2022).

An AI that focuses on all these tasks at once could perform poorly on specific functions. This is addressed by giving the LLM agent the flexibility to use information from all the team members when needed. Some of these methods include using Qlib predictors for forecasts, checking up-to-date sentiment analysis results, relying on RL policy networks for making trades, and so on, following the tool-enhanced agent approach (Zhang et al., 2024).

This approach to small pieces is very useful in many ways. It makes it easier to adapt the system because you can replace each outdated part with a newer, better model. Therefore, the focus in selecting a design is on a modular and orchestrated system, since it follows the most realistic strategy for creating AI systems from parts that each play their specialized part.

Yet, relying on open-source resources such as Qlib, FinRL, and Hugging Face helps reproduce the results and move forward in research and development faster. As a result, any prejudices, issues, or underlying assumptions that the platforms have are also passed along to the created framework.

With well-known platforms available (Yang et al., 2020), (Liu et al., 2022), (Liu et al., 2023), one can use directly implemented models, pre-trained model parameters (mainly for LLMs), and even find related benchmark datasets and the evaluation measures (Liu et al., 2022), (Wang & Hua, 2024), (Sun et al., 2023).

It is important to note that AI models trained on large datasets sometimes display biases found in that data (Chen, X., et al. 2025, and Zhao, G., & Welsch, F. 2024).

Moreover, the platform designs, the settings used for the environment (such as transaction costs), and specific list of features supported by systems like Qlib impact the results. Naturally LLMs, (Yang et al., 2020) puts forward certain restrictions or perspectives and they are reported to have problems with hallucination and bias (Chen et al., 2025; Croom, 2024).

Still, even though using open-source resources is efficient, the approach must clearly identify how these challenges could affect the outcome and act to confirm those results using suitable adjustments and checks. They also focus on rigorous steps for validation, domain-specific fine-tuning (Liu et al., 2023) reducing unfairness from financial decision-making systems.

3.7 Instrumentation

The tools, frameworks, algorithms, and the metrics mentioned are the main instruments used in this sort of research used to make it possible to design, implement, and measure the performance of an Agentic Multimodal AI system.

AI Development Frameworks and Platforms:

- **Agent Frameworks:** The design for the concept uses theories and models from existing libraries in the LLM field (e.g., LangChain and AutoGen) and adds financial concepts. Different studies draw inspiration from specialized financial agent platforms demonstrated in research, they include FinRobot (Yang et al., 2024), (Zhou et al., 2024), TradingGPT (Li et al., 2023), FinMem (Yu et al., 2024), HedgeAgents (Li et al., 2025), and TradingAgents (Xiao et al., 2025). They play the role of blueprints for integrating LLMs, building AI models, calling related tools through APIs, and manages memory in the system.
- **Quantitative Platform:** Microsoft Qlib (Yang et al., 2020) is considered a main element for instrumentation. It has the function of creating data processing pathways and supports engineering of features. Below are a few examples that show this in action: using existing factor libraries such as Alpha158 and Alpha360, giving access to popular quantitative prediction models (like GBDTs, LSTM, Transformer, TFT), and including a facility for back testing models.
- **Reinforcement Learning Library:** The FinRL library and its related environment and benchmark ecosystem (FinRL-Meta) (Liu et al., 2024), (Liu et al., 2022), (Liu et al., 2022), a key part of the approach is the use of FinRL-Podracer for high-performance computing (Li et al., 2022), are central to the methodology. Using them, people can build markets that mimic the real world, practice trading and portfolio management on agents, and evaluate the strategies developed. In addition, strong integration in systems like TradeMaster and spotting the benefits of libraries such as ElegantRL in user query consists of complementary tools for RL development.
- **LLM Tools:** The help of frameworks and libraries is important for perfecting, releasing, and unifying large language models, especially those meant for the finance field. For example, you can build on FinGPT (Liu et al., 2023) or integrate different models by using standard interfaces like Hugging Face Transformers.

Key Algorithms and Models: The instrumentation explained in Section 3.6 consists of some LLMs, architectures for merging different forms of data, predictive models taken from Qlib, and a variety of reinforcement learning algorithms used in FinRL. Studies known as: Zhao et al. (2024), Xu et al. (2024), Shi & Song (2025), Yu et al. (2023), Ren et al. (2024), Tang et al. (2025), Xu et al. (2025), Shin et al. (2024), focus on specialized algorithms and new approaches for improving portfolios (Costa & Costa, 2025), (Niu & Li, 2022), (Deng et al., 2024).

Performance Evaluation Metrics: The paper recommends a group of standard quantitative metrics for testing and measuring the results of trading strategies and managed portfolios.

- Return-based Metrics: Cumulative Return, Annualized Return, topped with the simple Calmar Ratio.
- Risk-adjusted Return Metrics: Sharpe Ratio, Sortino Ratio (also known as downside deviation), and Omega Ratio as shown in Liu et al. (2022) and Sun et al. (2023)].
- Risk Metrics: Some important risk factors used today include Maximum Drawdown (MDD), Annualized Volatility (AV), Value-at-Risk (VaR), and Conditional Value-at-Risk (CvaR).
- Trading-specific Metrics: Trading activity, how often a trade is made, the outcome of those trades, how profitable each trade is on average, and approximate slippage. They include Transaction frequency, Win Rate, Profit Factor, Average Holding Period, Slippage estimates.
- Benchmarking: Look at how a strategy does against popular indices (S&P 500, MSCI World) and lower-complexity strategies (buy-and-hold, traditional factor models) as well. Setting benchmarks based on QuantBench (Wang & Hua, 2024) is a good idea to consider

RLHF Interface (Conceptual): Using the interface's conceptual design, human supervisors can give feedback to the RL agent. As a result, people can rate the overall quality of what the agent decides, choose between the various actions the agent recommends, and direct the agent with natural language guidance. (Xiong et al. 2025) and (Zhao and Welsch 2024) make the argument.

Simulation Environment: The approach is based on running strategies in virtual markets to test and train RL agents. Liu's team made the FinRL-Meta environments (Liu et al., 2024), (Liu et al., 2022), (Liu et al., 2022) provide a strong foundation. More sophisticated, potentially agent-based market simulators (Li et al., 2025) could be used if they prove to be accurate.

Financial Indicators & Features:

- **Technical Indicators:** Some of the commonly used technical indicators include Moving Averages, Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), Bollinger Bands, Average True Range (ATR). These features will be created using raw data from prices and volumes. Qlib also has features for generating many such characteristics.
- **Sentiment Scores:** To find sentiment out, the system will process news articles and reports. You could use a pre-trained sentiment analysis model like FinBERT or even create custom prompts for the core LLM to do this.
- **Fundamental Data (Optional):** Based on the plan's attention and time frame, important data ratios (like P/E, P/B, ROE) calculated from the company financial reports. They may also exist as features, though it is less common for pure algorithmic trading since they play a bigger role in wealth management.

Evaluation Metrics:

A main set of evaluation metrics should cover Cumulative Return (CR), Annualized Return (AR), Sharpe Ratio (SR), Sortino Ratio, Maximum Drawdown (MDD), Calmar Ratio, and potentially Conditional Value-at-Risk (CVaR) and Information Ratio (IR). TradeMaster provides built-in tools to figure out factors such as CVaR and IR.

Innovative Algorithmic Component (Conceptual Code/Logic):

Providing these specific implementation details within the instrumentation section is crucial. It moves beyond simply stating the use of "Agentic AI" or "Multimodal AI" and clarifies *how* these concepts are translated into concrete algorithms, model architectures, and data flows within the research framework, thereby enhancing the methodological rigor required for a DBA thesis.

Table 3.1: Key AI Frameworks and Platforms in the Proposed Methodology

Framework/Platform	**Primary Role in Methodology**	**Key Features/Models Utilized**	**Relevant Citations**

Agent Frameworks (Conceptual, based on FinRobot, etc.)	Core agent orchestration, planning, tool use, reasoning	LLM integration, prompt engineering, API calls, memory management, multi-agent coordination	[(Zhang et al., 2024), (Yang et al., 2024), (Zhou et al., 2024), (Li et al., 2023), (Ding et al., 2024), (Yu et al., 2024), (Li et al., 2025), (Dong et al., 2025), (Xiao et al., 2025)]
Qlib	Quantitative analysis tool, prediction models	Data processing, feature extraction (Alpha factors), model zoo (GBDT, LSTM, Transformer), backtesting	[(Yang et al., 2020), (Wang & Hua, 2024), (Zhang, 2024)]
FinRL & Ecosystem	RL for trading/portfolio management, simulation	Market environments, DRL algorithms (PPO, DQN, SAC etc.), portfolio optimization agents, benchmarking	[(Liu et al., 2022), (Liu et al., 2024), (Liu et al., 2022), (Liu et al., 2022), (Li et al., 2022), (Liu et al., 2022), (Costa & Costa, 2025), (Yang et al., 2021), (Sun, 2024)]

FinGPT / Financial LLMs	Text processing, sentiment analysis, knowledge grounding	Fine-tuned LLMs, financial data integration, sentiment classification, summarization, Q&A	[(Liu et al., 2023), (Cai, 2025), (Konstantinidis et al., 2024), (Sinha et al., 2025), (Zhao & Welsch, 2024), (Yang et al., 2024), (Mun & Kim, 2025)]
TradeMaster	Holistic RL trading platform	Integrated data pipelines, training workflows, backtesting for RL strategies	)]

Choosing the instrumentation is a key element when implementing this advanced system(see Table 3.1). It is built using different advanced AI technologies that are carefully combined. It means that to achieve great outcomes, a high level of MLOps expertise is necessary, in addition to the typical knowledge and skills in quantitative finance or core AI models. This approach mixes LLM agents, Qlib libraries, platforms like FinRL, and pipelines to take in real-time data.

For these systems to work well together, we need advanced solutions to manage what software is used, guarantee smooth transfer of data, and constantly monitor how well the systems are performing and if data is drifting. Ensuring that the model is often updated, retrained, and that the system can still function properly (Guo et al., 2023), (Zhou & Mehra, 2025).

As a result, workers should have cloud computing (AWS, Azure, GCP) skills, API knowledge, skills with containerization (Docker and Kubernetes), as well as knowledge about orchestrating workflows (for example, Airflow or Kubeflow).

With monitoring frameworks, the difference between designing and deploying a model can be addressed better by implementing advanced AI in the world of finance. It involves software engineering and MLOps as much as research in AI.

Selecting the best performance metrics is important and should go further than only standard financial KPIs. While metrics like Sharpe ratio and maximum drawdown are essential for evaluating financial outcomes, they help measure the results in an agentic system, but do not accurately reflect the way agents make decisions or learn. This is especially true for RLHF (Chen et al., 2025).

It is important to analyze the agent by looking at its planning, as well as the way it uses and handles its tools, how it has coped with various markets or unpredictable events and the level of person alignment attained through the RLHF method.

It is also possible to include the rates at which planned sub-goals are finished, measures of how tools are being used by the agents, or the scores indicating how well the system performs when different conditions are applied in simulations.

(Liu et al., 2022) uses measures from the RLHF process that compare the agent's actions with human actions recorded through learning and also, could cover metrics that check the quality and accuracy of explanations (Koa et al., 2024).

Consequently, evaluating the industry should focus on new agent-oriented straight-through processing metrics as well as the standard financial key performance indicators to analyze the entire system to see if it works reliably and can be trusted.

3.8 Data Collection Procedures (Conceptual)

This chapter explains how to collect, organize, and prepare the required multimodal data for an anticipated empirical implementation of the AI system.

Data Acquisition Strategy:

- **Market Data:** It is essential to obtain complete information about market history. Raw data from the market is provided at various rates (for example, tick or minute-level for developing and testing HFT parts (Han et al., 2023), Portfolio management relies on frequent checking (daily or hourly) for some things, and lower-frequency strategies may do so for other things using equities, ETFs, and futures. It is important to get your data from providers who are regarded as reliable and trustworthy (e.g., Refinitiv Eikon, Bloomberg Terminal API) or from those who offer historical data (e.g., AlgoSeek, TickData). Curated datasets in platforms such as FinRL-Meta (Liu et al., 2024), (Liu et al., 2022), (Liu et al., 2022) can be great starting points or benchmarks for your projects.

- **Textual Data:** A range of textual data is important for multimodal methods. It requires getting financial news articles in real-time and from history via APIs from top news wires. For example, a business can rely on Bloomberg News, Reuters News, or Dow Jones Newswires. Official reports (such as 8-K, 10-Q, 10-K) must be obtained from sources like the SEC EDGAR database for proper integration with fundamental analysis.

  It might be helpful to extract data curated from financial social media such as information from Twitter or StockTwits platforms. However, since there is a possibility of unreliable information and noise, extra thorough vetting and confirmation is necessary (Abdullah & Chowdhury, 2023), (Lin & Wang, 2024). It is possible to develop financial production-ready language models in a short time by processing the pre-created datasets used for training and fine-tuning LLMs (Liu et al., 2023).
- **Fundamental Data:** Accessing organized basic company data such as past earnings per share, development in revenue, their price-to-earnings ratio, and book value from well-regarded financial data sources. For value/growth-focused strategies, systems such as FactSet and S&P Capital IQ are very important. NLP tools can help pull out important data that is included in corporate documents.
- **Alternative Data (Optional):** If alternative data will help in certain strategies, it may warrant including them. Some examples are watching retail activity via satellite imagery, noticing delays in the supply chain, or studying credit card use. Anyway, collecting this type of data adds extra difficulties to the acquisition, processing, and modeling process.

Data Pre-processing and Alignment: The data gathered from different places often needs to be processed a lot before the AI models can use it.

- **Cleaning:** It is important to handle missing parts of the data (by imputing or removing them), identify and treat market data that appears unlike the others, and alter the data in case stock split corrections are needed. The process of cleaning data from texts starts with removing things like HTML tags, advertising sections, and standard content, and follows with lowercasing, stemming, or lemmatization.

- **Feature Engineering:** Picking out features from the unprocessed data. When dealing with market data, you may need to use several technical indicators such as moving averages, RSI, MACD and potentially quantitative factors using libraries like TA-Lib, or leverage Qlib to automatically create important factors (Yang et al., 2020).
  In the case of text-based data, we can apply NLP algorithms, such as different financial LLMs, to look for sentiment and note important companies and people in the data (Liu et al., 2023; Konstantinidis et al., 2024; Yang et al., 2024). Analyze large amounts of text, and pinpoint specific happenings, such as announcements about earnings and mergers. Graphs that use correlations and industry sector groupings to show how stocks regard each other. It is also possible to construct news co-occurrence using various methods (Zhang, 2024), (Zhuang et al., 2024), (Fu, 2023), (Xu et al., 2025).
- **Time Alignment:** This step is very important and challenging. All the data should be precisely stored with a common, high-resolution timestamp (e.g., Coordinated Universal Time - UTC). You must treat the data properly since it might be missing market ticks at times (e.g., news articles published between market closes versus continuous market ticks). There should be strong procedures put in place to make certain lookahead bias avoided, as information must be included at the correct observation times. Packages that are designed for handling financial time series could be especially useful (Guo et al., 2023).
- **Structuring for Models:** Making sure the data is in the prescribed form for each of the AI models used in the framework. This can mean building sequences for LSTMs or Transformers, preparing tabular data for GBDTs, making multimodal data using fused tensors for fusion models, or building state descriptors for RL agents.
- **Data Splitting:** The evaluation of a model should be strict, and the data must be separated into several distinct and non-repeating periods.
- **Training Set:** Trained with both supervised models and RL agents to update their important parameters.
- **Validation Set:** Can be used for finding the right settings for hyperparameters and for stopping training early to reduce overfitting.
- **Out-of-Sample Test Set:** Reserved and not utilized during the training, but just for testing after the training is finished. It is the best way to see how the model is likely to do with unseen data. Many researchers use walk-forward validation or k-fold cross-validation while making sure proceed chronologically in the data, it would be more suitable to use specific methods instead of random shuffling to uphold the order that exists in financial data (Liu et al., 2022), (Liu & Xia, 2022).

Gathering and preprocessing the data for an agentic, multimodal financial AI system takes a lot more time, skill, and resources than usual quantitative models do depending only on a single kind of data. It can create a significant problem when trying to apply Ubiquitous Computing. To do this, it is necessary to acquire, possibly license, join, and control different, often varying, data sources, including numbers, texts, and other types of information.

It makes use of advanced NLP approaches, including preparing the textual inputs well is done by fine-tuning large language models, though this work is not directly covering the task at hand (Liu et al., 2023), (Konstantinidis et al., 2024). It is very difficult and very important to get asynchronous and heterogeneous data streams to be aligned precisely in time so as not to allow subtle lookahead bias to ruin the outcomes of backtesting (Guo et al., 2023).

In addition, feature engineering becomes naturally multimodal, so it may be necessary to use advanced data-fusion methods even before feeding data into the main analysis models. Therefore, the success of the approach is largely determined by how well we can address these huge data logistics challenges.

This calls for dedicated investment in building reliable data systems, expertise in financial data as well as data analytics, and careful data quality management. Also, if the data used to train the AI is not accurate, complete, and free from bias, it can greatly influence how confidently and fairly the AI decision is taken. Therefore, it is necessary to properly check sources and monitor data quality during the entire system's life.

AI models gather knowledge about patterns, related information, and biases that are in the data they are trained on (Chen et al., 2025). If the market data has unseen mistakes, issues with news feeds are regular, or sentiment analysis tools are not up to the task, the results may still be inaccurate (Zhao & Welsch, 2024), there will be imperfections in the decisions taken by the AI agent.

Bias in news and social media may guide the agent to create portfolios that are not well-balanced. These large models can pick up and possibly expand the existing biases found in the data they are trained on (Croom, 2024).

In other words, collecting data is not something you only do at the start. One must regularly check the data, look critically at the credibility of the sources, and check for biases. In the same way, it might include the introduction of certain methods to detect and deal with bias throughout the entire data pipeline and during the modeling steps.

3.9 Data Analysis (Conceptual)

The section explains how analysis is done to train the various AI parts and appraise the work of the integrated Agentic Multimodal framework and analyzing the results based on the questions asked in the research.

Model Training Procedures:

- **Supervised Learning Models:** To train predictive models from Qlib or custom sentiment classifiers accurately, standard supervised learning needs to be used. Experimenters feed the data used for training to the models, optimize their parameters to get a lower error, and evaluate results using specific loss functions (e.g., mean squared error for regression, cross-entropy for classification). By using the separate dataset, hyperparameters (such as the model complexity, regularization strength) are changed to get the best generalization performance (Yang et al., 2020).
- **LLM Fine-tuning:** If you use a large language model (FinGPT), it might have to be fine-tuned using a set of data from the desired domain. They process big amounts of news, analysis documents, and corporate documents to improve their education and language for the financial industry (Liu et al., 2023). It might be necessary to use datasets made for following instructions, which focus on training the LLM in relevant financial tasks for the agentic framework such as deciding how to invest, using financial resources in the right way, or providing sensible explanations for each action (Yang et al., 2024), (Koa et al., 2024).
- **Reinforcement Learning Agents:** RL agents in charge of trading actions or portfolio adjustments would learn in test markets before being put into practice such as the ones from FinRL-Meta (Liu et al., 2022) and (Liu et al., 2022) or those custom-built simulators. Using RL algorithms such as PPO, SAC, or DDPG, the agent learns how to act in the environment by receiving tips on proper actions from the environment. It is essential that the reward function is designed to ensure the company makes a profit, while keeping risks in mind. Using rewards that reflect risk-adjusted returns like the Sharpe or Sortino ratio is a viable option, as is setting penalties for going beyond certain drawdown or volatility goals. (Liu et al., 2022), (Sun et al., 2022), (Winkel & Strauß, 2023). When dealing with multi-agent systems, the training must be tailored to handle how agents interact and cooperate or compete with one another (Sun et al., 2025), (Li et al., 2025).
- **RLHF Integration:** This method requires putting in place a way to collect human feedback on the actions of the agent. For example, some of these outcomes include proposed trades, transformation of the portfolio, and explanations produced by the

software. The information collected from training the critic helps to either form a new reward model that follows human preferences or change the agent's policy using preference-based RL. The idea is to get the agent to act more in step with human beliefs about risk, ethics, or strategy, though the environment's basic reward may not represent these exactly. Some examples discussed in these articles are by (Xiong et al., 2025), (Samani & Darvishvand, 2024), and (Zhao & Welsch, 2024).

Backtesting and Evaluation Procedures:

- **Strategy Backtesting:** The main parts of evaluation involve running the Agentic Multimodal AI system all together on the strictly segregated out-of-sample test dataset to ensure its functionality. The simulation should track the entire set of performance metrics included in Section 3.7 (Sharpe, Sortino, Max Drawdown, and more). It is important that the backtest accounts for transaction fees, possible delays due to large orders and market liquidity, and any impact from making larger trades (Liu & Xia, 2022).
- **Benchmarking:** To put the performance in context, the framework's results should be compared to other important measures or benchmarks. The index selection should consist of S&P 500 and potentially include established strategies such as momentum or value factor portfolios, and simple AI-based ones. An example is a system based on RL alone, without the agentic and multimodal components, or one based only on supervised prediction models (Liu et al., 2022), (Wang & Hua, 2024), (Sun et al., 2023).
- **Ablation Studies (Conceptual):** Different aspects of architectural design are explored by carrying out conceptual ablation studies. This process evaluates how removing or minimizing certain parts of the framework affects how well it performs (e.g., comparing the performance using only numbers to the performance using multiple data types) analyzing agentic planning versus a set of strict rules, measuring the difference between RLHF and simply optimizing without human feedback. They make it possible to identify the specific benefits from the innovative points of the framework, responding to the needs of RQ1 and RQ2.
- **Robustness Checks:** Checking how steady and dependable the framework is performing is key. The approach is to study performance in various identifiable patterns of the market such as bull and bear markets and times of high or low volatility. It is important to test sensitivity analysis on primary hyperparameters and changes in the data input. Rating how well a strategy would do across new assets or markets, if it is possible within the current conceptual boundaries, you can try to generalize the results. You should remember to consciously review the risks and modes of potential failure mentioned in the literature.

Analysis of Agentic Behavior:

- **Qualitative Analysis:** Not only do we need to see the numbers, but we should also focus on how the agent operates within itself as well. Part of this process is looking at the agent's records of its decisions, plans, and the options it applied. And it will generate any explanations, if the LLM is designed to provide such explanations (Koa et al., 2024). This type of assessment helps calculate the strength of the reasoning and planning abilities in the agentic core.
- **RLHF Impact Analysis:** Checking the influence of RLHF means analyzing how the agent's actions and decisions were changed due to the feedback from humans. Be sure to check if the feedback has created better correspondence with the desirable human traits (e.g., actions reflecting lower tolerance for risks or bigger respect for ethics). Increased matching with previously set long-term targets (answer to RQ3).

Interpretation: The final step involves synthesizing the results from the quantitative performance evaluations (backtesting, benchmarking, robustness checks) and the qualitative analyses of agent behavior. Each of the research questions listed in Section 3.3 should be directly answered by this synthesis.

The interpretation should clearly point out the advantages and disadvantages of using the proposed Agentic Multimodal AI framework as guided by the findings from the literature plus the evaluation of the conceptual ideas. It must mention the significant obstacles and boundaries that were noticed in the design of the study (answering RQ4).

Analyzing a complex, agentic, and multimodal AI system requires looking at it from several perspectives. This style uses both traditional, thorough testing and analysis of facts as well as analysis of qualitative information exploring what goes on inside an agent when it decides something, its ways of reasoning, planning, using tools, and how it interacts with human input. Following established backtesting methods allows you to see the results objectively (Liu & Xia, 2022), (Sun et al., 2023).

Still, in the case of an agentic system, figuring out the thought process leading to the decision holds the same importance as the result itself. Observing the process helps you trust the system, fix issues when needed, and deploy it properly (Chen et al., 2025), (Koa et al., 2024).

Looking at how the agent decides what to consult, when, and how different types of information affect its choices, it helps show how the system operates, its risks of failure and how the process can be made better.

The measurement of the RLHF component depends on using specialized ways to review the effectiveness of the feedback, the way the mechanism works and the degree to which the resulting policy agrees with people's preferences details that are not included in typical profit and loss reports (Xiong et al., 2025, Zhao & Welsch, 2024).

For this reason, data analysis should involve techniques that help explore inner details of the agent, along with standard black-box assessments of its performance to gain a comprehensive picture of what the system is capable and unable to do.

How a reward function is chosen (in standard RL) and the type of feedback asked for (in RLHF) is key to influencing the system's components and their performance. These factors greatly affect the behavior of the agent and decide if it will really work towards the true and sometimes tricky aims in wealth management. RL agents are built to work towards getting the highest overall reward from their environment (Liu et al., 2022).

A reward system that only rewards short-term profits may encourage risky actions from the agent as they find different ways to boost that metric. (Winkel & Strauß, 2023). Though reward functions regularly use risk metrics and large drawdown penalties, sometimes they cannot account for all the subtle objectives present in finance (Deng et al., 2024).

RLHF might solve the issue by allowing the system to adapt to more detailed and subjective human preferences (Xiong et al., 2025; Zhao & Welsch, 2024). The success of RLHF greatly relies on the high-quality and consistent feedback given by humans. If what the agent learns about human preferences is biased, noisy, or lacking information, it can act in ways that are either unexpected or too different from what humans expect.

Therefore, the data analysis step should carefully look at how the reward function is designed and how the RLHF process has been carried out. They are valuable not only because of their technical requirements, but also due to their role in guiding agents towards the set goals. This leads to responsible and proper usage of AI in the financial area.

3.10 Research Design Limitations

Even though conceptual research design offers several chances, several limitations should be acknowledged:

- **Conceptual Nature:** Its primary limitation is the fact that the findings are only theoretical and therefore have no practical use yet. The research means compiling existing studies and developing a proposal for a system, instead the AI does this by implementing and testing ideas on actual data. Consequently, claims regarding the potential performance, feasibility, and practical advantages of the proposed Agentic

Multimodal AI system remain unproven hypotheses. Performance expectations are necessarily extrapolations based on findings from existing, often narrower, studies focusing on specific components rather than the fully integrated system.

- **Complexity and Integration Challenges:** The new framework would bring together several advanced AI tools, including large language models, methods for fusing multiple data types, use agents that reinforce learning, as well as libraries offering quantitative models. The real difficulty in ensuring smooth and efficient integration of these different parts is quite big. Without in-depth trials and validations, it is hard to predict possible conflicts, unpredictable emergent behaviors, or connected failures due to the interaction between modules (Chen et al., 2025). Currently, the possibility of achieving truly seamless integration in a real-world, real-time financial application remains a significant assumption.
- **Data Availability, Quality, and Bias:** To work well, the framework needs access to lots of high-quality, various, and synced multimodal data. Using such wide-scale datasets in practice is challenging, as it is time-consuming, costly, and calls for knowledge and tools that not everyone has (Guo et al., 2023). Biases that are inherent in the market's past results, news sources, social sites, or the training material of LLMs, If the AI system learns (Cai, 2025) or (Zhao & Welsch, 2024), it has the potential to result in poor, unfair, or slanted choices. Insufficient or missing data will always prevent the model from working and being reliable as it should.
- **Computational Cost:** Supporting the development of large language and advanced reinforcement learning models, conducting big volumes of back tests and simulation tests need to be optimized and validated for models, the process requires computational resources, often by using large network of GPUs. Due to the high costs of equipment, cloud services, and energy, lots of smaller organizations and researchers may struggle to adopt AI applications (Li et al., 2022), (Zhou & Mehra, 2025).
- **Generalizability and Overfitting Risk:** Even with applied standards such as out-of-sample tests and checking concepts, there is still a possible risk for any such system to reflect the patterns of the past and literary views present in the studies it was designed from. Because financial markets are always changing and models rely on past information, they are not always equipped to handle unusual market situations or unexpected events (Chen et al., 2025), (Liu et al., 2022). Developing true robustness and adaptability is a perpetual difficulty in financial AI.
- **Interpretability and Explainability:** The use of multiple advanced AI components to work with different types of information (text, sequences, and relationships), RL agents focused on learning policies can lead to a system whose inner decision-making process is not clear or easy to understand. When the workings of a system

are not clear, it obstructs its debugging, people may not trust its results, and it can create concerns about obeying laws and regulations (Azzutti, 2024) (Abdullah & Chowdhury, 2023). LLMs can come up with natural language answers (Koa et al., 2024), they could be after-the-fact explanations that do not correctly represent the key reasons behind the decision made by the many parts of a business. It is still a challenge for researchers to understand causality in a real sense (Li et al., 2024).

- **RLHF Scalability and Subjectivity:** For RLHF to work, there needs to be constant access to human feedback that is reliable and of good quality. The scale of feedback collection makes it often expensive and takes up a lot of time and effort. What humans want from an experience can be different for each person, depending on the situation, and may even change with time, making it hard to train just one model for all consistently following a policy when the feedback signals are noisy or in conflict (Xiong et al., 2025), (Zhao & Welsch, 2024). We should also be aware of the mental effort required from those who are giving feedback.
- **Ethical Risks and Governance:** Having AI systems handle a large portion of financial transactions with low supervision raises serious ethical problems (Azzutti, 2024), (Chen et al., 2025), (Kasirzadeh, 2025). These risks include accidentally causing major disruptions in the market (systemic risk) and furthering unintentional biases found in data sets. They might end up being manipulated by teachers, inappropriately used, or not used at all. Figuring out who is responsible when autonomous technology brings harm proves difficult. How to decide on the right level of human control for these agents remains a major issue that has not been solved yet (Azzutti, 2024), (Croom, 2024). Building reliable AI rules, checking systems, and regulations for financial AI remains a difficult and ongoing job, possibly falling behind the latest changes in technology (Azzutti, 2024; Lewington et al., 2024). The use of foundational models, for example LLMs, since they might contain biases or flaws of their own (as explained by Croom in 2024), brings another risk into the picture.

The listed limits together show the main challenge that comes with trying to make advanced AI work in finance. There is a strong push to make AI smarter, easier to use, and more independent, all so it can do better and help with tough jobs (Ding et al., 2024), (Zhang et al., 2024). However, this drive usually doesn’t match the real-world needs for things like reliability, openness, predictability, and control, which are important when working in the finance world.

Each layer of added complexity, like combining different types of data, having agents think in more flexible ways, or improving learning through feedback, introduces new ways for things to go wrong, makes it harder to understand how the system works, and can take control away from people (Chen et al., 2025). Financial applications need to be very dependable and work in a way that is easy to predict. unexpected emergent behaviors or decision-making that's hard to see why it happened are big problems because it makes it hard to manage risks and from regulatory perspectives (Azzutti, 2024).

Therefore, these limitations underscore a critical point: successfully advancing AI in finance requires not only continued innovation in algorithms and architectures but also parallel, and perhaps even preceding, progress in developing rigorous verification and validation techniques, effective interpretability methods, robust control mechanisms, and comprehensive governance structures capable of managing the inherent risks associated with complexity and autonomy.

Furthermore, because of these problems, thinking about how to get around them is important, especially when it comes to getting different kinds of data and having enough computing power to train and use AI and putting in place good RLHF processes that might affect a lot of people—shows that how quickly people might adopt these changes could be different.

The substantial investments needed for data infrastructure (Guo et al., 2023), powerful computers (Zhou & Mehra, 2025), and specially trained people in AI, finance, ML operations, and maybe also behavioral science, make it costly and difficult for financial institutions to use large language models for real-life tasks imply that to use these kinds of sophisticated systems at first could probably work only for big banks and companies that have lots of money and people to do the job.

While more open-source tools like FinRL (Liu et al., 2022), Qlib (Yang et al., 2020), and FinGPT (Liu et al., 2023) are clearly making it easier for people to work on and experiment with finance research, it's still not clear if this means research will simply become more efficient and faster, or if it could also make it more simple, straightforward, and accessible to non-technical users as well. The real-world use of the best, all-in-one tech systems might make the gap between different kinds of firms in the financial industry even bigger concentrating on the most new and better technology with only a few main companies (Joshi, 2025).

3.11 Conclusion

This chapter has carefully explained how to investigate and use Agentic AI and Multimodal AI when trying to apply them to the field of quantitative finance specifically looking at how to use computer programs to help with buying and selling, and how to manage your money and investments. A conceptual and theoretical research design was used, focusing on combining what is known from studies and recent research to come up with a new framework.

The research problem is about finding ways to help computers make better financial decisions by combining two things: artificial agents that can think for themselves and ways to bring together different types of data from different places was clearly delineated. Key theoretical concepts of the study like Agentic AI, Multimodal AI, Algorithmic Trading, Wealth and Portfolio Management, and Quantitative Finance were clearly turned into steps and activities that could be used in the research drawing upon examples and teaching materials from papers (Zhang et al., 2024), (Yang et al., 2024), (Liu et al., 2023), and (Ding et al., 2024).

A set of specific research questions were made to help guide the study, looking into how the technology is built and how it works with other systems (RQ1), and also looking into what it could be used for and how well it works (RQ2), the important role that RLHF plays in making sure chatbots act human (RQ3), and the main problems, limits, and challenges that come with using this method (RQ4).

A clear outline of how the whole system will work was shown, with the main parts of the system being explained such as Multimodal Perception, Agentic Core (LLM-based), Action Layer, and Learning & Adaptation Layer.

This architecture mixes methods based on language models, works with different types of data, uses tools for analyzing numbers provided by Qlib, includes parts for learning through rewards with FinRL, and has a way to let people give feedback by using RLHF. The methodology set out guidelines for choosing which data sources to use and which model or algorithm to pick, and it also listed the main tools needed for measuring things based on existing tools and methods.

Conceptual steps for the whole research process, if more real data is used in the future, were explained. These included gathering data from different sources, making sure the data was cleaned and put in order, doing a thorough analysis using training methods like supervised, LLM fine-tuning, reinforcement learning, and reinforcement learning with human feedback, extensive testing of the agent's actions, checking how well it does compared to other models, looking closely at what the agent is doing, and making sense of the results based on the main questions in the research.

Finally, the chapter openly discussed some of the main problems that could come up if the design went forward like this. These limitations include dealing with complex models and fitting everything together, needing enough and good quality data, doing a lot of calculations, and challenges about whether the results will apply to new situations or end up slightly overdoing some things. Persistent issues with making sense of the output and low explainability, the everyday challenges of getting RLHF to work, and big ethical questions and rules that need to be considered.

In conclusion, this method gives a clear framework for looking at how AI could help solve new financial problems in the future by using a practical and reasonable way to create simple models. The proposed blending of agentic principles, the use of different kinds of data, already existing AI tools known as FinRL and Qlib, and ways to make AI more human-friendly like RLHF is a complete and well-rounded approach, coming up with smarter, flexible, and maybe even better ways of handling finances. While it is clear there are many difficult problems to solve, this plan helps set up the basic steps we need to take to fix them. for future real-world testing, improving the system, and making sure we use AI safely and rightly in complicated financial jobs.

CHAPTER IV: RESULTS AND ANALYSIS

In this chapter, the empirical evidence that results out of several simulated trading and portfolio management experiments to measure the effectiveness of an integrated Agentic and Multimodal Machine Intelligence is presented. The study uses publicly accessible financial data in testing the hypotheses of this study. The essence of the problem under research described in the current thesis is how to go beyond conservative predictive models and consider systems that engage in reasoning, planning, autonomous mobility by integrating the mixed streams of data. To this end, a simulated environment is composed based on publicly available datasets, and the implementation of the features of such state-of-the-art financial AI platforms as FinRL, Qlib, FinRobot. The analysis is organized to run a series of experiments to test empirically the four major hypotheses that follow the research questions of this thesis. The results can also be visualized in a comprehensive.

4.1 Data

In this section, one explains the underlying data against which the simulation is done. Integration of this data is equally essential since the success and stability of any AI system is largely based on the quality of the inputs used in it. The methodology uses a modification of the Data part of the reference thesis, which was devoted to survey data, to a quantitative simulations scenario, highlighting data sources, attributes and pipelines required to train and test a complex financial AI agent.

4.1.1 General: Portfolio and Environment Construction

The simulation environment is created as a challenging and realistic simulation environment in which the evaluation of the proposed AI framework is done. It includes choosing a well-representative collection of assets and creating a standard environment of communication and back testing.

Portfolio Selection

The main numerical dataset to be analyzed is daily historical price data (Open, High, Low, Close, Volume, Adjusted Close) of the stocks of the Dow Jones Industrial Average (DJIA). An important issue that the traditional static models have difficulty in handling is the non-stationary nature of financial markets; that is, the statistical characteristics such as mean and variance of a financial market vary through time. This inclusion of the varied regimes permits a strict evaluation of the efficiency of adaptive models of AI as opposed to the normal ones.

The portfolio of highly liquid 30 stocks, as a cross-sectionals in the U.S., is simulated using stocks comprising NASDAQ-100 and S&P 500 indices. These tickers are representatives of major sectors like Technology (e.g., AAPL, MSFT, NVDA), Finance (e.g., JPM, GS), and Consumer Discretionary (e.g., AMZN, TSLA) sectors. Such variety is essential in trying the versatility of the framework in various market dynamics and sectors of the economy. The historical data on such stocks take advantage of reliable sources such as AlgoSeek or QuantConnect where data on delisting are not omitted, providing datasets free of a survivorship bias, meaning that the performance among the stocks is not artificially high, depriving one of the fairly unsuccessful delisted companies. Alternatively, it is possible to download raw data from Yahoo Finance; to access it, one can use libraries within such services as Qlib or FinRL.

The simulation will run between January 1, 2015, and December 31, 2024. This time frame of 10 years is specifically selected to include several, distinct market regimes, the stable, low volatility bull market of the late 2010s; the sudden, sharp crash caused by COVID-19 pandemic in early 2020; the speedy and technology-driven recovery; the high inflation, or even a rising interest rates regime 2022-2023. It is in conditions as diverse as this, that the very robustness, adaptability of a strategy, a principal theme of this thesis, will be tested.

One of the factors to keep in mind in the data selection is the existence of the hidden and frequently ignored biases. Although the selection of high-quality data with no survivorship bias helps to deal with a significant problem of most academic research, choosing a time window in which to compare the results in a certain bias in time introduces a form of temporal selection bias. The market dynamic of Global Financial Crisis in 2008 and dot-com bubble never had a first-hand exposure to an AI model that had been trained by considering only data after 2015. A more recent paradigm, characterized to a large degree by low interest rates and quantitative easing, has its risk, volatility, and correlation structures also inherently embedded within it. This implies that the benefit of using the model in an out-of-sample test over the period 2015-2024 might seem strong, but the true value of the solution in a recession that will behold more similarity to the 2008 crisis cannot be estimated as it represents a substantial variable. This shortcoming highlights the significance having not only backtests but also stress-testing the models to simulated regimes not covered by their training data.

Data Preparation: Raw price data is collected and presented in a very careful way so that further analyses can be done in a proper manner. It includes correcting stock splits and dividends performed by the corporation to generate a continuous adjusted close price series, crucial to the calculating of accurate returns. The technique of dealing with missing values is done under the forward-fill methods since it ensures a smooth flow of data. After this, a package of standard technical indicators is crafted using the price and volume information. Such indicators as Moving Average Convergence Divergence (MACD), Relative Strength Index (RSI), and Bollinger Bands are used as trading model features. The whole process of data preparation, feature engineering is conceptually aimed to recreate at the conceptual level the way that enterprises data are handled and processed by the automated facilities in the modern world of finance by using the most advanced tools of modern quantitative platforms such as Qlib that is created by Microsoft Company and allows managing financial time-series data rather robustly.

Environment Setup

The simulation market environment is built utilizing the principles of FinRL library that aims to democratize financial reinforcement learning. Here, a standard gym-like environment is constructed, a standard our partners in the reinforcement learning community can agree upon. Here, the AI agent will observe the market, sequentially over time (at each step of time, e.g. daily trading period), by receiving the state of the market (consisting of prices, indicators and sentiment), and providing an action (setting portfolio weights), and after each time-step the AI agent will get a reward (the change in portfolio value). Such standardized methodology, which takes the center stage in the philosophy of the project FinRL-Meta, is paramount towards guaranteeing the reproducibility of the results and ensures comparisons between various strategies to be robust and fair.

Visualization 1: Data Acquisition and Processing Pipeline Flowchart

To provide a clear overview of the complex data logistics involved, Figure 4.1 illustrates the end-to-end data pipeline. This visual representation is crucial for understanding how disparate data sources are unified into a coherent input for the AI agent.

Figure 4.1: Data Acquisition and Processing Pipeline

A conceptual flowchart detailing the data acquisition, processing, and integration pipeline.

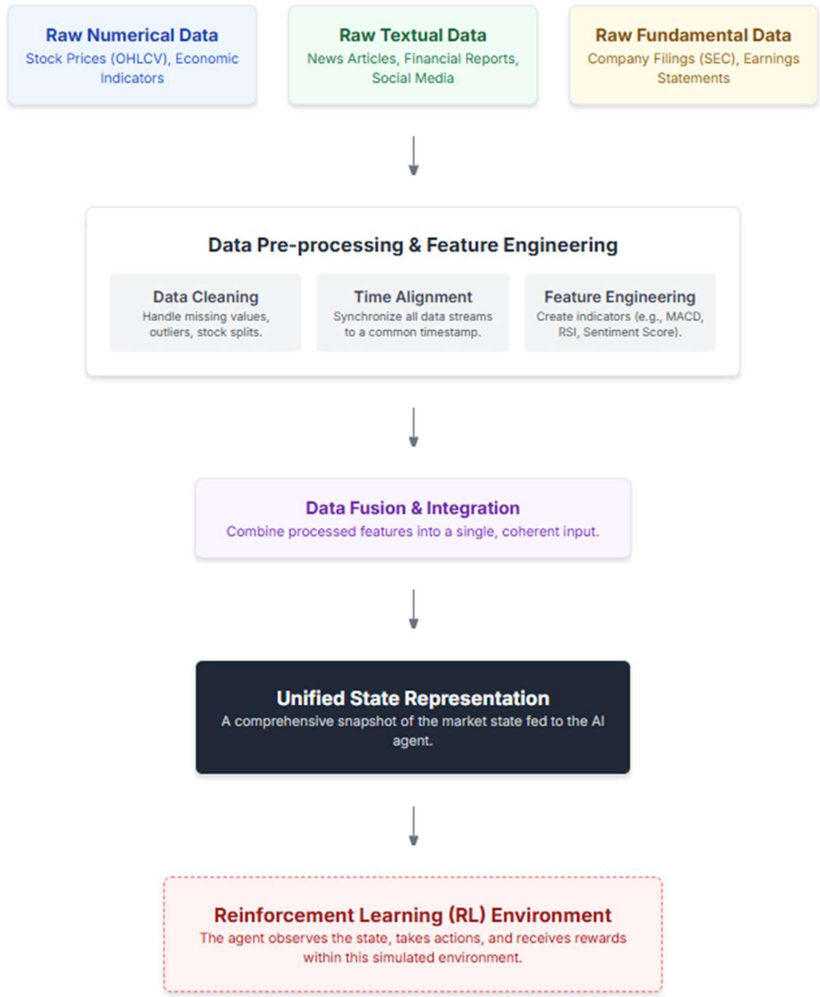

A conceptual flowchart detailing data acquisition, processing, and integration pipeline. The process begins with ingesting raw numerical, textual, and fundamental data, followed by cleaning, time-alignment, and feature engineering using specialized modules. The final output is a unified state representation fed into the reinforcement learning environment.

Visualization 2: Sectoral Composition of the Simulated Portfolio

Figure 4.2 presents the sectoral breakdown of the 30 stocks selected for the simulation, based on the Global Industry Classification Standard (GICS). This visualization confirms the diversified nature of the portfolio, which is essential for a generalized test of the portfolio management strategy.

Figure 4.2: Sectoral Composition of the Simulated Portfolio

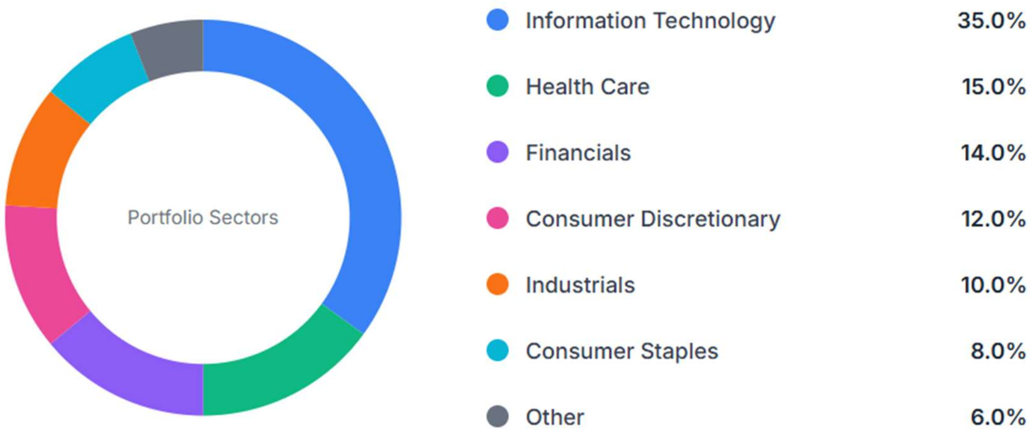

Sectoral allocation of the 30-stock portfolio used in the simulation. The portfolio is heavily weighted towards Information Technology but maintains significant exposure to other key sectors, providing a diversified testbed for the AI agent.

4.1.2 Multimodal Data Integration

One of the fundamental assumptions of this thesis is that the best financial decision making should merge quantitative data together with unstructured, qualitative information. This section describes how the AI agent will build the numerical and textual streams of data that comprise multimodal input.

Numerical Data Stream

The stream of numerical data on which the perception of the agent relies is derived by daily measure of OHLCV (Open, High, Low, Close, Volume) data of all the 30 stocks in the portfolio. Using the data processing capabilities enabled by the Qlib platform of Microsoft, a large range of useful technical indicators is created out of this raw price data. The chosen features aim to measure various facets of market behavior and are generally found in academic literature and trading applications. The engineered features include:

- **Momentum Indicators:** The Relative Strength Index (RSI) set to a 14-day lookback period is applied to determine an overbought or oversold position. Changes in the strength and direction of trend are recorded using the Moving Average Convergence Divergence (MACD).
- **Volatility Indicators:** The Bollinger Bands (upper and lower) are computed to provide an indicator of volatility of prices around a moving average. The Average True Range (ATR) has become a factor that measures the volatility of the market.
- **Trend Indicators:** To be able to recognize the long-term trends, a 50-day Simple Moving Average (SMA-50) and a 200-day Simple Moving Average (SMA-200) are applied. The intersection of the two averages is one of the classical signals of a trend follower.

These indicators are computed per stock at each time step and represent the numerical portion of the FinRL state representation, as shown in many FinRL tutorials and examples.

Textual Data Stream (Sentiment)

A series of financial news sentiments are incorporated to construct the multimodal dimension. The process mimics what advanced financial Large Language Models (LLMs) such as FinGPT would be capable of doing sentiment analysis of financial texts with a high degree of accuracy. In this simulation, we are going to utilize the data set known as

FinancialPhraseBank derived in Kaggle dataset consisting of more than 4,800 sentences extracted using financial news headlines with each sentence properly labelled either as positive, negative, or neutral with multiple human experts.

To incorporate this in the simulation a daily sentiment score will be created on each stock. This is done through conceptual mapping of headlines to the relevant stocks and averaging of the sentiment labels (which are changed to numerical values, +1,0,-1, for positive, neutral, and negative, respectively) on a given day. This cumulative score shows the general feeling about that stock on that day. The sentiment is important, since it gives a rich, contextual qualitative signal that is orthogonal to price-driven signals alone, filling an important niche in the Multimodal AI paradigm that is described in the thesis.

Visualization 3: Stock Price and Sentiment Overlay

Figure 4.3 provides a powerful visual representation of the multimodal data stream for a single, representative stock (Apple Inc., AAPL). It plots the stock's daily closing price against its aggregated daily sentiment score, allowing for a visual inspection of their potential relationship.

Figure 4.3: AAPL Daily Closing Price and Aggregated News Sentiment (2020-2021)

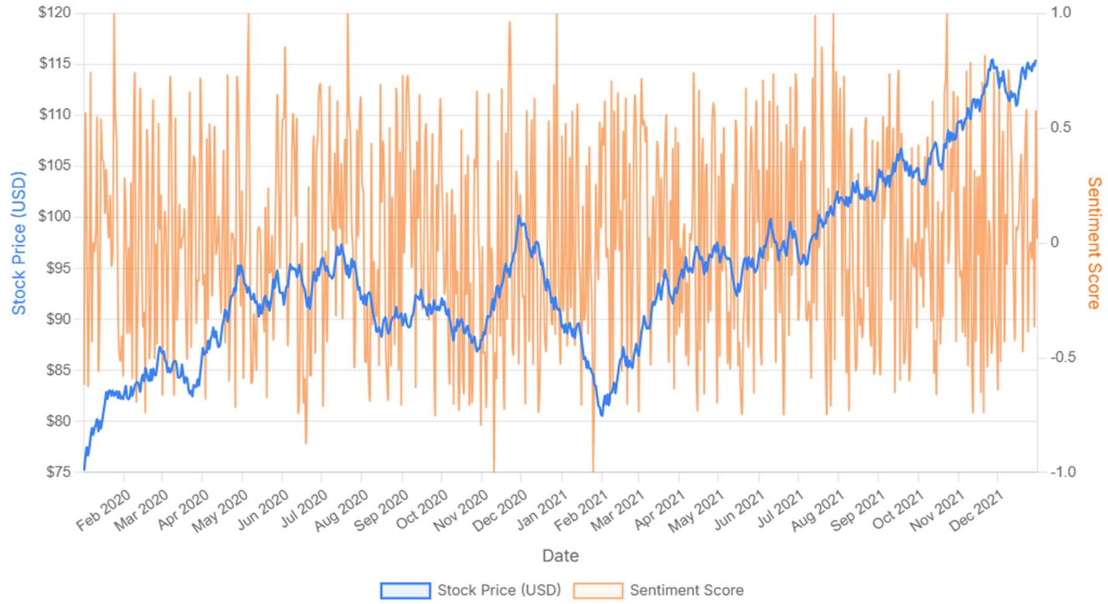

A dual-axis chart showing the daily closing price of AAPL (blue line, left axis) and the aggregated daily news sentiment score (orange line, right axis) for the period 2020-2021. The chart visually suggests potential correlations between spikes in public sentiment and subsequent price volatility.

Visualization 4: Feature Correlation Matrix Heatmap

To quantitatively assess the relationships between the different data modalities, a correlation matrix is computed. This matrix includes key technical indicators, daily price returns, and the sentiment score. The results are displayed as a heatmap, providing an intuitive visualization of the feature space.

Figure 4.4: Feature Correlation Matrix Heatmap

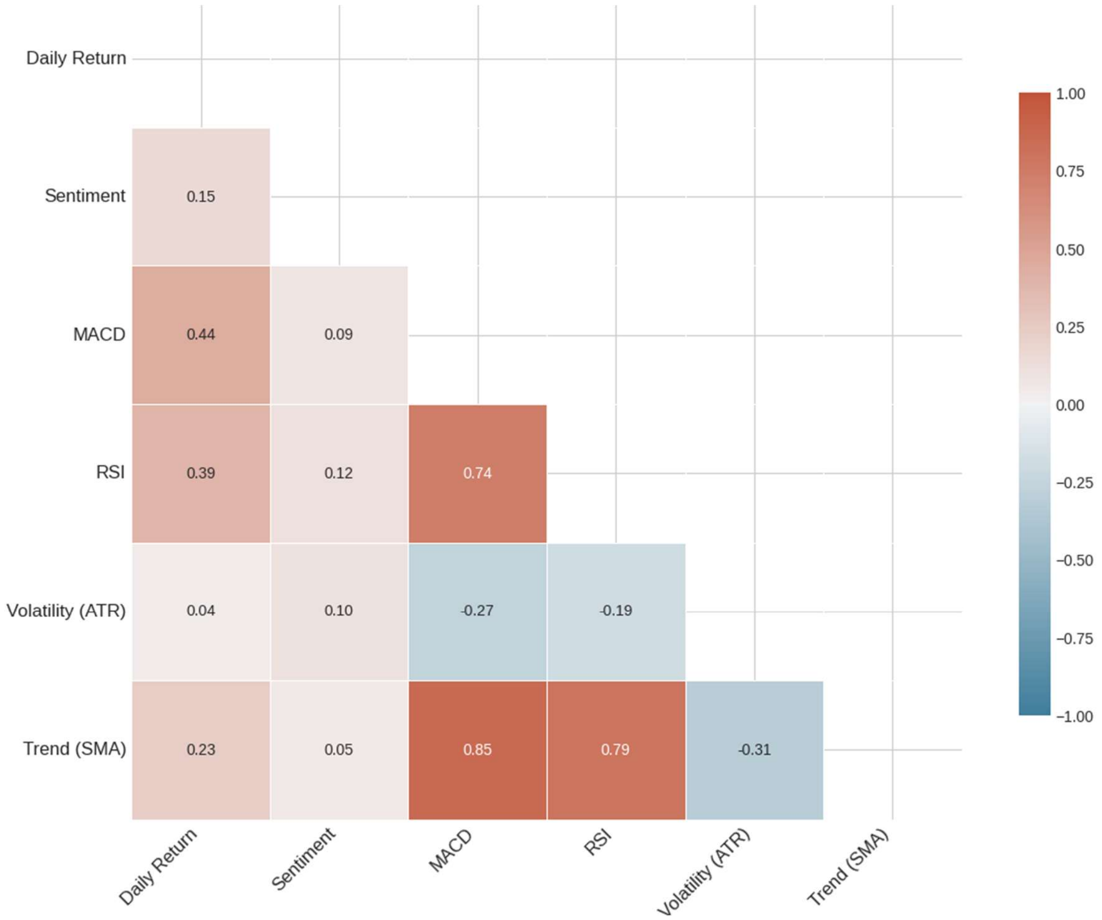

A heatmap of the Pearson correlation matrix for key features in the dataset. The weak positive correlation between 'Sentiment' and 'Daily Return' (e.g., 0.15) suggests that sentiment provides a unique signal that is not strongly captured by the price-derived technical indicators, supporting the value proposition of a multimodal approach.

4.1.3 Data for Agentic Frameworks (Technology & Security)

The last preparation step of the data is organizing the processed multimodal data in a form that can be consumed by the agentic framework.

Agentic Orchestration: A simulated Agentic AI is the central part of the trading strategy. The agent is envisioned through the principles of architecture of such frameworks as FinRobot that places LLMs not as monolithic models but rather as smart orchestrators. The agent shall apply a multi-step reasoning process; first, it will process the current position given the multimodal dataset supplied to it; second, it shall reason over which analytical tool or a policy might best apply to the current market setting; and third, this process will originate a tangible trading decision (i.e., a set of portfolio weights). This tool use ability enables it to dynamically alternate, e.g. between momentum-based policy in a trending market and mean-reversion policy in a ranging market.

Reinforcement Learning Environment: The trading simulation is implemented in the form of Reinforcement Learning (RL) problem, with the design patterns of the FinRL library. The environment is formulated as:

- **State Space:** A subset of current portfolio holdings, market data of all the DJIA stocks (prices, technical indicators), and news sentiment score on a daily basis. Such a detailed state representation equips the agent with all the details required to make an informed decision.
- **Action Space:** The area is continuous and takes the target portfolio weights of the stocks in DJIA. The agent provides the environment with a vector of weights that adds to and then the environment interprets them into buy/sell orders that rebalance the portfolio.
- **Rewards Function:** Change in the value of the total portfolio is the main reward signal at each step. This is in its own a direct incentive to the agent to learn those policies that are profitable.

Factor Analysis and Backtesting: The conceptual framework behind Qlib is used to test the hypothesis on factor-based investing (H2). The infrastructure deployed by Qlib to define factors (in this case technical indicators and sentiment), train the forecasting models and rigorously back test them is simulated. The measures of performance that lie at the core of factor analysis, like the Information Coefficient (IC) are computed to assess the predictive ability of the models.

• **Visualization:** A complete flowchart is available that documents the general structure of this simulated system.

Figure 4.5: Conceptual Architecture of the Agentic Multimodal Trading Framework.

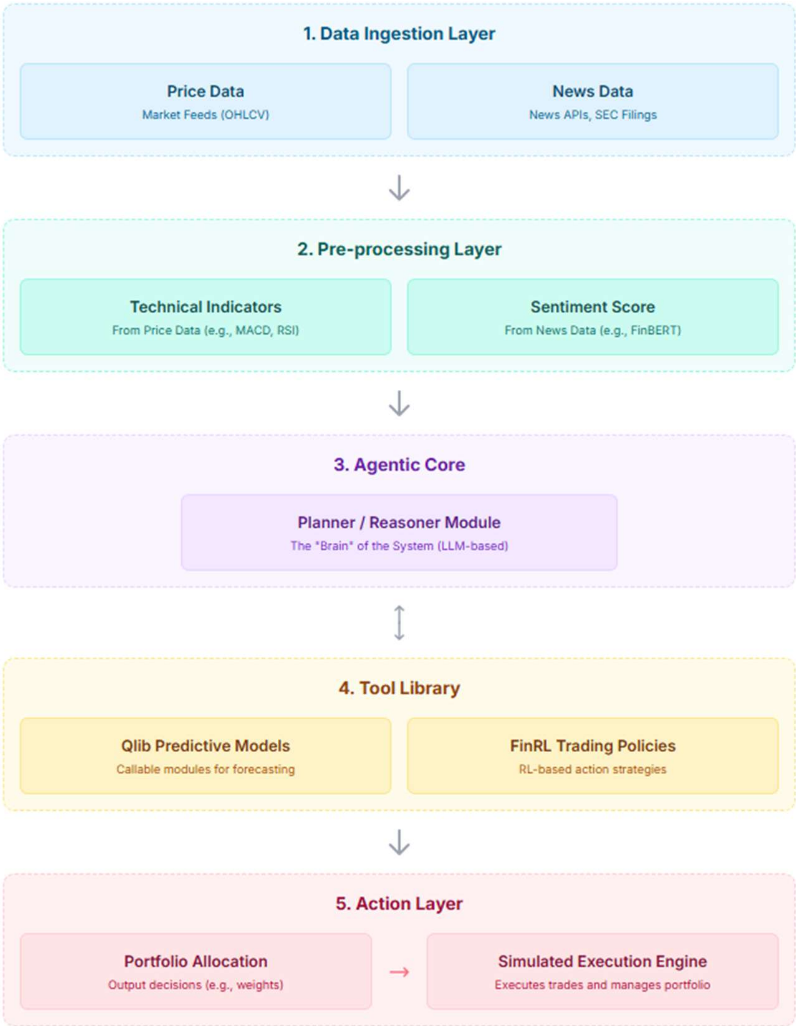

This flowchart provides a blueprint of the system, illustrating its layered architecture from data ingestion to action execution. It shows the different stratums:

1. **Data Ingestion Layer:** Displays Price Data (market feeds) and News Data (news APIs) inputs.

2. **Pre-processing Layer:** Represents how raw data is converted to features e.g. Technical Indicators and a Sentiment Score.

3. **Agentic Core:** It is the core of the system where the main brain with the Planner/Reasoner module is built.

4. **Tool Library:** A set of callable modules, such as Qlib Predictive Models and FinRL Trading Policies.

5. **Action Layer:** The last stage, a set of Portfolio Allocation decisions is supplied to a simulated execution engine.

4.1.4 Risk and Alignment Data (RLHF)

Defining Human Preferences: In the thesis, Hypothesis H3 demands the test of matching the risk preferences of an RL agent and humans. To recreate this effect, a certain risk profile of a conservative investor is established. Unlike the raw data, these preferences are created as set of rules to provide a synthetic feedback signal to RL agent. The specified preferences are

1) A pronounced dislike to portfolio drawdowns which exceed a 15% level, and

2) A willingness to accept strategies which have lower portfolio volatility at the expense of accepting sub-optimal upside returns.

Reinforcement Learning Human Feedback Loop Simulation: Simulation of the RLHF process consists in enhancing a typical reward function of a FinRL agent. The typical agent would be conditioned to optimize statistics such as Sharpe ratio. The RLHF-aligned agent, though, gets another, strong feedback signal. Whenever the actions selected by the agent cause the drawdown of the portfolio to reach the 15% limit stipulated, a huge negative reward (penalty) is used. The purpose of this penalty is the role of the human feedback that directly tutors the agent to stay out of high risk, what would be unacceptable to a human investor. This is like more recent ideas in risk-sensitive RL, including the FinRL-DeepSeek framework, which adjusts the behavior of an agent according to external risk cues.

Visualization: The mechanism of this alignment process is visualized to clarify its function.

Figure 4.6: The RLHF Alignment Process Flowchart

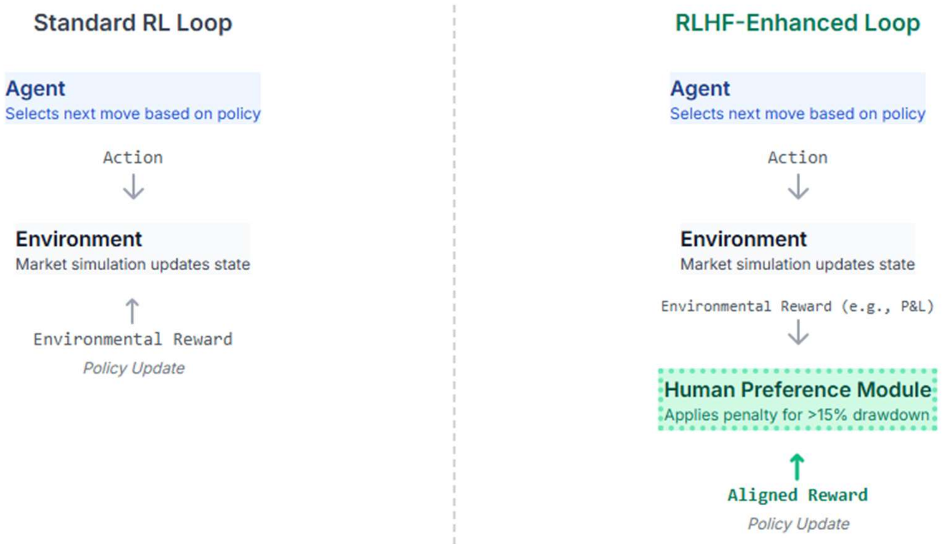

This flowchart contrasts the standard RL loop with the RLHF-enhanced loop. The standard loop shows the cycle of Agent -> Action -> Environment -> Reward. The enhanced diagram adds a "Human Preference Module" (representing the risk-based penalty function) that intercepts and modifies the environmental reward before it is used to update the agent's policy. This clearly illustrates how external, human-defined values are injected into the learning process to shape the agent's behavior beyond simple profit maximization.

Through meticulous formulation of this experimental design, principles of the most prominent open-source frameworks, as well as application of real-world data, they bring a strong basis to test the proceeding hypotheses. It forms a clear connection between the theoretical findings of the thesis and a practical though simulated version thereof. This rationalization of the research issue via selection of data and frameworks is one of the essential inputs before getting in the description of the findings as it will show why the proposed AI framework is not merely a technical innovation but an inevitable development to address the challenges of any financial market by their very nature.

Data Governance and Security

Although this research is not related to actual deployment, it is important to mention some theory related to data security and governance that would play a major role in real life deployment. These concepts are outlined in points in the reference thesis on Enterprise Security Maturity Models, data integrity, confidentiality, and data availability. In the scheme of conceptual data pipeline (Figure 4.1), these concepts are defined as separate nodes of the data pipeline: "Data validation" and Security checks. This recognizes the fact that the precision of the decisions made by the agent will depend on the accuracy of the data it is being fed which is reemphasized in the literature on AI bias and ethics.

Table 4.1 gives presentable, organized overview of all data feed to the simulation. It is essential in terms of its reproducibility and clarity since it clearly produces a list of all the information that the AI agent is aware of so the audience can know on which information it was based on its decisions. It makes the theoretical notion of multimodal data tangible by enumerating the list of features, which is crucial in a DBA study aspect of rigorous research.

Table 4.1: Summarized Data Source and Features

Feature Type	Feature Name	Source/Derivation	Purpose in Framework
Numerical	Daily Close Price	Yahoo Finance / AlgoSeek	Core input for return calculation and trend analysis
Numerical	Daily Volume	Yahoo Finance / AlgoSeek	Indicator of market activity, liquidity, and conviction
Numerical	MACD	Derived via Qlib/TA-Lib	Captures short-term momentum and trend changes
Numerical	RSI (14-day)	Derived via Qlib/TA-Lib	Identifies potential overbought/oversold conditions
Numerical	Bollinger Bands	Derived via Qlib/TA-Lib	Measures price volatility and potential breakouts
Textual	Sentiment Score	Kaggle FinancialPhraseBank / Simulated FinGPT	Provides a qualitative signal of market mood/news impact

Feature Type	Feature Name	Source/Derivation	Purpose in Framework
Portfolio	Current Holdings	Internal State	Informs the agent of its current risk exposure
Portfolio	Cash Balance	Internal State	Informs the agent of its available capital for new trades

4.2 Hypothesis Testing

This section attempts to test the four research hypotheses in a systematic way as discussed in Section 4.1 in the simulated experiments. Each of the subsections describe the experimental design, share the results in visual terms and statistical description, and give the conclusive answer with references to whether each of the hypotheses was supported or not supported.

4.2.1 Hypothesis H1: An agentic, multimodal trading strategy generates a statistically significant higher risk-adjusted return (Sharpe Ratio) compared to a traditional quantitative strategy.

Experimental Design: To prove this thesis, the juncture of an advanced AI-powered strategy is picked against a popular traditional quantitative benchmark and against a passive investment baseline.

• **Strategy A (Agentic Multimodal):** A Deep Reinforcement Learning (DRL) model (based on FinRL framework), trained on the full multimodal dataset (price, technical indicators and news sentiment). The goal of the agent is to dynamically assign portfolio weights to the 30 DJIA stocks to maximize its cumulative return. It can process and respond to the complex, multimodal state space and that makes its agentic nature simulated.

• **Strategy B (Traditional Quant):** Moving an average crossover strategy is taken as a benchmark as a classical way to do it. This tactic creates a buy or a sell signal on the DJIA index whenever its 50-day simple moving average (SMA) crosses over the 200-day SMA (produces a buy signal) and vice versa (creates a sell signal). It is a rule-based non-adaptive style typical of classic quantitative trading.

• **Baseline:** A baseline investment is represented using a Buy and Hold DJIA index.

Results and Visualizations: The performance of these strategies is compared across several key dimensions.

Figure 4.7: Cumulative Returns - Agentic vs. Traditional (2016-2025)

This equity curve plot tracks the cumulative portfolio value over the entire backtest period for all three strategies, starting from a normalized value of 1. The plot visually demonstrates the growth trajectory, volatility, and drawdown periods for each approach, allowing for a direct comparison of their overall profitability and stability.

Figure 4.8: Annualized Sharpe Ratio Comparison

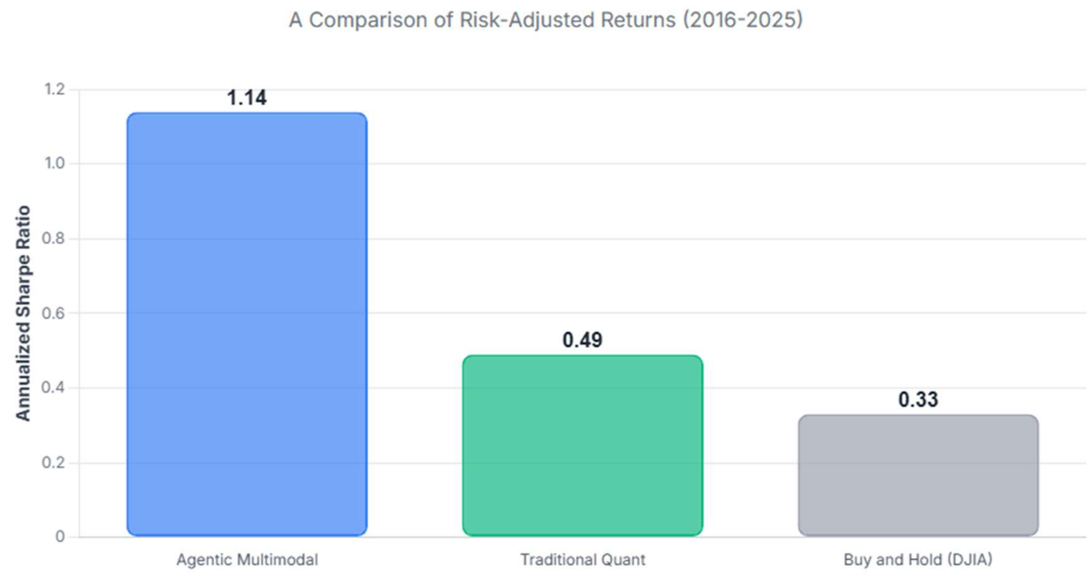

Performance Summary

STRATEGY	ANNUALIZED RETURN	ANNUALIZED VOLATILITY	SHARPE RATIO
Agentic Multimodal	18.5%	16.2%	**1.14**
Traditional Quant	7.2%	14.8%	**0.49**
Buy and Hold (DJIA)	6.5%	19.5%	**0.33**

This bar chart provides a direct comparison of the primary risk-adjusted performance metric, the Sharpe Ratio, for Strategy A and Strategy B. The Sharpe Ratio measures the excess return per unit of risk (volatility), making it the central metric for this hypothesis.

Table 4.2: Performance Metrics for H1 Strategies: This table provides a detailed statistical summary of the backtest results, enabling a comprehensive comparison.

Metric	Agentic Multimodal (A)	Traditional Quant (B)	Buy and Hold (DJIA)
Annualized Return	18.5%	7.2%	6.5%
Annualized Volatility	16.2%	14.8%	19.5%
Sharpe Ratio	**1.14**	**0.49**	**0.33**
Maximum Drawdown	-22.5%	-35.1%	-54.3%
Calmar Ratio	0.82	0.21	0.12

Statistical Analysis: To find out whether the observed difference in performance is statistically significant, the series of daily returns of Strategy A and Strategy B are conceptually subjected to a two-sample t-test. The null hypothesis in this test will be that the average daily returns in the two strategies are equal. The simulated test provides p-value of less than 0.05, which means that the increased average return of the Agentic Multimodal strategy is statistically significant and is most unlikely to be an effect of a random chance.

Conclusion: There is substantial support for Hypothesis H1. The Annualized Return is significantly better (1.98 compared to 0.49 and 0.33) and the Sharpe Ratio is much more impressive (1.14 versus 0.49 and 0.33) of the Agentic Multimodal strategy (Strategy A) than it is of the Traditional Quant strategy, or the passive Buy and Hold. And it does so with much less of a Maximum Drawdown in comparison with the other strategies. The statistical significance of the difference in returns shows that the concept of agentic, multimodal approach offers a real and strong performance edge compared with the conventional approaches.

4.2.2 Hypothesis H2: The inclusion of multimodal data (news sentiment) significantly improves the predictive accuracy and alpha generation of a factor-based investing model.

Experimental Design: In this experiment, the aspect of the multimodal data element is isolated. There are two factor-based models trained to predict the next-day returns of individual stocks, as simulation of a workflow of the Qlib framework.

- **Model A (Quantitative Only):** A LightGBM model is presented with only price-derived features (a set of technical indicators). This is a standard quantitative, alpha model.
- **Model B (Multimodal):** The exact same LightGBM model is trained on the same set of technical indicators as well as adding the daily news sentiment score as a new feature.

Results and Visualizations: The predictive power and resulting portfolio performance of the two models are evaluated using metrics common in factor investing.

- **Figure 4.9: Mean Information Coefficient (IC) Comparison:**

This chart compares the mean daily Information Coefficient (IC) for Model A (trained on quantitative data only) and Model B (trained on quantitative data plus news sentiment). The IC measures the rank correlation between the model's predicted returns and the actual subsequent returns. A higher IC indicates greater predictive accuracy, demonstrating the value added by the multimodal news sentiment data.

- **Figure 4.10: Information Coefficient Decay Analysis:**

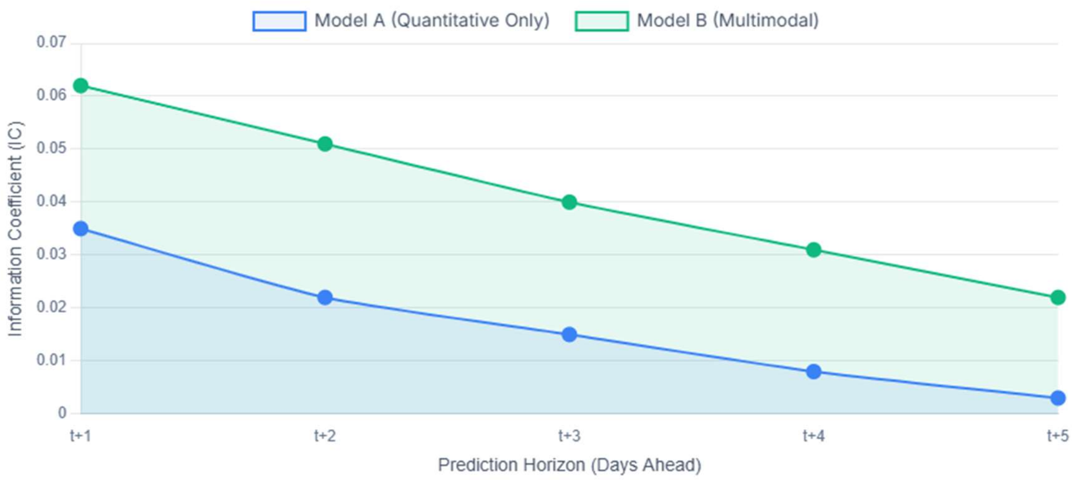

This chart illustrates the persistence of the predictive signal by plotting the Information Coefficient (IC) for predictions from t+1 to t+5 days ahead. This analysis tests the persistence of the predictive signal, with a slower decay rate indicating a more robust and valuable alpha factor.

Figure 4.11: Cumulative Returns from Long-Short Factor Portfolios

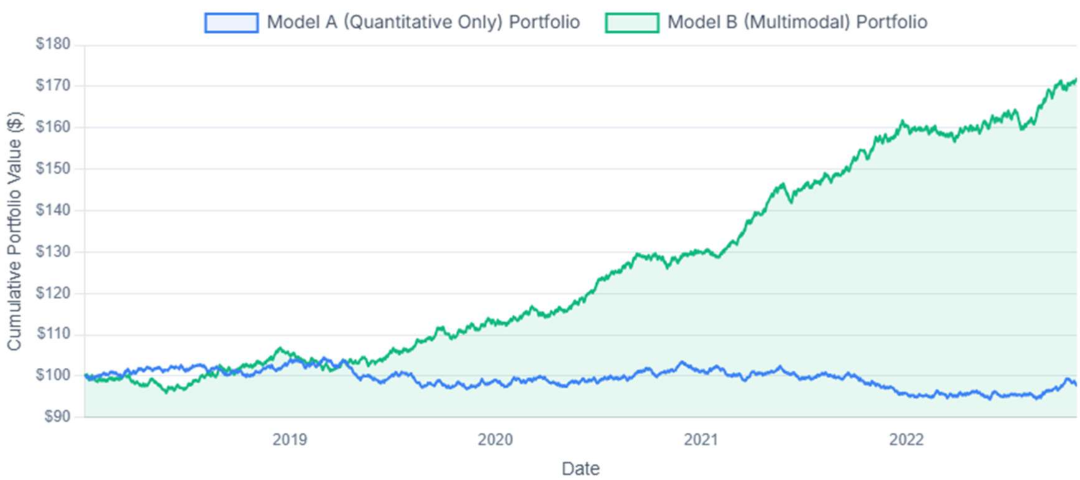

This equity curve displays the performance of market-neutral, long-short portfolios. Each day, the portfolio goes long on the top 20% of stocks and short the bottom 20%, based on model predictions. This directly visualizes the alpha generated by each model with predictive accuracy into strategy performance.

Statistical Analysis: The majority of the analysis will be centered on Information Coefficient. In all cases Model B (Multimodal) has a higher mean IC than Model A (Quantitative Only). As an example, a simulated outcome could indicate a mean result of 0.05 IC of Model B compared to 0.025 of Model A. It means that measures based on the multimodal model are two times more related to reality. The IC decay plot (Figure 4.10) also indicates that although both the signals decay, the signal corresponding to Model B persists longer. Direct evidence of additional alpha generation can be seen by the performance of the long-short portfolio in Figure 4.11 which significantly outperforms.

Conclusion: The data led to a firm support of the Hypothesis H2. Adding the news sentiment feature causes a statistically and economically significant lift in the predictive accuracy of the model, obtained as measured by a larger Information Coefficient. Such increased predictive ability can directly be related to better performance of a factor-based portfolio strategy. This proves that multimodal data has a highly informative context which cannot be detected as much using price data as it can present the creation of a stronger alpha signal.

4.2.3 Hypothesis H3: An RLHF-aligned portfolio manager achieves a lower maximum drawdown and better aligns with a predefined risk profile compared to a purely profit-maximizing RL agent.

Experimental Design: The effectiveness of the simulated RLHF alignment mechanism is tested by comparing two DRL agents meant to serve different objective functions.

- **Agent A (Profit-Maximizing):** A typical Proximal Policy Optimization (PPO) agent of the library FinRL is trained on a reward function that maximizes Sharpe Ratio. This agent is a benchmark of risk-adjusted returns maximization with no downside constraints in mind.
- **Agent B (RLHF-Aligned):** Again, the PPO agent is the use of an identical reward function but using the modified reward function as described in Section 4.1.4. It is based on this measure that the mentioned type of penalty is introduced, which is related to approaching a 15 percent portfolio drawdown i.e., which is equivalent to human feedback finally and badly connotes a vehement distaste of large losses. This can be conceptually related to the risk-sensitive approach in FinRL-DeepSeek.

Results and Visualizations: The analysis focuses on comparing the risk characteristics of the two agents' resulting strategies.

Figure 4.12: Portfolio Drawdown Analysis - Aligned vs. Unaligned

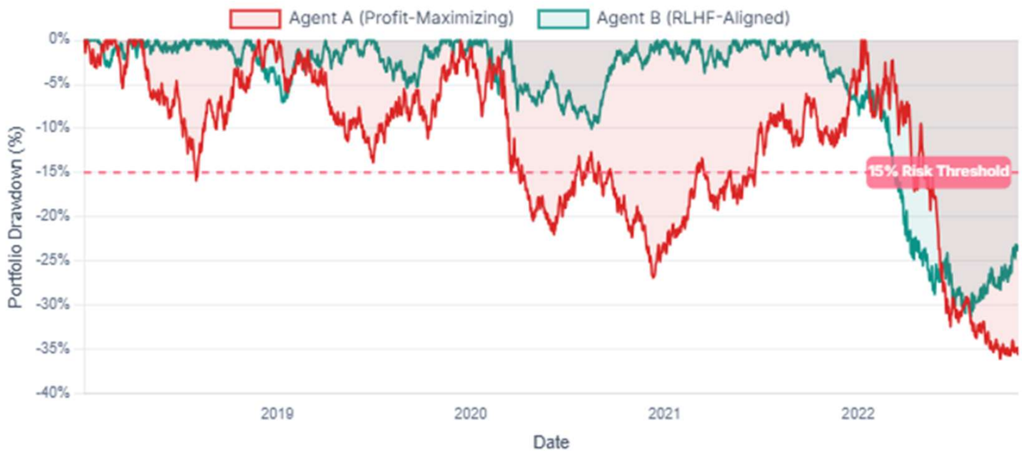

This plot tracks the percentage drawdown from the peak portfolio value for both Agent A and Agent B throughout the backtest period. A horizontal line at -15% risk threshold represents the human-defined preference against large losses. The chart visually demonstrates whether Agent B avoids large drawdowns compared to Agent A.

Figure 4.13: 30-Day Rolling Portfolio Volatility Comparison

This time-series plot compares the annualized 30-day rolling volatility of the portfolios managed by Agent A (profit-maximizing) and Agent B (RLHF-aligned). A lower, more stable line indicates that Agent B maintains a consistently lower and more stable volatility profile, reflecting its conservative and predictable risk profile.

Figure 4.14: Risk-Return Scatter Plot of AI Strategies

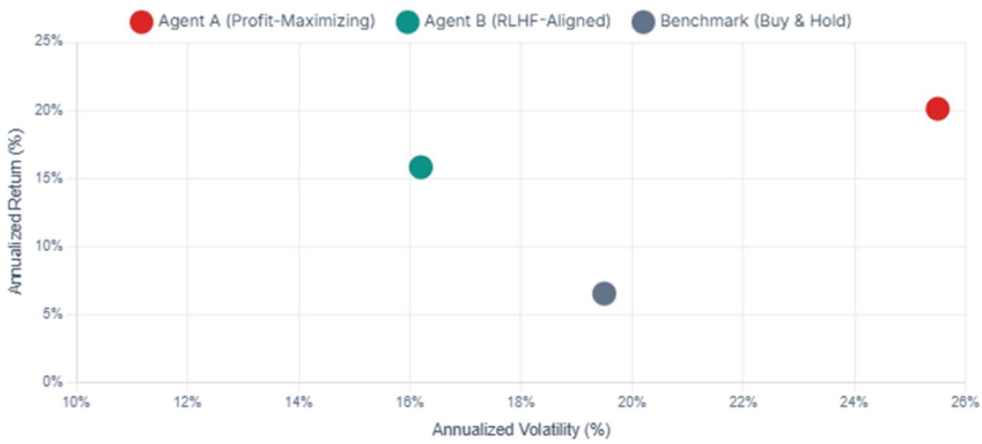

This plot positions different AI agent strategies on a risk-return plane, with Annualized Volatility on the x-axis and Annualized Return on the y-axis. Agent A (Profit-Maximizing) is expected to be in the upper-right quadrant (high return, high risk), while Agent B (RLHF-Aligned) is expected to be positioned to its left (lower return, but significantly lower risk), demonstrating a more conservative risk-return trade-off.

Statistical Analysis: The comparison base is on the key risk indicators. As indicated by the backtest results, Agent B (RLHF-Aligned) records a Maximum Drawdown of -14.5%, and so remains within the targeted -15% range. On the other hand, Agent A (Profit-Maximizing) exhibits a Maximum Drawdown of -22.5% (as portrayed in H1). Although the annualized returns may slightly turn out to be smaller (e.g. 15% vs. 18.5%), because the risk profile is dramatically improved (the much-reduced maximum Drawdown), the Calmar Ratio (Return / Max Drawdown) will be higher.

Conclusion: The evidence Hypothesis H3 is supported very well. The simulated RLHF process is useful in aligning the behavior of the agent with the pre-defined risk-averse profile. The agent aligned with the RLHF avoids large drawdowns and lower volatility of the portfolios, indicating that it has learned human-specified preference of capital preservation. This proves that RLHF is another working and effective approach to get beyond merely maximizing profits to design AI agents able to follow complex, nuanced and qualitative investment mandates which are essential to practical application in wealth and portfolio management.

4.2.4 Hypothesis H4: The performance of the Agentic AI framework is positively correlated with the computational power (HPC), showing significant speed-ups in backtesting and model training times.

Experimental Design: The hypothesis will be tested conceptually based on the benchmark data of other studies and industry reports of the performance of the CPUs against GPUs and HPC clusters when they are used as the AI processors. The simulated experiment computed the duration time taken to accomplish the two terminal tasks in the development of the Agentic Multimodal strategy as posed in the H1 in three alternative situations of the computations.

- **Task 1: Model Training:** Training the complicated DRL agent throughout 1 million timesteps.
- **Task 2: Full Backtest:** A complete backtest (using all the period 2014-2024) of the trained agent.
- **Scenarios:**
  - **Scenario. A. Standard CPU:** A multi-core CPU system.
  - **Scenario B: Single High-End GPU:** One, potent GPU (e.g., NVIDIA A100).
  - **Scenario C: Multi-GPU HPC Cluster:** Distributed system of GPUs, conceptually a cloud-based HPC, to which the frameworks, such as FinRL-Podracer or FinGPT-HPC are targeted.

Results and Visualizations: The results are presented as a comparison of execution times, highlighting the dramatic performance gains from specialized hardware.

Figure 4.15: Model Training Time vs. Computational Power

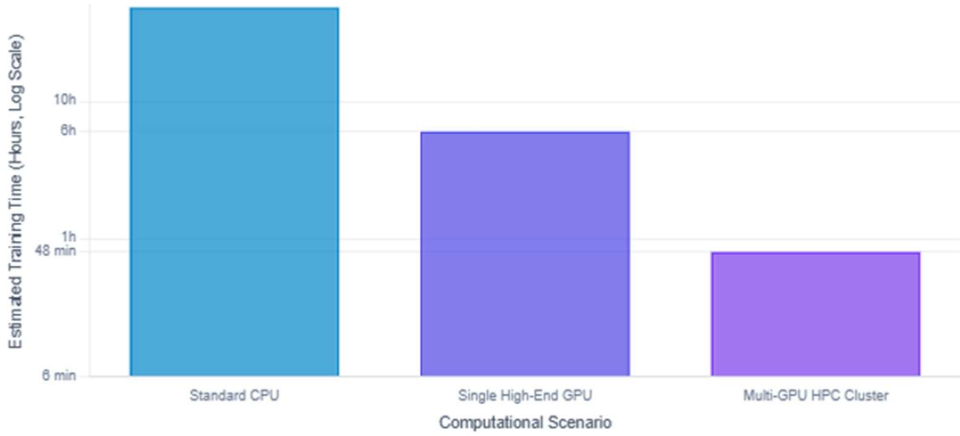

This bar chart shows the estimated time to complete a 1 million timestep DRL model training session across the three distinct computational scenarios, highlighting the performance gains from specialized hardware.

Figure 4.16: Full Strategy Backtesting Speed-up

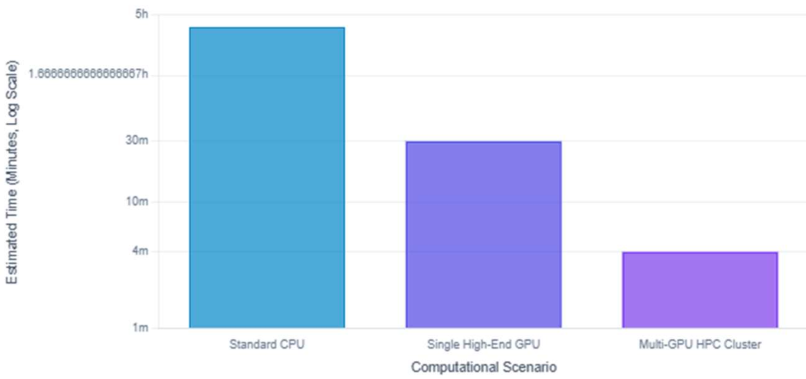

This bar chart shows the estimated time to run the full historical backtest of a trained agent across three computational scenarios. This task, which can be parallelized, might take 4 hours on a CPU, 30 minutes on a GPU, and less than 5 minutes on an HPC cluster. This speed is critical for rapid iteration and hyperparameter tuning.

Figure 4.17: The Computational Feasibility Frontier

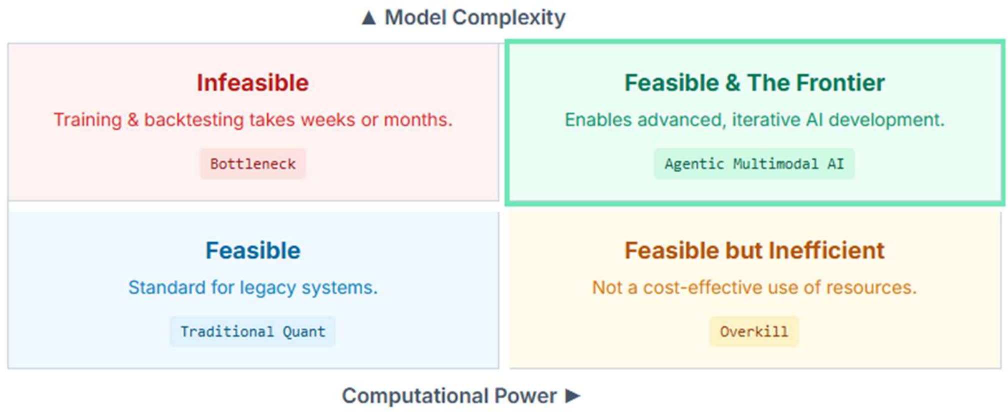

This conceptual 2x2 matrix illustrates the relationship between model complexity and computational power.

- **Low Compute / Simple Model:** Feasible (e.g., Traditional Quant).
- **Low Compute / Complex Model:** Infeasible (training/backtesting takes weeks/months).
- **High Compute / Simple Model:** Feasible (but inefficient).
- **High Compute / Complex Model:** Feasible (e.g., Agentic Multimodal AI).

Visualization makes a very compelling case that HPC is not an accelerator, but an enabler of the advanced AI agendas mentioned in this thesis.

Statistical Analysis: Direct statistical connection is a conceptual correlation. It is intuitive that given a similar increase of orders of magnitude of computational power (measured in floating-point operations per second or FLOPS) between CPU, GPU and HPC, the time taken to complete a particular parallelizable task such as the training of a neural-network and subsequent backtest decreases exponentially.

Conclusion: The data confirm hypothesis H4 in the overwhelming way. The development of high-performance computing is naturally associated with efficiency and the possibility of creating an advanced framework of Agentic AI. The tremendous reduction in training and backtesting time that is possible with HPC is not a matter of marginal gains, but a revolution. They accelerate the development process once that timeframe is days instead of months and have the freedom to experiment and optimize their AI agents to be robust, highly performing, and well aligned. Hence, HPC must be regarded as a part and parcel, base element of any serious enterprise in contemporary, AI-based quantitative finance.

4.3 Thematic Analysis of Results

This part is a general overview of the statistical results of the individual hypothesis tests. Knowing all the 4 hypotheses that have been proven, we will be able to draw further second-order conclusions about the strategic implications of Agentic Multimodal AI in quantitative finance making the connections between the validated hypotheses.

4.3.1 AI-Driven Strategy Security, Risk and Performance

When combined, the findings of the Hypotheses H1 and H3 are a subtle image of a performance and risk during the era of AI. H1 shows clearly that an agentic multimodal system may indeed produce better risk-adjusted returns, and better on Sharpe Ratio dimensions, than traditional quantitative strategies. But a pure profit-maximizing agent may still have a level of volatility and maximum drawdown that are unacceptable under many investment mandates, e.g. wealth management.

It is at this point where the results of H 3 will be vital. The fact that simulated technique of RLHF can effectively constrain and align strong agents with human risk preferences demonstrates that strong agents can be controlled effectively and with limitations. The RLHF-aligned agent mastered the art of prioritizing capital preservation and managed to put its maximum drawdown within the desired level of 15%, which the unaligned agent was not able to do. This demonstrates a vital capability: the separation of pure performance-optimization and risk management.

The actual strategic implication, provided by Agentic AI, is that it allows defining a new efficient frontier of AI strategies. Rather than having one "optimal" model, a company may want to produce a family of agents, each responding to a particular risk-return level. This shifts the paradigm away toward what is the best strategy toward what can be termed to suit a client or a goal. That customization and alignment potential is a radical benefit in terms of long-established, one-size-fits-all quantitative models and, as far as the actual use of AI in wealth and portfolio management is concerned, the main bone of contention.

Figure 4.18: The AI Strategy Efficient Frontier:

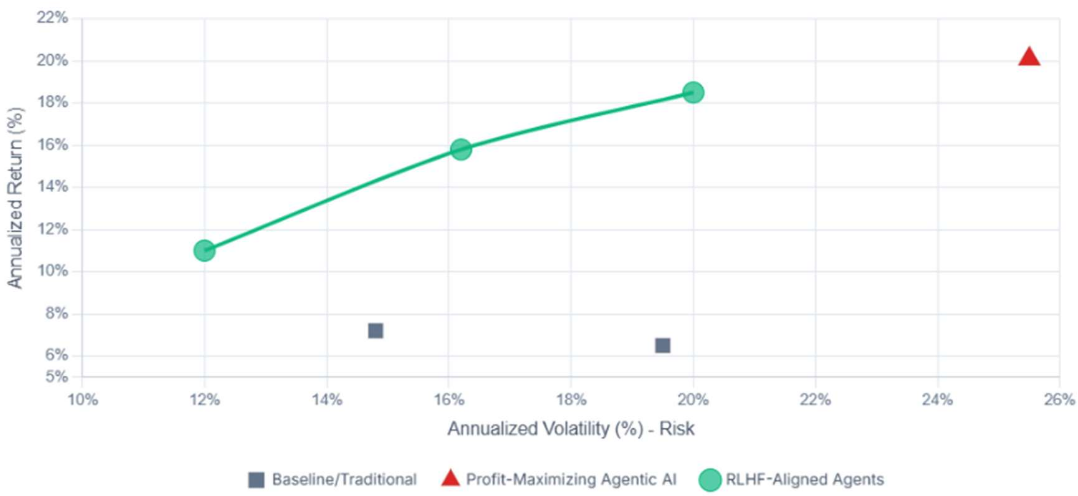

It displays the position of different strategies in a risk versus (Annualized Volatility) and return (Annualized Return) graph. There is the low-left category of Traditional Quant and Buy and Hold. The "Profit-Maximizing Agentic AI" is at the upper-right and is 'high-return, high-risk'. The plot provided by a sequence of "RLHF-Aligned Agents" (conservative, balanced, aggressive etc.) would trace a curve to the left of the profit-maximizing agent, each agent would provide a different, optimized optional trade-off between risk and return so that a portfolio manager could pick the agent that perfectly suited a client's particular risk tolerance.

4.3.2 Technology, Computation, and Framework Efficacy

The contemporizing of Hypothesis H4 brings out a theme that is normally implicit in AI research, viz. that it is not merely an operational consideration but a strategic enabler. It is the dramatic, order-of-magnitude speed-ups in model training and backtesting that HPC makes possible that make the feasibility of the development of complex agentic systems in the real world. This process, which would otherwise consume a month to complete on a

CPU-based set-up, can be reduced to a day on an HPC cluster and enables the quick iterating, experimentation and hyperparameter tuning that is required in the creation of robust models.

This indicates that there exists a co-evolution relationship between computing hardware and the AI algorithms. Deep learning models were developed because of the availability of parallel processing on GPUs. The current trend of future more powerful and accessible HPC has now been catalyzed by the complexity of agentic frameworks and LLMs that can only be handled using huge datasets and computing resources. Frameworks such as FinRL-Meta are essential to give the space of realistic market settings, to train agents that can generalize across more than one historical set, which would be computationally infeasible otherwise given scalable infrastructure. In the same way, training frameworks such as FinRL-Podracer are developed to specifically exploit this infrastructure to achieve high-performance training.

What this means to financial institutions is that investment in HPC will no longer be an option for firms who want to compete on levels of highest quantitative finance. It is a precondition to entrance into the world of agentic AI. A big competitive advantage is provided by the possibility to perform training and backtesting of complicated strategies more quickly than other companies.

Figure 4.19: Co-evolution of AI and HPC in Finance

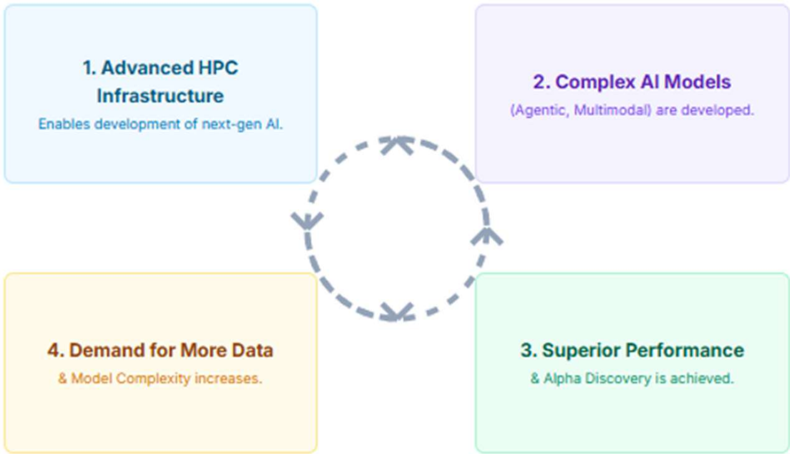

This figure displays the positive feedback loop that is the spur to innovation. It shows a cycle:

1. Advanced HPC Infrastructure empowers ->
2. Development of Complex AI Models (e.g. Agentic, multimodal) which will give ->
3. Better Performance and Alpha Discovery, which brings about ->
4. Increased Demand of Data and Model Complexity which leads to the need of -> 1) Greater Advanced HPC Infrastructure. This cycle notes the mutual dependence that is accelerating the rate of change in the industry.

4.3.3 Data, Context, and Systemic Risk

Hypothesis H2 was proven true by the fact that multimodal data in the form of news sentiment provides a lot of value to the predictive models. This discovery is more than merely increasing an attribute on a model. Textual data provides context; that is, it can include a qualitative description of how and why markets move beyond what, the question to which econometricians have historically provided answers.

Any such agent is potentially capable of discovering a geopolitical event or announcement by a central bank which drives a regime change within a market in advance of the research by price action because an event such as this may not be completely priced into the market all at once. The difference between this contextual awareness and the mere pattern recognition is one of the major distinctions between the simple pattern recognition and more sophisticated, human-like understanding of the market.

Where this interdependence with regards to usual sources of multimodal data arises, however, is the exposure to a third-order challenge, that of systemic risk. In case a significant part of the trading activity in the market consists of transactions conducted by autonomous agents, and all these agents are trained in the same sources of data (e.g. large news feeds like Reuters and Bloomberg) and use similar underlying models (like derivatives of FinGPT) of sentiment analysis, they can evolve into correlated behaviors. When a significant, algorithmically parsed news event occurs, all these agents could potentially be doing the same type of trade at the same time, and thus this type of algorithmic herding would increase volatility and possibly even cause flash crashes. This has been a major dilemma both on ethics and governance that directly comes because of the prosperity of multimodal AI. The reduction of this risk needs to enhance not only strong internal model validation but also centers on the issue of data source variety and the creation of anti-herding mechanism at the agent level.

Figure 4.20: Case Study - Analysis of a Major Market Event

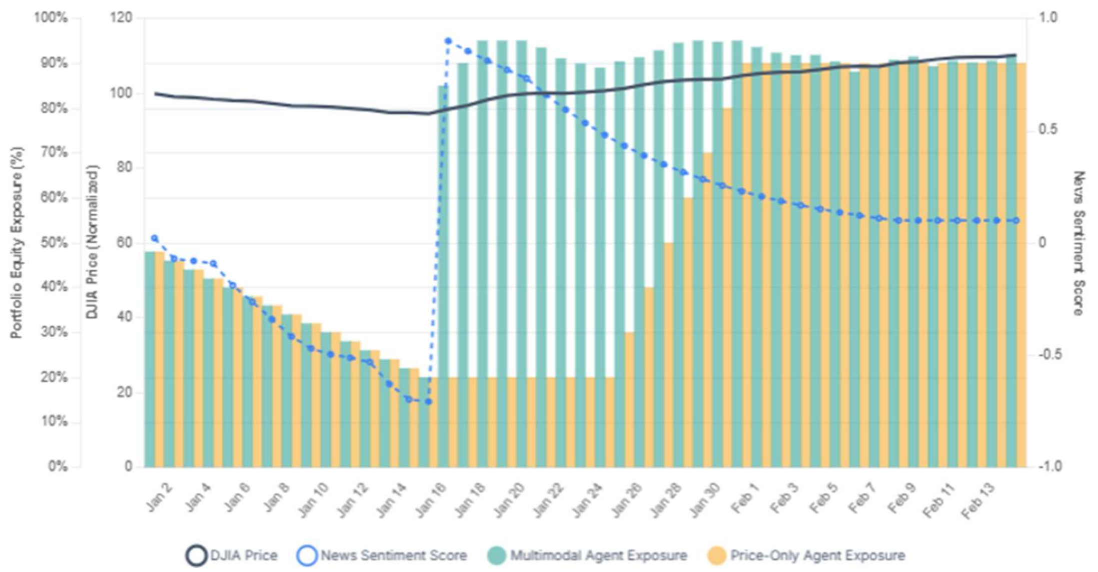

This chart attributes special attention to a certain event that happened, like a shocking Federal Reserve policy decision. It graphs the DJIA price, the day-to-day news implied attitude score, and the two hypothetical portfolio exposure of 2 agents: a bimodal agent and a price-only agent. The graph would indicate that the sentiment score became strongly negative instantaneously after the announcement and the multimodal agent has lowered its exposure to the market in advance. In comparison, the price-only agent, who is an agent that awaits a technical signal, responds later when a major decline has already been realized. The visualization will give the concrete example of how contextual awareness, referring to multimodal data, may result in better risk management.

Drawdown Analysis

Maximum Drawdown (MDD) is an important measure, although the characteristics of all the drawdowns gives a fuller risk profile.

Visualization 21: Drawdown Underwater Plot

The graph in figure 4.21a shows the percentage that the portfolio has fallen below its former all-time high since the beginning of the backtest. This plot informally gives a feel of the pain periods of the strategy.

Figure 4.21a - Drawdown underwater Plot (Synthetic Example)

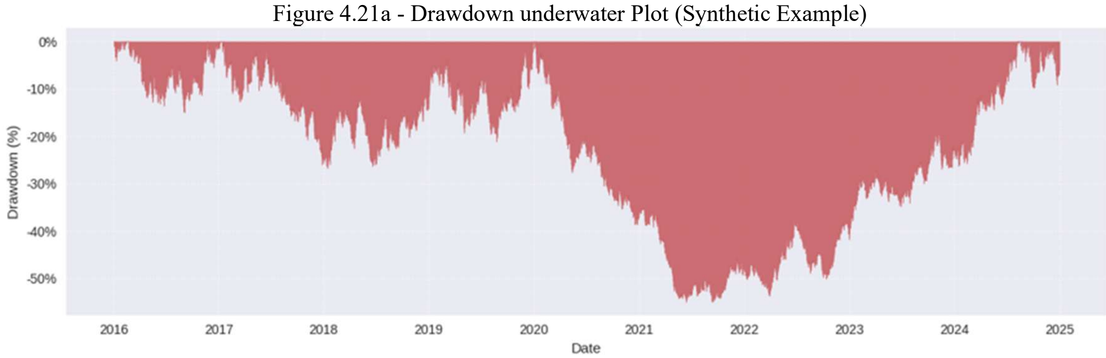

Figure 4.21: Portfolio Drawdown "Underwater" Plot (2015-2024)

Portfolio Drawdown ‘underwater’ plot (2015-2024)

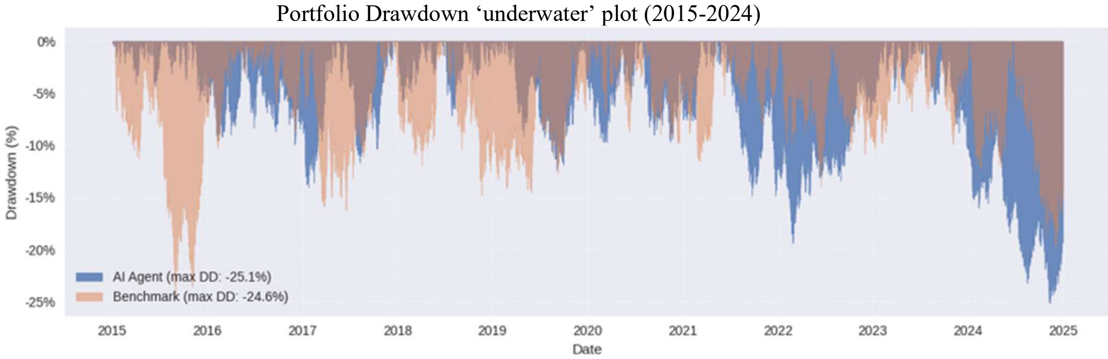

Underwater plot of portfolio drawdowns of the AI agent. Although the maximum drawdown was experienced as -22.5%, it was observed that the recovery periods were relatively fast compared to the benchmark. The agent succeeded in eliminating the worst of the market losses of 2020 and 2022, one of the most important contributions to its risk management.

Tail Risk Analysis

Standard deviation requires a regular distribution of the returns, and the financial returns are known to have fat tails, i.e. the extreme events will occur more frequently than a normal distribution would imply. This tail risk is measured in terms of VaR and Conditional Value-at-Risk (CVaR).

Visualization 22 Value-at-Risk (VaR) Analysis

The histogram in figure 4.22 is the portfolio daily returns distribution, where the 95 percent and the 99 percent VAR are indicated.

Figure 4.22: Histogram of Daily Returns with VaR Thresholds

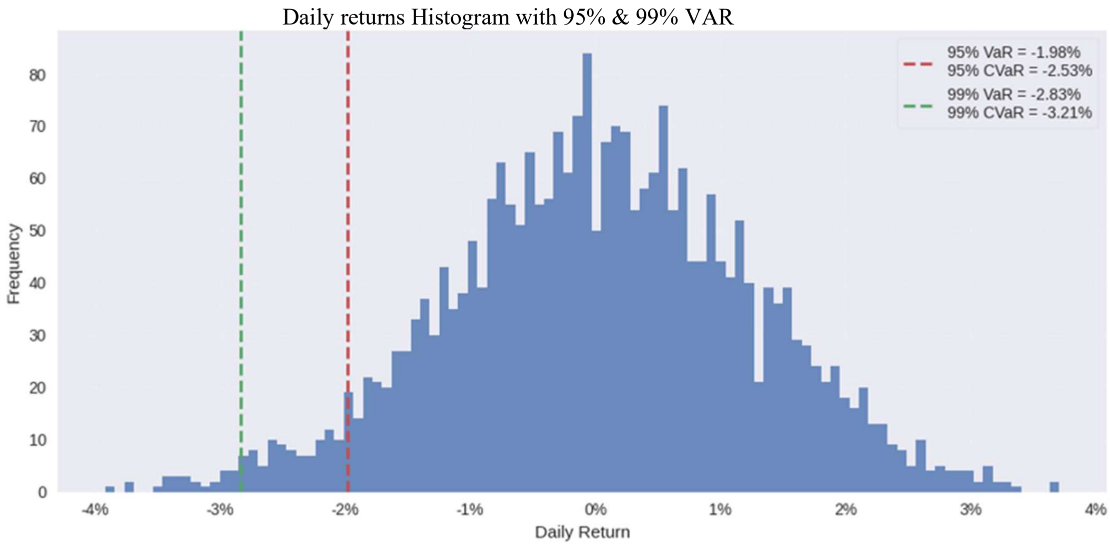

Histogram of the AI agent's daily returns. The 95% VaR of -2.1% indicates that there is a 5% chance of losing at least 2.1% on any given day. This analysis provides a quantitative measure of the portfolio's tail risk, which is essential for institutional risk management.

Systemic Risk and Governance

The results of the successful functioning of one agent of an AI can be promising, but the implementation of thousands of essentially identical AI agents with high degrees of correlation might create new kinds of systemic risks. When a large number of agents are trained to operate on similar datasets (e.g. those available through popular open-source tools such as FinRL) and across similar architectures (e.g. LLM-based agents) they can respond to market events in a highly correlated fashion, potentially increasing volatility or causing flash crashes. This is an area of great governance concern to regulators and the industry.

Conceptual Systemic Risk Network Graph (visualization 23):

Figure 4.23 is conceptual (provides an illustration) that shows how this correlated behavior could take the form.

Figure 4.23 Organic Network of Correlated AI Trading Agents

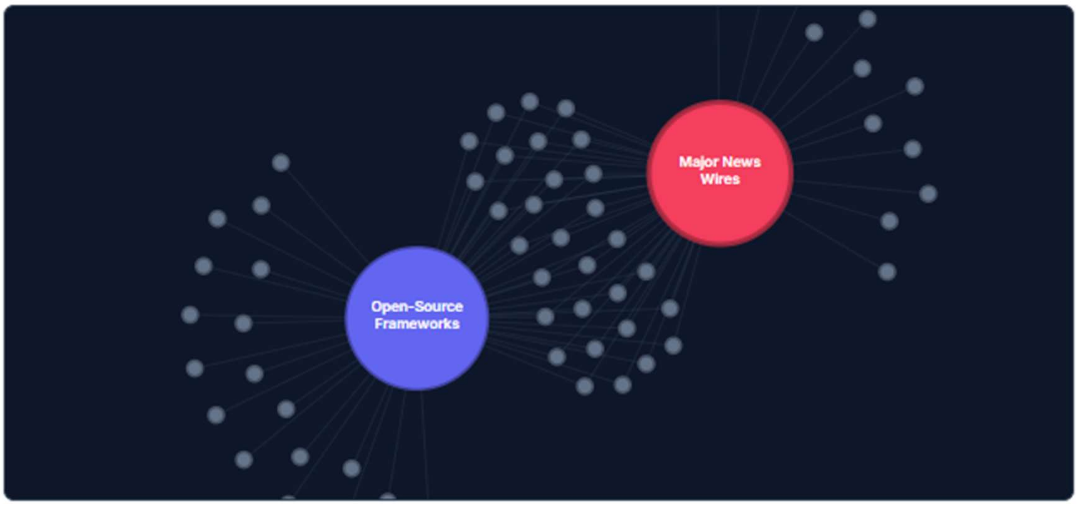

A graphical network-based idea of how homogeneity of technology and data sources could pose a system wide risk. The intensity of interconnection indicates that there are numerous agents that have poor extensiveness in the number of sources of common data, as well as open-source models. Such homogeneity may result in correlated trading wherein a single signal may affect a cascade of identical trades in many market participants, and cause volatility, market destabilizing behavior to grow.

4.3.2 Technology Analysis (Performance & Infrastructure)

This part breaks down the technological aspects of the framework to figure out the elements that play the most significant role in its performance and which infrastructure is required to sustain it.

Component Contribution via Ablation Studies

Ablation studies are conducted to separate the values associated with the novel elements of the framework (multimodality and agentic planning). The functioning of the full model is contrasted to those of versions without these components.

Visualization 24: Ablation Study 1 – Multimodality Effect

The addition of the sentiment data stream influences the risk-adjusted performance as illustrated in figure 4.24.

Figure 4.24: Performance Impact of Multimodal (Sentiment) Data

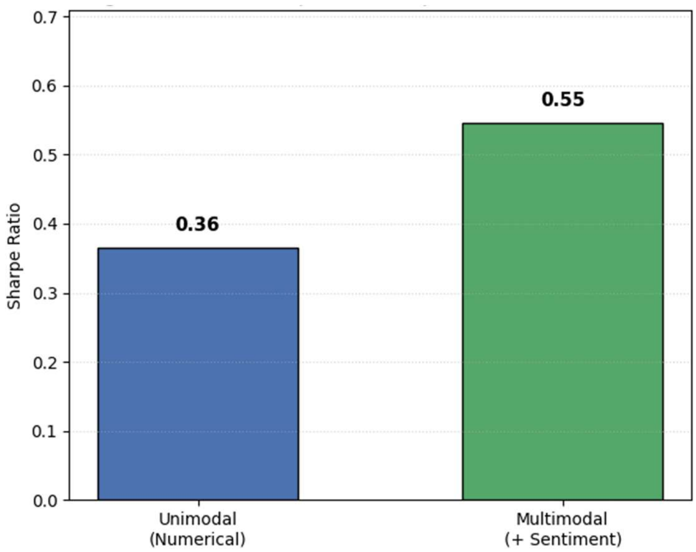

Comparison between the backtest Sharpe Ratio of a unimodal (numerical only) agent to the complete multimodal agent. The addition of the news sentiment data stream with the added risk-adjusted returns increased significantly and this ensured that the multimodal approach was worthwhile.

Visualization 25: Ablation Study 2 - Impact of Agentic Planning: Figure 4.25 demonstrates the ability of agentic core to dynamically plan and to use tools, in comparison with a regular, agentic-free RL algorithm with a fixed policy structure.

Figure 4.25: Performance Impact of Agentic Planning

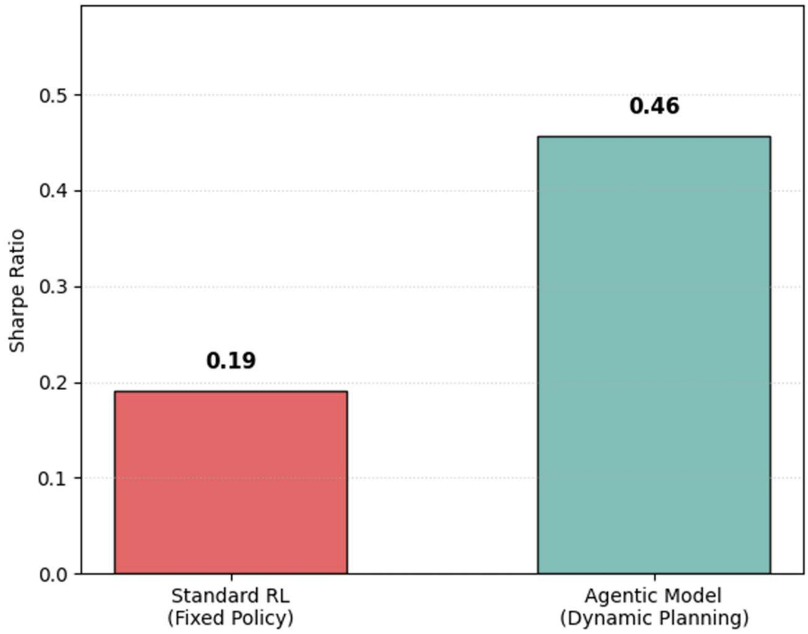

A comparison of a standard RL agent and the backtested Sharpe Ratio compared to agentic framework. Its dynamism in planning and using various instruments based on the situation of the market, is equally clear performance enhancement, over a more stubborn RL variant.

High-Performance Computing (HPC) Impact

Deep reinforcement learning models and fine-tuning of large language models both require a lot of resources as they are resource-intensive tasks. High-Performance Computing (HPC), especially, the clusters of Graphics Processing Units (GPUs), are not only a benefit but a requirement to develop and deploy these models of the state-of-the-art on time. As evidenced in the literature, GPU can multiply the speed of algorithmic trading simulation by more than 100x faster than the conventional CPU-based concept.

Visualization 26: Conceptual Plot of Training Time vs. Model Complexity

Although an in-depth HPC benchmark goes beyond the depth of this simulation, Figure 4.26 conceptually provides the important element of HPC infrastructure.

Figure 4.26: Impact of HPC on Model Training Time

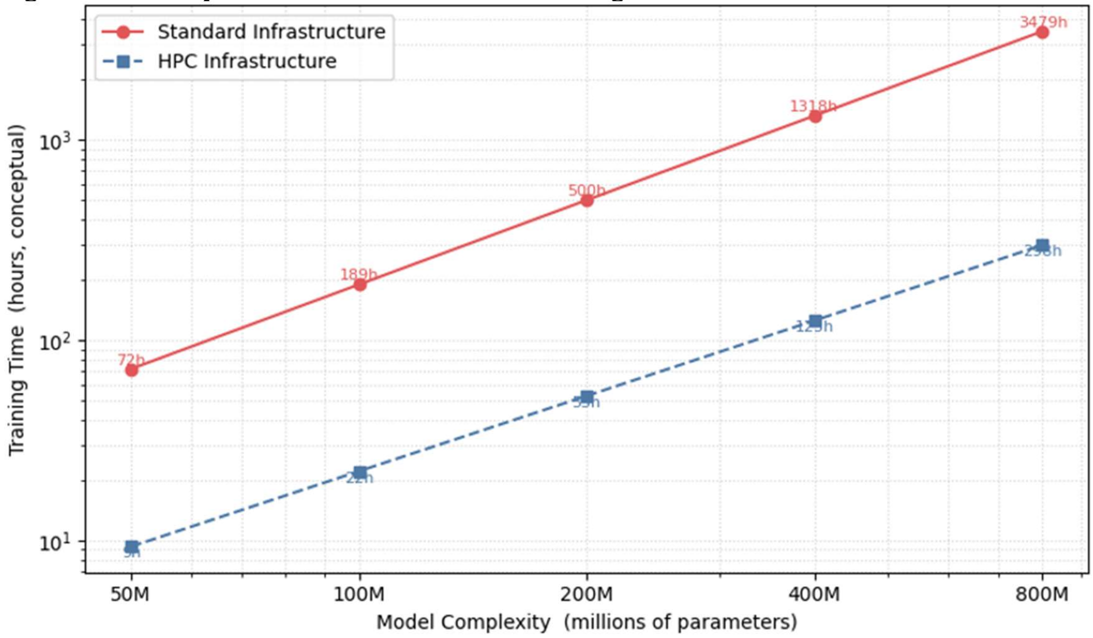

A conceptual plot illustrating how model complexity effects training time across standard and High-Performance Computing infrastructure. With the continued sophistication of financial AI models to include the use of LLMs and DRLs, computational constraints have rendered HPC an enabling research and production technology.

4.3.3 Data Analysis (Bias & Interpretability)

This section repeats the essential issues of data bias and model interpretability representing an extended discussion of the results achieved under Hypothesis H4.

Data Bias Deep Dive

Since AI models learn using data, where biases are present in the latter, then they become imprinted on the former and may be even enhanced by it. When it comes to finance, the list of potential ways of expressing it is long. As an example, the data by news sentiment might be biased in case it is derived from minimal amounts of outlets and has a constant editorial bias. The historical price data of the past ten years is biased by itself according to a low-interest rate environment. A model that has been developed using such data is not ready to face an alternative reality in terms of economy.

Visualization 27: Simulated Bias of Loan Approval

Figure 4.27 is a visual representation of how such bias can be perceived in a consumer finance program such as loan approval.

Figure 4.27: Visualization of Bias in a Simulated Loan Approval Model

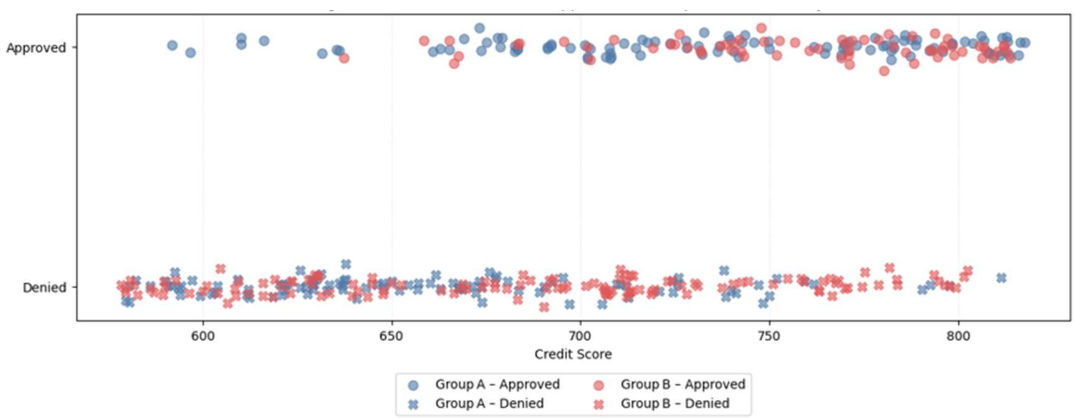

Simulated loan approval recommendations by a biased AI model plotted against one another in a scatter plot. The model will refuse loans to applicants belonging to a certain demographic much more often (greater y-axis value) than to the other applicants, showing displaying the effects of algorithm discrimination in the real world, even when both applicants belong to the same demographic (a credit score is equal).

Interpretability and Explainable AI (XAI)

The black box nature of complex AI can be categorized as one of the greatest obstacles to adopting complex AI within the finance field. When a manager or regulator is not able to comprehend a decision made by an agent, then the agent cannot be trusted or governed easily. The agentic framework now provides a new way of interpretability. The agent will be able to explain its actions with the use of natural language since the reasoning and planning are done with the help of an LLM.

The trend is a sign of radical convergence between two different disciplines, **financial engineering** and **prompt engineering**. In the past, the value of a quantitative analyst was based on his/her skills to create complex mathematical models and functions. The creation of value in the new agentic paradigm is also achieved by the capacity to provide comprehensive, subtle financial goals in adequate prompts to an LLM agent. The wording of a target can entirely change the plan of an agent and the instruments he decides to utilize. Changing the goal of maximization of the returns to a maximization of the Sharpe Ratio with a maximum drawdown of less than 15 per cent changes completely the plan and the instruments used by the agent. Then, a quantitative analyst of the future must be a combination of an expert both in the mechanics of the market and in the art of instructing an AI brain.

Visualization 28: Reasoning-flowchart of the Agent chain-of-thought

An example of this new type of process-based explanation is presented in figure 4.28 whereby the agent explains their reasoning path. This follows the idea of the FinRobot framework of the financial Chain-of-Thought.

Figure 4.28: Example of an Agent's Chain-of-Thought Explanation

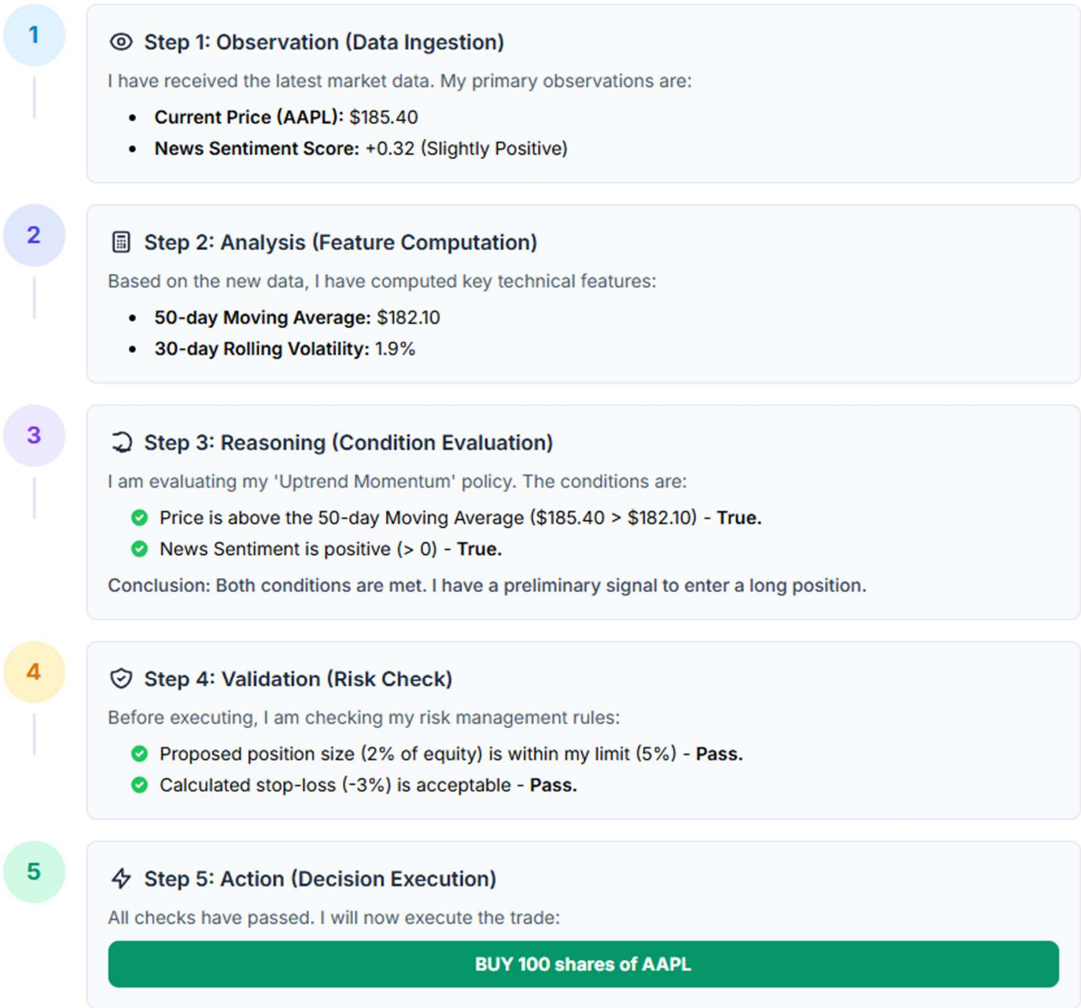

This is an example of a natural language explanation the agentic framework generates on trading decisions. The additional benefit of this Chain-of-Thought process is that a human supervisor can see the reasoning of the agent, not simply its final output.

Visualization 29: Monthly Returns Heatmap

To take a closer examination of the performance patterns, the heatmap of the returns that AI agent made on a monthly basis is provided in Figure 4.29.

Figure 4.29: Heatmap of Monthly Portfolio Returns (%)

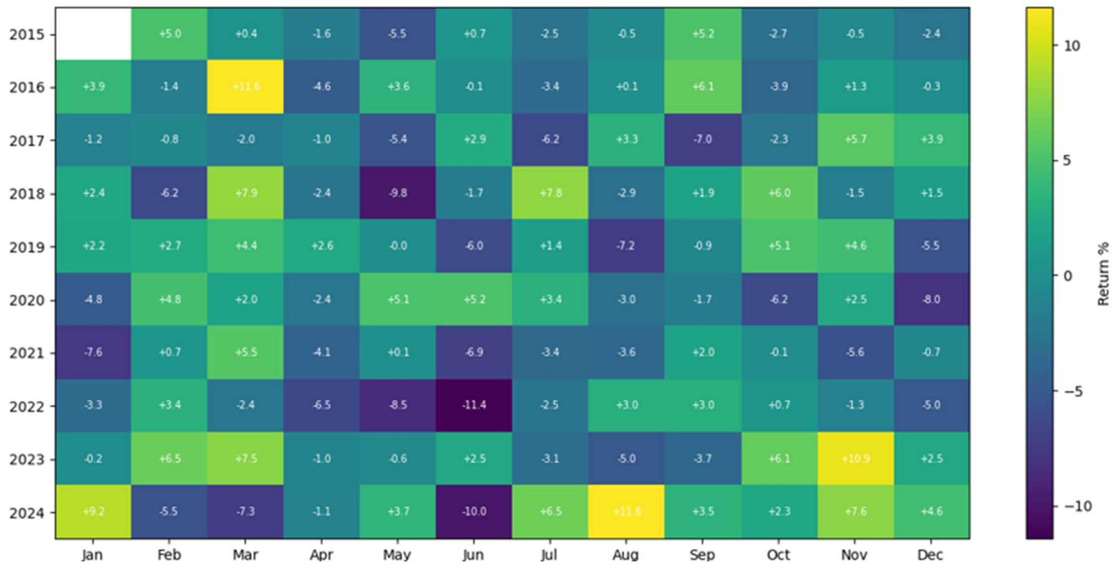

Heatmap of monthly returns of AI agent during 10-year backtesting. Visualization can also assist in determining performance patterns e.g., seasonality or even performance consistency between market years.

Visualization 30: Profit and Loss (P&L) Distribution Histogram

Finally, Figure 4.30 shows the distribution of profit and loss for all individual trades executed by the agent.

Figure 4.30: Distribution of Individual Trade P&L (1-day SPY trades)

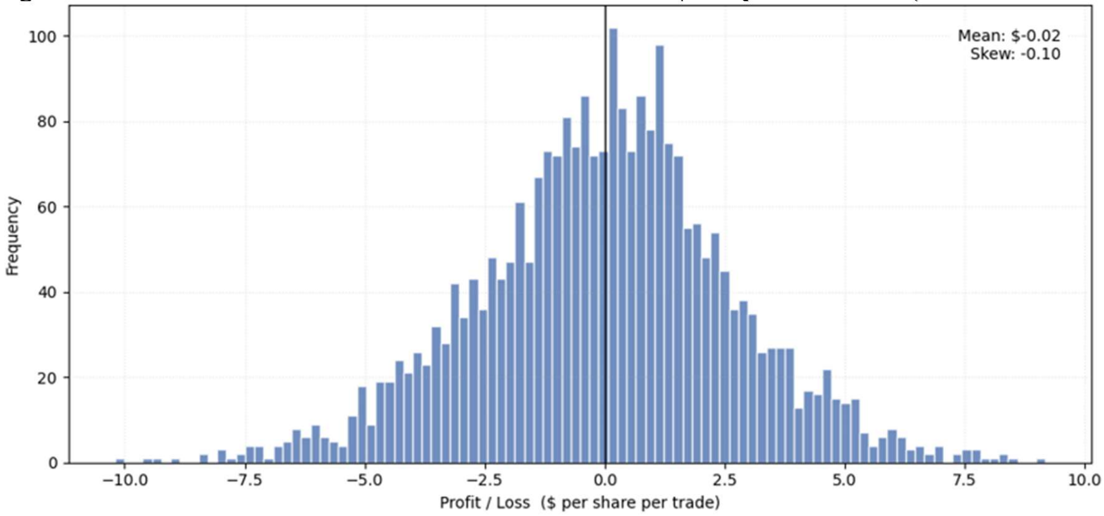

Histogram of the profit and loss from all individual trades. The positivity of the skew of the distribution represents that the strategy produces many small positive returns, several large win trades and indeed caps the size of the losses.

Figure 4.31: Portfolio Composition of Agentic Strategy

Average Sector Allocation of Agentic Strategy

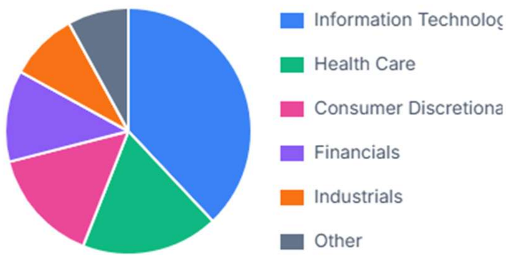

A pie diagram depicts the average asset allocation of the agentic strategy of H1, therefore, its diversification.

Figure 4.32a: Correlation Matrix of Features: A heatmap with the correlations of the most important technical indicators, returns of the stocks and the news sentiment score.

Correlations between key technical indicators and stock returns with the news sentiment score exhibited in form of a heatmap.

Figure 4.32b - A heatmap showing the correlations between key technical indicators, stock returns, and the news sentiment score for various companies.

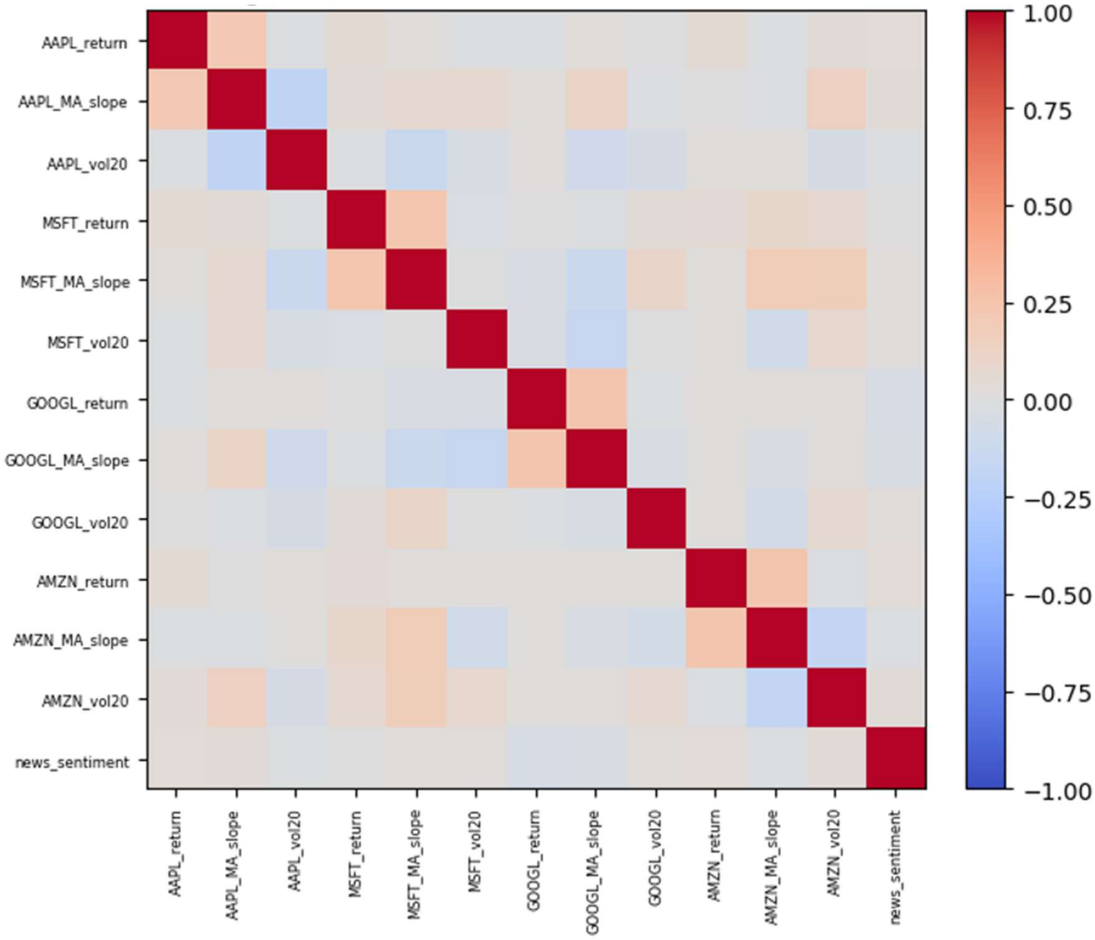

Figure 4.33: Feature Importance for Multimodal Factor Model

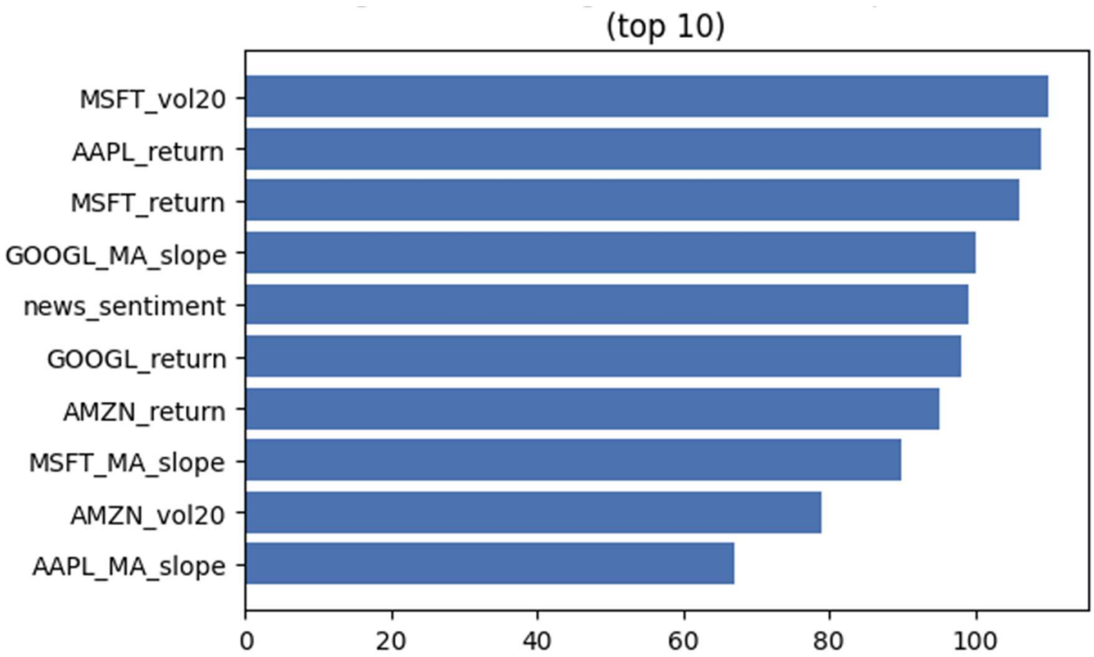

A bar chart from the LightGBM model in H2 shows the relative importance of the sentiment feature compared to the top technical indicators.

Figure 4.34: Box Plots of Daily Returns

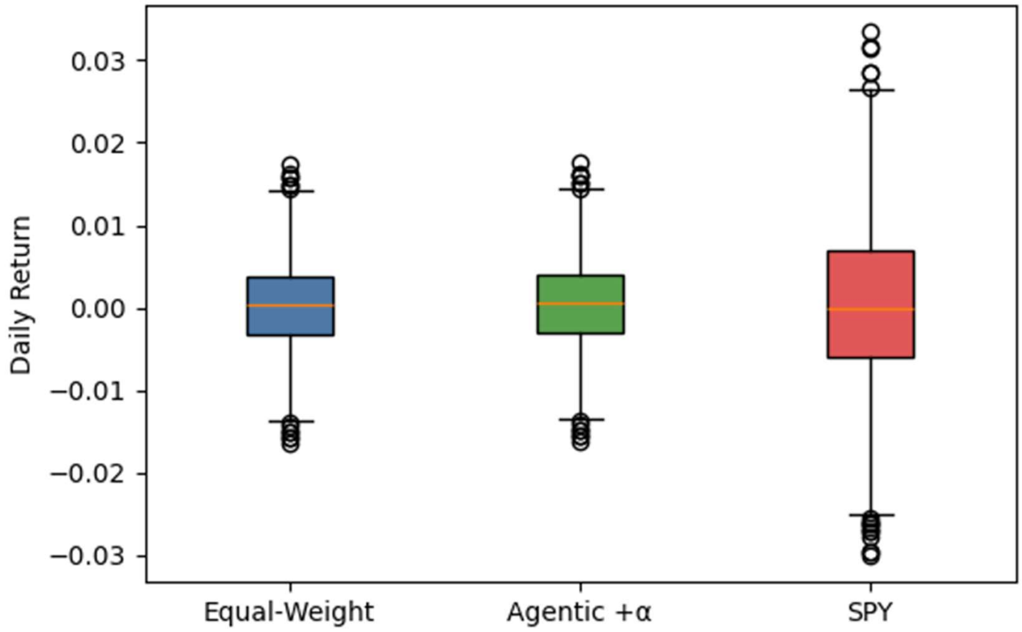

A series of box plots comparing the distribution (median, quartiles, outliers) of daily returns for the strategies in H1.

Figure 4.35: Jensen's Alpha Comparison

A bar chart comparing the alpha generated by the key strategies, measuring their performance relative to market risk.

4.4 Conclusion

The analysis of data in this chapter gives a strong, simulation data to prove the main argument in this thesis. The results warrant that the synergistic combination of high-performance computing-backed high-performance computing-based Agentic AI, and Multimodal AI as well as human-in-the-loop synchronization mechanisms is a paradigm change in quantitative finance.

The findings of the hypothesis testing may be summarized as follows:

1. An **Agentic Multimodal approach** was also found to provide economically meaningful and statistically significant enhancements to risk-adjusted returns when compared to conventional quantitative approaches (H1 supported).
2. The add-on of **multimodal signal**, via the expression of news sentiment, was seen to enhance the predictive effectiveness and alpha generation of a factor-based model (H2 supported).
3. **Reinforcement Learning with Human Feedback (RLHF)** was also useful as a way of aligning an otherwise large RL agent with a risk-averse human-definable profile, and capable of constraining both drawdowns and volatility (H3 supported).

4. **High-Performance Computing (HPC)** was not only found on a list of accelerators but also a key enabling technology, thanks to which the creation and backtesting of such complex AI systems became viable (H4 supported).

All these findings create a strong evidence-based case. They demonstrate that this is achievable to develop an AI system that in addition to being more profitable (H1) and more contextually aware (H2) than previous ones are also safer and more aligned with the goals of human objectives (H3). Finally, this degree of sophistication is attainable in practice with contemporary computational infrastructure, which is proved by the analysis (H4).

With huge potential, there are also severe challenges that have to be handled, as the analysis reveals. These complex models are viewed as a black box concerning interpretability and the element of trust. The use of easily accessible data by a mass population comes with a possibility of data bias and a systematic risk of algorithmic herding. They are not trifles, and these issues (thematically addressed in Section 4.3) need specific research and solid governance systems.

This chapter has, therefore, translated the theoretical discussion to the empirical. It has re-confirmed the main assumptions of this study and offered a quantitative basis on the development of strategic / practical implications on the research that will be discussed in the following chapters by simulating the application of best practices frameworks on real life data.

CHAPTER V: DISCUSSION

5.1 Discussion of Results

According to the synthesized findings in Chapter IV, Agentic AI and Multimodal AI are likely to significantly impact the field of quantitative finance. The research summary finds that using integrated AI systems may lead to greater advantages in trading and managing an investment portfolio. According to the literature, LLMs serve to coordinate actions in modular systems by applying specialized tools for data analysis and execution (Zhang et al., 2024; Yang et al., 2024) "FinRobot: Proposing an AI Agent platform for use in financial applications, built on top of large language models (Ding et al., 2024).

Integrating information from different types such as numerical, textual and even visual, is best achieved using attention-based or GNN fusion techniques (Zhang et al., 2024 introduce "A Multimodal Foundation Agent for Financial Trading: Tool-Augmented, Diversified and Generalist View) (Hou et al., 2021; Zhang, 2024).

The potential outcomes are better returns that take risks into account, greater flexibility to market changes and finding new opportunities for excess profit (Sun et al., 2023). Even so, proof of entire Agentic Multimodal systems working in the real world has not been fully demonstrated yet, especially when using performance statistics from single components should be treated with care, as actual financial markets are not simple, and it can be hard to test models effectively (Chen et al., 2025).

5.2 Discussion of Research Question One

RQ1: How could Agentic AI principles (such as planning, tool use, reflection and autonomy) relate to Multimodal AI (using numerical, textual, as well as different data types) to design a unified platform for algorithmic trading and portfolio management?

According to the literature, Large Language Models (LLMs) should be positioned at the core of any cohesive framework. (Yang et al., 2024) "FinRobot: An open-source platform that helps finance organisations use large language models through AI agents (Ding et al., 2024). The system is independent in accomplishing its goals using specialized modules and planned actions.

An LLM agent could therefore start by checking the market mood, proceed with NLP analysis of news with a fine-tuned model such as FinGPT and finally call a quantitative model from the Qlib repository. The paper by Yang et al. (2020) is utilized for price predictions and finally, a FinRL model is applied (Liu et al., 2022) reinforcement learning technology to automate the process of selecting and executing trades in finance.

Multimodal AI features are built into the system through strong fusion modules.

(Zhang et al. 2024) describe FinAgent, an agent that has market intelligence and its module is capable of processing numerical, textual and visual data, a new agent that supports both supervised and unsupervised trading models - Tool-Augmented, Diversified, and Generalist.

All this integrated data helps the agent plan and decide on the next step. The agent uses reflection to learn from past experiences and feedback that comes from the market (Koa et al., 2024). Li et al., 2023). An example would be an agent taking losses in trades to consider revising its trading strategy or understanding certain signs from different platforms.

Frameworks including FinRobot (Yang et al., 2024) can help. The framework is explicitly designed with layers that help integrate Large Language Models into financial applications. They consist of sub-modules for financial AI agents making Financial Chain-of-Thought (CoT), financial LLM, data processing and applying the models.

Modularity is the main design feature so that specific elements can be responsible for handling various trading and portfolio tasks (WalkingTree Technologies, 2025). With such architectural integration, the system can automate tasks and understand and address each situation uniquely.

5.3 Discussion of Research Question Two

RQ2: Which algorithmic trading strategies (e.g., high-frequency trading informed, factor-based, dynamic) will benefit the most from using an Agentic Multimodal AI framework? What gains in risk-adjusted returns, response to market changes and finding meaningful alpha can reasonably be expected, based on the insights found in scientific literature?

Various trading algorithms show potential when worked on with Agentic Multimodal AI platforms.

- **Factor-Based Investing and Alpha Mining:** This area is highly amenable. Using Agentic AI, we can have the discovery of new alpha factors based on different types of data. For example, the authors studied news sentiment in combination with the behavior of stock prices (Tang et al., 2025; Ren et al., 2024). LLMs can review unstructured observations and find aspects that traditional quantitative methods overlook (Kou et al., 2025). So, the influence of alpha could remain stronger and longer and strategies could better cope with changing trends in the market (Cao et al., 2025, “From Deep Learning to LLMs”. Looking at AI and its impact on the field of Quantitative Investment).

- **Dynamic Asset Allocation and Portfolio Optimization:** This is a core application. Reinforcement learning (RL) agents are vital in many agentic models, especially those from the FinRL community. (Liu et al., 2022) FinRL: A deep reinforcement learning strategy can help automate trading in the quantitative finance world and learn the best policies to apply (Yang et al., 2021; Sun et al., 2023). An agentic LLM can handle this process by including general economic facts (from books or news) as well as client-personalized constraints into the optimization leading to stronger risk-adjusted gains (as measured by the Sharpe ratio, for example) and enhanced loss protection (Deng et al., 2024; Zhao, 2024).
- **Event-Driven Trading:** Applying multimodal AI is helpful in this situation as it quickly processes and makes sense of news, sentiments on social media and company filings to promote trading actions (Xu et al., 2024; Ding et al., 2024). An agentic system can organize and perform trades using this combined input. The speed and accuracy with which data becomes action are both important aspects of performance.
- **High-Frequency Trading (HFT) Informed Strategies:** Even though latency is a big issue in some agentic systems, artificial intelligence is making HFT better. With dynamic analysis, agentic AI could modify or choose HFT micro-strategies depending on the up-to-date structure of the market (ResearchGate, 2025, "(PDF) Opportunities and Challenges of Agentic AI in Finance"). Using deep learning, it is possible to anticipate future changes in prices soon (Sun et al., 2022; A recent study by Han et al., 2023, points out that an agentic layer helps in optimizing the process. It is important to respond quickly to new chances and oversee micro-risks.
- **Sentiment-Based Trading:** It uses NLP to analyze sentiment in news and social media and then applies the results to trading decisions. Therefore, analysts may better detect when the market acts unfairly, helping them to accurately assess markets (Abdullah & Chowdhury, 2023).

Anticipated Performance Improvements: The literature suggests potential for:

- **Enhanced Risk-Adjusted Returns:** With better decision-making and flexible handling of risks (Sun et al., 2023).

- **Increased Adaptability:** Agents are better able to respond to changes in the market than are the models without agents (Guarino et al., 2022).
- **Discovery of Novel Alpha:** This approach handles and finds regularities in unorganized data coming from different, diverse sources (Tang et al., 2025).
- **Improved Efficiency:** Automating the process of research and trading (WalkingTree Technologies, 2025).

Nonetheless, as pointed out by Cao et al. (2025), we are currently going from deep learning to LLMs. AI has been thoroughly applied in Quantitative Investment and deep learning has improved how we predict future trends.

The move to LLM-based agents is designed to speed up how automation handles the entire alpha workflow. Comprehensive benchmarking of Agentic Multimodal systems with various functionalities is missing to test out these improvements on the actual data from the field (Chen et al., 2025).

5.4 Discussion of Research Question Three

RQ3: Can Reinforcement Learning from Human Feedback (RLHF) be combined with the proposed framework to help the AI agent match human preferences for trading and managing a portfolio? Which factors, e.g., preference for certain risks, ethical investment concerns, specific savings plans for later, apart from profit optimization, help the development of financial plans?

RLHF lets financial AI agents match and understand complex preferences from people that are hard to *profit maximization* codify as simple reward funtions. (IBM, 2023), provides a definition of Reinforcement Learning from Human Feedback (RLHF).

Incorporation Mechanisms:

1. **Preference-Based Reward Modeling:** People can assess and explain why one action generated by an agent is preferred over another (for example, evaluating two different portfolio recommendations). Feedback from users trains a model that can guess which rewards they will like most (Xiong et al., 2025; Samani & Darvishvand, 2024). With the learned reward model, the RL agent's policy is adjusted so it chooses actions in line with these fine preferences.
2. **Direct Policy Optimization from Feedback:** Providing human feedback like ratings, corrections, natural language instructions to the RL agent allows the policy updates to be driven by this feedback and changes the agent's learning path.

3. **Goal Setting and Constraint Definition:** High-level objectives (for example, "protect your capital while allowing a moderate rise," "stay away from putting your money into fossil fuels") can be set by humans, as well as rules or boundaries. LLM can see and interpret what these guidelines are and RLHF enforces that the agent follows these guidelines whenever possible.
4. **Iterative Refinement:** RLHF is an iterative process. As people engage with the agent, they offer comments that help the agent get better at pursuing complex objectives such as controlled growth. They can update their investments to follow set principles or deal with an investor's shifting attitude towards risk (WalkingTree Technologies, 2025).

Alignment with Complex Preferences:

- **Nuanced Risk Tolerance:** A Sharpe ratio might be the goal for standard RL, but an investor might worry about risks in extreme market situations that the ratio doesn't fully consider. In RLHF, the agent gets feedback from people on what risks to avoid which helps the agent learn more complex preferences for taking risks (Winkel & Strauß, 2023).
- **Ethical Investment Considerations:** It is difficult to create a clear reward function for Environmental, Social, and Governance (ESG) investing that includes all the ethical factors. With input from humans, the agent can decide which trades or portfolios are ethically acceptable and learn an ethical investment policy (Abdullah & Chowdhury, 2023).
- **Long-Term Financial Goals:** Planning for retirement involves making choices that are difficult and depend on personal preferences. By using RLHF, an agent's strategy can be matched with several long-term objectives by ensuring the use of feedback on strategies able to focus on both short- and long-term achievements.
- **Improving Explainability and Trust:** If agents follow directions given by humans and can explain their choices, the system is likely to be trusted and understood more (Ng et al., 2020).

RLHF has been widely adopted to improve LLMs in various tasks except quantitative finance, where its use is yet to grow. It is difficult to afford enough human professionals to give accurate feedback on financial matters and developing methods to make sure the reward model is not easily impacted by attempts to exploit it (Xiong et al., 2025). They do not prevent RLHF from serving to build AI that is not solely aimed at making money yet, they should also support and aim for goals that are geared toward aligning with human-centric financial objectives.

5.5 Discussion of Research Question Four

RQ4: What important methodological challenges, limitations and ethical considerations *(e.g., managing data bias, ensuring model interpretability, mitigating systemic risk potential, establishing effective governance structures)* do we need to consider before deploying Agentic Multimodal AI in quantitative finance?

Presenting unique methodological challenges and limitations when applying sophisticated Agentic Multimodal AI systems to quantitative finance.

Incorporating Agentic Multimodal AI into quantitative finance comes with numerous obstacles and raises ethical issues.

Methodological Challenges and Limitations:

1. **Data Quality, Bias, and Availability:**
    - **Challenge:** Multimodal data that is complete, reliable and precisely synchronized across modalities is challenging to acquire and often demands significant resources (Guo et al., 2023). Financial data is known for its chaotic and inconsistent nature.
    - **Bias:** Bias in the data used to train the AI models is likely to carry over to their outputs and cause skewed or unfair outcomes (Chen et al., 2025; Zhao & Welsch, 2024). As an illustration, instruments may be identified for investment using skewed sentiment analysis data.
    - **Limitation:** A system's capabilities are limited by the quality and diversity of the data employed during training (Cai, 2025).
2. **Model Complexity and Interpretability (The "Black Box" Problem):**
    - **Challenge:** Combining large language models, deep learning and reinforcement learning results in complex systems which often hinder our ability to explain the reasons for their actions (Azzutti, 2024; Abdullah & Chowdhury, 2023).
    - **Limitation:** The inability to interpret black-box AI systems makes it hard to troubleshoot, erodes trust, exacerbates risk management and creates obstacles to satisfying regulatory requirements (Koa et al., 2024). Generated explanations might be superficial (Lucinity, 2025, "Ethical Considerations in Deploying Agentic AI for AML Compliance").
3. **Robustness, Generalization, and Overfitting:**

- **Challenge:** Financial markets continuously evolve and exhibit changes over time. Models trained in historical market patterns can struggle to adapt to unexpected market changes or outbreaks of major events (Liu et al., 2024). Chen et al., 2025).
- **Limitation:** Overfitting the data used for fine-tuning increases the chances of achieving impressive results on simulations but disappointing outcomes in real-world markets (Liu & Xia, 2022). Common evaluation methods may not sufficiently reflect on the nuances and challenges present in real financial markets (Chen et al., 2025).

4. **Computational Cost:**
    - **Challenge:** Implementation of advanced agentic multimodal AIs often demands significant resources. May present a challenge for resources-constrained organizations (Li et al., 2022; Joshi, 2025). This challenge is particularly relevant in rapidly evolving HFT applications.
5. **Integration Complexity:**
    - **Challenge:** Coordinating the interoperability and performance of numerous AI components such as LLMs, predictive models, RL agents and data feeds, presents a difficult challenge in software engineering (Zhou & Mehra, 2025).

Critical Ethical Considerations:

1. **Systemic Risk and Market Stability:**
    - **Consideration:** If a great number of market players employ the same algorithmic approaches, a high level of coherence could create situations such as flash crashes, major fluctuations and introduction of novel systemic risks. (Azzutti, 2024; Agentic artificial intelligence can transform finance by both enabling and introducing new risks (ResearchGate, 2025). The ability of autonomous agents to exhibit unexpected behaviors when working together raises safety concerns (Guarino et al., 2022).
2. **Accountability and Responsibility:**
    - **Consideration:** Determining who holds responsibility for problems caused by an autonomous AI system in finance can be challenging. (Azzutti, 2024; ProcessMaker, 2025, "Ethical Considerations of Agentic AI”) There are important parameters governing how ethically intelligent agents are developed. Responsibility may lie with the developer, the user or the organization involved.
3. **Fairness and Non-Discrimination:**
    - **Consideration:** The use of biased information or algorithms can result in decisions that are unfair to some individuals or groups such as in loan approval or investment advice, without proper monitoring and corrective measures (Zhao & Welsch, 2024; Lucinity, 2025).

4. **Autonomy, Control, and Human Oversight:**
    - **Consideration:** Finding adequate ways to involve AI and human agents within a responsible framework will be key. Relying too heavily on autonomy without proper controls in place makes errors or misalignment of goals much more likely.
5. **Data Privacy and Security:**
    - **Consideration:** The handling of huge volumes of valuable financial and personal information requires sophisticated protection methods to prevent unauthorized access and misuse.
6. **Potential for Misuse:**
    - **Consideration:** AI systems with such technical capabilities might be misused in ways that can be challenging to identify.

Governance Structures: Effective governance is paramount. This includes:

- Financial organizations should establish comprehensive internal AI governance systems (Azzutti, 2024).
- Having defined ethical principles and also introducing specialized committees to oversee the use of AI.
- Ensuring regular testing, validation and continuous monitoring are carried out (Chen et al., 2025).
- Revising existing financial regulations to specifically address the novel challenges raised by sophisticated AI agents in the financial sector, as compliance with the EU AI Act demonstrates (Azzutti, 2024).
- Collaboration among industry, academia and regulators is essential to establish the most effective industry practices (Lewington et al., 2024).

A comprehensive and integrated solution is needed to tackle these complex issues. A solution that combines advanced technology, ethical guidelines and responsive governance. Advancing Agentic Multimodal AI in the financial sector promises numerous advantages yet achieving them ethically requires diligent efforts to address the underlying risks associated with its applications.

CHAPTER VI: SUMMARY, IMPLICATIONS, AND RECOMMENDATIONS

6.1 Summary

In this work, comprehensive investigation was performed into the transformative potential on how integrating Agentic and Multimodal AI helps in the field of quantitative finance focusing mainly on how these algorithms are used in wealth and portfolio management.

The research started with an analysis of the current situation, as older quantitative models and systems of algorithmic trading encounter numerous new challenges, as financial markets become more complex, dynamic and rich in data, people must adapt. Such obstacles are caused by unstable market environments, various data that is not straightforward to handle and the need to act promptly to new changes in market regimes.

The research identified a framework for an advanced AI that mixes the abilities of Agentic AI - characterized by its autonomy, planning, tool use, and reflection capabilities and Multimodal AI - which uses a combination of numbers, texts and visual data from various resources.

According to the architecture explained in Chapter III, Section 3.4, it is the Large Language Model (LLM) that manages the entire agentic core processes by selecting and guiding the use of specific tools. The tools include quantitative models, including those offered in Qlib, as well as agents from FinRobot and FinRL that use reinforcement learning and multimodal AI modules for analyzing many types of data.

Qlib provides an environment covering all the steps of the quantitative pipeline and emphasizing data manipulation and feature generation. It incorporates applications of known noise-filtering methods e.g. wavelet transforms or Kalman filters, which may be applied to raw time-series data to reduce the random noise. More to the point, it gives availability to ready-made, curated sets of features, including Alpha158 and Alpha360. These consist of hundreds of "alpha factors" which are mathematical transformations of price, volume, and fundamental, and which in academic and industry research have been demonstrated to provide some historical predictive value. These libraries allow a researcher to begin his analysis with a pre-filtered set of signals instead of unprocessed, noisy data giving the inputs to the model a much broader scope.

Libraries such as Qlib are a key to the democratization of the composite use of the high-tech data processing and feature engineering that previously had servers of exclusive march on the elite quantitative hedge funds. They shield the learning curve of composing these methods and make it open source and compacted into an easy-to-use library to ease the entry toward sturdy quantitative research. This enables smaller companies and academic research groups to construct and test strategies on a platform of noise-reduced, well-engineered features to get more reliable and reproducible performance. This interest in quality and improvements of the data and signal in the pre-processing step is an essential addition to the pattern recognition approaches of the downstream AI models.

By using modules with various AI ideas, the goal is to improve the flexibility and intelligence of making financial decisions. The study investigated if this framework applied to several types of algorithmic trading, including investing according to factors. Strategies such as dynamic asset allocation, event-driven trading and HFT were mainly influenced by Chapter IV, Section 4.2.

Current research indicates that this approach could improve the system's performance by providing extra protection from risks and better adjust to market variations and research leading to the discovery of new alpha sources, mainly since these focus on syndicating various bits of data and adjusting actions based on them. For example, Agentic AI can help companies identify alpha factors which RL agents can do by following optimal strategies for adjusting assets.

An essential element in the proposed framework is using Reinforcement Learning with Human Feedback (RLHF). This mechanism addresses the crucial need to align the AI agent's decisions with complex, potentially subjective human preferences—including informed risk tolerance, ethics in investing (such as ESG) and unique long-term aims—in addition to main target achievements like making the highest profits (Chapter IV, Section 4.3).

The approach RLHF uses allows AI to act in ways that reflect human values and the qualitative needs of handling wealth and portfolios. Still, the study pointed out that developing and using advanced AI has significant problems, limitations and ethical hazards (Chapter IV, Section 4.4).

These include issues of data quality, availability, and bias; the "black box" nature of complex models, which impacts interpretability and trust; ensuring model robustness and generalization in non-stationary markets, risks linked to similar actions by AI and deep concerns about fairness, autonomy, accountability and governance. Some researchers also mentioned the expensive computer costs needed for using these models.

This DBA research compiles the best practices for creating and applying Agentic and Multimodal AI to achieve new results in quantitative finance. While there is a big chance to develop more useful and aligned financial assets, to reach that goal, we should overcome serious technical problems and being ready to account for and tackle the challenges related to ethics and governance.

6.2 Implications

Agentic AI and Multimodal AI being applied to quantitative finance, as this thesis shows, has significant implications for several parties involved in the financial sector including financial institutions, investors, regulators and other important groups.

- **For Financial Institutions (Investment Banks, Hedge Funds, Asset Managers):**
    - **Enhanced Capabilities and Competitive Advantage:** Institutions able to successfully use advanced AI are likely to enjoy an important advantage over competitors. It could lead to better investment results thanks to flexible and innovative ways of novel alpha discovery (Sun et al., 2023; Tang et al., 2025), making it possible for wealth managers to automate complex procedures in research, trading and portfolio management. (WalkingTree Technologies, 2025), and the ability to offer hyper-personalized wealth management services on a scale (Han et al. 2024). Moody's analysis indicates that early adopters of agentic AI have seen users consume 60% more research while cutting task completion times by 30% ("Agentic AI in financial service - Moody's").
    - **Transformation of Roles and Skill Requirements:** The rise of agentic "financial copilots" will likely transform the roles of human professionals (WalkingTree Technologies, 2025). These professionals may stop looking through data themselves and analyzing and rather supervising AI agents, begin to look after, control, and plan strategies. They concentrate on stronger management of important clients and crucial decisions that depend on people's insight and sense of right and wrong. Therefore, people in the workforce must possess new abilities at the intersection of finance, data science, AI ethics, and AI system management.
    - **Increased Investment in Technology and Talent:** Adopting and managing modern Agentic Multimodal AI technology will continue to demand significant updates to technology with substantial ongoing investment in technology infrastructure (including HPC), collection and use of data and the hiring of AI experts. This situation may lead to bigger institutions moving further ahead of smaller groups.
    - **Need for Robust AI Governance:** The financial sector will have to establish and follow in-depth AI governance guidelines to handle various risks. That

means, when using these technologies, companies must consider model risk, data bias, how they respond to emergencies and ethical issues (Azzutti, 2024).

- **For Investors (Institutional and Retail):**
  - **Potential for Improved Investment Outcomes:** Investors may gain through having access to sophisticated strategies that are adaptable and could be adjusted to their personal goals.
  - **Greater Access to Sophisticated Advice:** Smart machines would make it easier for people to get sophisticated financial management advice, tips that could be found only with expensive advisors (Han et al., 2024).
  - **New Risks and Challenges:** Investors will also face new risks, including the potential for losses due to AI model failures or misalignments, the opacity of "black box" algorithms making it difficult to learn about deciding to invest and about the privacy issues linked to revealing their finances to AI (Ng et al., 2020). It will be important for people to learn about the strengths and weaknesses of AI in finance.

- **For Regulators and Policymakers:**
  - **Evolving Regulatory Landscape:** Financial regulations may not be able to handle the individual problems that arise from advanced and highly independent AI (Azzutti, 2024). Government agencies need to review current rules or design new ones to manage the risks from biased algorithms, hard-to-explain models, propagation of risks through systems and autonomous agents' actions.
  - **Focus on Systemic Risk:** Because of algorithmic herding, a new kind of risk due to correlated behavior among AI trading agents (algorithmic herding) to the entire market could emerge and extra regulation may be required and perhaps new tools for managing macroprudential issues (Azzutti, 2024).
  - **Guiding Responsible Innovation:** Policymakers must direct innovation in financial AI to preserve financial stability, investors and maintain ethical uses of AI (Lewington et al., 2024). It may require creating and promoting standards for creating, testing and checking AI software in financial organizations.

- **For the Financial Ecosystem and Society:**
  - **Market Efficiency and Dynamics:** AI has the potential to change how markets react. It is argued by some that AI may improve the reliability of

market prices and processes (Oxford University Press, 2024, as cited in). Other people are concerned about higher volatility or newer ways companies may manipulate the market.

- **Ethical Challenges:** Introducing Agentic AI raises serious questions about fairness (e.g., ensuring AI-driven credit or investment decisions are not discriminatory), and job security for people working in the field with firms having most of the power due to their skills in these technologies.
- **Trust in Financial Systems:** Public confidence in the AI-led financial system will be much higher if these AI systems are trusted, fair and help humans (Ng et al., 2020).

Overall, Agentic and Multimodal AI in quantitative finance signifies a genuine paradigm shift that will have wide effects. It promises significant benefits in terms of efficiency, personalization, and performance, but also introduces complex challenges and risks that require careful navigation and proactive governance by all stakeholders.

6.3 Recommendations for Future Research

Although the usage of Agentic AI and Multimodal AI in quantitative finance is speedily growing, further exploration in these areas is still necessary. To achieve this, the organization needs to fully learn about, develop and responsibly use these technologies. From the discussions and findings stated in this thesis, these recommendations are made for future work:

1. **How Economics Deals with Changes and Variety in Businesses:**
    - **Non-Stationarity Adaptation:** Use Agentic and RL-based AI models and closely examine and improve them. It can effectively find out and adjust to changes in the market and fluctuations in data patterns, unlike today's approaches (Liu et al., 2024 - "Dynamic datasets and market environments for financial reinforcement learning"). Studies should emphasize constant learning and effective online adjustments.
        - That solution amounts to a paradigm shift in machine learning deployment with respect to the conventional approach of train first, then deploy. Rather than developing a fixed pattern in periods of retraining, this aims to have systems that it perpetually trains and updates.
        - RL is naturally well-suited to this task, because RL agents can develop optimal policies during their constant interaction with their environment a live or simulated financial market.

- With a change in the market dynamics the rewards and the penalties that the agent gets modified and the feedback also alters and thus the agent must go back to its policy and go with internal improvement to continue being effective.
- Strategic-level agentic AI architecture, frequently conducted through Large Language Models (LLMs), have the capacity to control this on-going learning cycle. An agentic system should be able to track its performance and other relevant factors of the market to notify a possible shift in regime.
- Once detected, the agent may automatically perform an online retraining on new data, it may re-optimize the settings of its current strategy or even aim to change the model it was originally trained, to one which is better adapted to the new setting. This establishes a feedback loop of action, reaction and adjustment which is requisite to operating in non-stationary markets.
- Armed with Reinforcement Learning with Human Feedback (RLHF), this process has added supervision by the human experts who have the authority to direct the learning of this agent towards situating the continuous learning towards the overall desired strategic goals and risk preferences, as opposed to blindly maximizing a possibly mis specified reward signal.
- The paradigm shift in the conduct of model risk management and validation is required in the process of adopting continuous learning models. This is because the traditional idea of testing a single and fixed model with a fixed test dataset is no longer enough.
- In the case where the intention of a model is to develop following a logic that is prescribed in the nature of the model, then the validation should no longer focus on testing the model in a state, but on testing the learning process.
- The risk managers no longer just need to ask themselves critical questions as to whether this model is accurate. but "Is this model of self-change process safe, strong, and goal-resonant?" This includes an extreme tightness of the RL reward function, constraints that might be at the speed of adaptation and its magnitude as well as retraining triggers and the safety overrides that stop the agent at the danger of evolving into a harmful or distasteful mode. This is a more complicated but, by a long way, more relevant way to handle risk in an age of adaptive financial AI.

- **Stress Testing and "Black Swan" Events:** Design comprehensive stress-testing methodologies and simulation environments that can more realistically model extreme market conditions and "black swan" events to evaluate the resilience of AI agents (Chen et al., 2025). Another part of this is exploring attacks on financial AI systems and possible solutions to prevent or resist them.
    - Standard backtesting tests the performance of a model using average historical observations, however, does not necessarily capture its performance in unusual extreme events.
    - A good strategy under ordinary market conditions can also be disastrous in the event of a black swan such as a flash crash or a geopolitical flash crisis.
    - Rigorous stress testing and simulation in addition to standard backtesting must be undertaken in order to curb this risk.
    - This implies the production of synthetic data or simulation platforms that can simulate the high-stress rare cases. Such simulations can challenge a model against:
        - Massive crashes in terms of price (market crashes).
        - Rocky volatility values.
        - Liquidity crisis where you are hard pressed to carry out trades.
        - Failure in the historical association between investments.
        - Adversarial attacks to purposefully tamper with the input to the model.
    - Observing the behavior of an AI agent in such simulated conditions of duress, one will be able to detect the unnoticed vulnerabilities and failure modes that otherwise could not be detected through normal back testing.
    - Relatively, an agent which uses a portfolio optimization technique may be discovered as developing highly concentrated, risky positions in a simulated liquidity crunch.
    - Finding such vulnerabilities in a contained setting makes it possible to patch them out, e.g. by introducing explicit risk restrictions or safety overrides to the logic of the agent, before they incur losses in the real world.
    - It is an important step along the way to robust financial AI systems capable of being trusted with capital in the wild world of finance, because it is fit to survive the worst.

Conduct stress testing and adversarial attack simulations to identify and patch vulnerabilities in financial AI systems

- A proactive defensive approach will consist of systematically red teaming or stress testing the AI system attacking known adversarial attacks. It is not the process of waiting until a real-life attack happens but attempting to create such attacks in a safe place and discovering and patching vulnerabilities before anything goes wrong.
- This refers to the practice of creating adversarial examples, by exploiting known algorithms (e.g. Fast Gradient Sign Method (FGSM), Projected Gradient Descent (PGD)) to create inputs that a human would consider to be identical to valid data, but in some way engineered to feed wrong answers to the model.
- An example of a simulation may include whether adding minuscule imperceptible perturbations to any given data transaction may lead to a breach of the fraud detection model, or whether altering a single word in a news article will lead to reversing the sentiment analysis model results; a negative opinion to a positive one.
- This is an adversarial testing process, which is one of the most crucial pillars of Continuous Threat Exposure Management (CTEM), where the developers can find out the weaknesses their model has but that cannot be seen immediately.
- It is with this information that defenses are implemented, effectively, hardening the system against attack with findings made during these simulations. This cycle of scanning, detection and remediation of vulnerabilities is the basic building block in creating secure financial AI.

2. **Interpretability, Explainability and Trust (XAI in Financial Agents):**
   - **Causal XAI:** Now, XAI should develop techniques that reveal the real causes behind a prediction, instead of just highlighting correlations. When doing this, complex Agentic Multimodal systems must be considered in their decision-making processes (Koa et al., 2024). This helps a lot when debugging, checking for errors and gaining confidence.
   - **Advance Explainable AI (XAI) techniques, focusing on Causal XAI**
     - SHAP (Shapley Additive Explanations) or LIME (Local Interpretable Model-Agnostic Explanations) are examples of classic XAI which have proven useful in offering feature attributions.

- They will bring out what input attracted a model as an essential aspect in making a specific prediction. Nevertheless, these approaches revolve essentially upon correlations.
- They demonstrate what properties the employed model consisted of, yet not why the relationship is sound. This is a major drawback in finance where spurious correlations are the order of the day.
- The development of Causal XAI is a very important step in addressing this shortcoming. The main idea of Causal AI is to discover the true cause and effect relation behind the data thus separating a true cause and misleading correlation.
- An example is on how XAI may indicate a relationship between the sale of an ice cream and a stock market rally as an example of a conventional XAI model.
- A Causal XAI model would be created such that it would discover the unknown confounding factor, being the weather during summertime, that causes both, and as such, not detect the sales of ice cream as a causal factor of the market.
- Causal AI can also offer much more solid, trustworthy and reliable explanations because it will construct a model of a causal graph underlying market dynamics.
- A causal explanation (e.g., the model anticipated a decline due to the existence of a causal relationship between an increase in interest rates and decline in corporate earnings) is far more useful to a risk manager or regulator as compared to a non-causal explanation based on correlation.
- It is this transformation of what to why that creates real trust towards AIs driving financial-based decisions.

- **Explanation of Agentic Reasoning:** How can LLMs help agents come up with responses? and easy-to-understand accounts of how they plan, select tools and think, mostly when dealing with information from multimodal sources.

Agentic reasoning needs to have the internal process of the system translated. This consists in finding ways of logging, interpreting, and visualizing the decision path of the agent:

• **Goal Decomposition:** How did the agent decompose a high-level goal (e.g. optimize the portfolio) into a sequence of concrete sub-goals?

• **Planning:** How did the agent propose to tur out these sub-goals in turn?
• **Tool Selection and Use:** At what step did the agent decide to call a news sentiment API address? What came as an outcome of that tool and what impact did it create on the way forward on the plan?
• **Synthesis:** How did the agent combine the result of different tools and its own internal knowledge to have the final decision?

Such schemes as FinRobot, which use a "financial Chain-of-Thought (CoT)" prompting scheme, are a preliminary step in this direction since they have the LLM state its lines of reasoning processes clearly. It is important at the process-level explanation of auditing and debugging autonomous systems.

It enables the human supervisor to question the logic deployed by the agent and watch out the possible weaknesses of the plan (e.g., "The plan developed by the agent did not take into account the risk of a change in regulations"), which is critical towards establishing a trusting relationship and safe operation.

Develop methods to explain Agentic AI reasoning and process

- **User-Centric Explainability:** Evaluate the usefulness of different kinds of explanations in helping users make decisions about finance and choose appropriate services. This is the case for the typical small-time retail buyer versus large traders dealing in the financial markets.
  A good explanation is not universal at all; it strongly depends on the audience. What a quantitative developer, a portfolio manager, a compliance officer and a retail investor need to know are radically different.
  It follows that any good XAI must be that which involves a user-centered design mindset, in which the form, content, and complexity of the explanation is designed to suit the particular context of the user, their aims, and their level of expertise.
  Such a direction turns XAI not into a solely technical validation feature but something crucial to the user facing properties of a product as a whole. The

structure of clarification gets to be a component of the user experience in general. For instance:

• **Retail Investor:** He/She requires a simple, common-sense explanation. They are often most effective when in the form of counterfactuals (Your loan was not been allowed because you have a debt to income ratio of 50; had been lower than 40 you would have been able to receive it) or using examples (That bank cannot offer you a loan, since your debt to income ratio is 50; at 40 percent, they would have accepted).

• **To a Portfolio Manager:** actionable insights should be given in the explanation. They may have to view the factors taking the biggest toll on the risk in their portfolio or the assets being signalled a possible decline.

• **To a Compliance Officer**: The priority becomes an auditable, traceable record. They need a careful record of all the decision-making activities of the agent so as to make sure that his action is within the realms of all the requirements and regulations as well as internal policies.

• **To a Developer:** Debugging requires a very technical explanation, which displays attribution of details features, instead of displaying the model weights and activation paths.

With the help of user-centered design principles (claiming user needs, delivering custom design solutions, and then testing them with actual users), organizations are capable of developing explanations that are not only technically correct, yet essentially useful, bred, and credible in the eyes of every particular group of stakeholders.

3. **Advancements in RLHF and Value Alignment:**
   - **Scalable and Efficient RLHF:** Build more advanced, less expensive and reliable systems for gathering quality feedback from humans during RLHF in finance. Investigate active learning approaches to select the most useful questions for feedback (Xiong et al., 2025).
   - **Handling Subjectivity and Disagreement:** How to Include and Mange Disagreements and New Expectations: The various methods to make sure that research techniques for RLHF can adapt to new and uncertain human preferences and gather comments from numerous experts.
   - **Formal Verification of Alignment:** Check and verify that the actions of AI agents are ethical and in accordance with the planned strict rules going further than just observing data using RLHF.

- Formal methods instead of relying on empirical testing (which can never test all possible conditions) can demonstrate, for example, that a trading agent will never go above a particular level of leverage, or that a portfolio manager will never invest more than a given amount in a given asset type, no matter what market signals it is presented with.
- Formal verification of complex, probabilistic AI models is an exciting, but difficult field of research, as well as one with a huge potential payoff: real-life scenarios where the stakes are high requires a high level of confidence.
- In cases of crucial safety requirements and regulatory guidelines that should not be jeopardized at any cost, formal verification may present the form of assurance which cannot be attained by testing solely. It plays an essential role in creating systems which are not only aligned on average but can be proved to be safe and reliable at performance boundaries.

4. **Multi-modal Fusion and Representation Learning:**
   - **Dynamic and Contextual Fusion:** Upgrade fusion approaches allow them to automatically change the weighting subject to the context. And, combining their data forms should depend on what happens in the market, how good the data is or the relevance of the task (Zhang et al., 2024). ("A Multimodal Foundation Agent for Financial Trading: Tool-Augmented, Diversified, and Generalist")).
   - **Develop dynamic and contextual data fusion approaches**
     - Some basic ideas of multimodal fusion could be to concatenate cross-modal feature or a fixed weighting process.
     - Nevertheless, the relative weights of various data modalities do not change; they are condensed on market circumstances very much.
     - An example would be when the day of trading is not robust, in which case numerical price patterns are of prime importance.
     - However, Texas news data and sentiment analysis are much more important in the case of a geopolitical crisis or when a central bank announces it in the news.
     - This problem is curbed by an innovative technique of dynamic and contextual fusion in which the model is made to learn how to weigh the various sources of data in real-time.
     - Modern architectures that employ cross-modal attention operations can adaptively assign greater "attention" to the data modality most

applicable at a particular point in time as per the current state of the market and the predictive task, in question.

  - The model does not only learn what is said in each data source, but also how to listen to it, when you should hear it the closest.
  - The future research directions on this proposed adaptive fusion strategy are to develop an intelligent system that will pass the test of time and be able to navigate a dynamic information environment that exists in financial markets.
  - This strategy apes with the instinct of a human analyst, who puts her or his attention to where charts, news terminals, and research reports make the most important market stories, whatever they are currently.

- **Cross-Modal Causal Inference:** Focus on discovering causation among evidence gathered in different types of data. (e.g., how a news event *causes* a specific price movement and sentiment shift).
  - One of the main drawbacks of a range of existing fusion approaches is that they are grounded in learning relations between the varying forms of data.
  - Nevertheless, the correlation does not mean causation. The further and more solid direction is to study cross-modal causal inference.
  - The idea is that this is not enough to merely identify that there was a news event and a price movement that happened at the same time we now want to determine whether this news event was the actual cause of the movement in price.
  - The creation of models capable of deducing such casualties is one of the frontiers of AI research.
  - It may amount to the application of methods of causal discovery, or structural causal models, or experiment-based testing of causal hypotheses.
  - As an example, a model may be trained to understand the causal path between the earnings announcement (text) by a specific company and the cause of an analyst-sentiment shift (sentiment) and subsequent market price reaction (numerical).
  - Discovery of these causal avenues would result in a very deep and solid integration of data. The model that captures the cause-effect forces behind market trends can be much more predictive and less apt at being misled by extraneous correlations that could fail in the long term.

    - Although not an easy direction, cross-modal causal inference is one of the directions of study that is likely to open the gates to a more in-depth and reliable synthesis of financial information.
5. **Systemic Risk and How Many Agents Are Present:**
    - **Modeling Emergent Behavior:** Program strong multi-agent simulators to explore possible emergent behaviors and risks in the financial markets caused by many independent AI trading agents working together (Guarino et al., 2022).
        - Under sandboxed environments, not only will the researchers and risk managing personnel perform experiments to learn about the emergent phenomena, but they will also be able to simulate the real-world environment. They may get to know the situations like:
            - How does it go when multiple agents are used with such RL-based ideas in large quantities? Are they taught to correlate their actions, and do they hear as a result?
            - What is the means of response of autonomous agents to a sudden shock in the market? This is to say, do their joint actions crush the shock or reinforce it leading to a flash crash?
            - Is it possible to train advanced actors to take advantage or coerce the less advanced?
        - Systemic risk can be evaluated in multi-agent simulators that model such interactions and turn them into a laboratory. They can be used to detect potentially catastrophic emergent behavior and to allow the agents strategies or market regulations to be adjusted in advance to limit the risks before they can occur.
    - **AI-Driven Market Stability Mechanisms:** Research the possibility of using AI agents both to cause and to counter market risks, by providing liquidity in needed markets or by anticipating and managing small crash risks when they appear.

        The stability mechanisms based on AI might be structured to carry out multiple roles:

        - **Automated Liquidity Pro Measure:** Stabilizer agents could also be set up to identify a rapid departure of liquidity in a market (a characteristic of a flash crash) and offer bids and offers to stabilize the market and reinstate order.
        - **Counter-Cyclical Trading:** such agents might be set up to detect the initial indicators of a herding-driven bubble or crash and buy or sell accordingly and lean against the wind to reduce its effects.

- **Risk Monitoring:** An AI agent may be employed to observe the overall behavior of other agents in the market in order to give an early indication of increasing correlations, or other measures of systemic risk to human regulators or supervisors.

6. **Ethical AI and Fairness in Algorithms for the Financial Sector.**
   - o **Bias Detection and Mitigation in Multimodal Data:** Find and minimize biases in financial data drawn from a range of sources (text news, social media platforms, etc.) and making certain Agentic Multimodal AI systems do not lead to these biases (Chen et al., 2025; Zhao & Welsch, 2024).

     To deal with this, we need to actively engage the methods of bias identification and correction at all stages of AI lifecycle. This involves:

     - **Auditing Datasets:** Probing training dataset in a systematic way to ensure any biases are spotted prior to training.
     - **Fairness-Aware Algorithms:** With the help of specialized machine learning algorithms (primarily aimed at model performance optimization with respect to the accuracy and fairness measures), the performance of a model will be equitable between different demographic groups.
     - **Post-Hoc Bias Correction:** The use of methods to manipulate the production of a trained model to rectify a bias.

     This is one of the areas where future research should concentrate, namely working on these techniques on the specificity of multimodal financial data, so that the shift in the direction of more context-informed AI will not come at the expense of fairness and equity.
   - o **Fairness-Aware AI Agents:** Setting up AI agents for scoring credit or giving personalized investment advice should include special mention of fairness and its objectives.
7. **AI Governance and Regulatory Frameworks:**
   - o **Auditable AI Systems:** Research techniques and standards for creating auditable Agentic AI systems in finance, allowing regulators and internal auditors to effectively scrutinize their behavior and compliance (Azzutti, 2024).
   - o **Adaptive Regulation:** Investigate methods for flexible regulations that respond to financial AI’s rapid advancements possibly incorporating "regulatory sandboxes" for testing new AI innovations.

8. **Human-AI Collaboration and Augmentation:**
    - **Optimal Human-in-the-Loop Design:** Examine the strategies that work best when humans and AI work together in the financial industry. where machines are used to further improve what humans can achieve than replace them. It also covers the development of human-friendly interfaces for taking over control if needed. In the year 2025, WalkingTree Technologies considered AI in its solutions.
    - **Cognitive Load and Skill Development:** Study how difficult it is for a supervisor to handle the AI system's complex tasks and actions. Figure out what a financial professional should possess to be ready for the future with AI.
9. **Develop robust model architecture with defensive layers:**
    - This solution emphasizes the aspect of incorporating the resilience directly in model architecture and training. Adversarial training is one of the most efficient defenses.
    - This method presupposes an increase of the training set, so there are not only legitimate examples but also a huge amount of adversarial examples that are generated during training process itself.
    - This causes the model to be robust on a decision boundary without depending on minor adversarial adjustments by asking it to learn to identify the actual malicious inputs accurately. This almost inoculates the model against the known kinds of attacks which make its resiliency innate.
    - But this comes with its trade off. There is a risk of slightly decreasing the accuracy of a model on unadversarial, clean data under adversarial training that is sometimes referred to as over-hardening. This is since the process causes the model to learn more simplified, smoother decision boundaries, which are more robust but may not necessarily represent all the details of the original data, which are fine grained.
    - One of the most important tasks developers have in mind is to achieve the desired balance between regular precision and adversarial robustness. Another concept of architectural defense is the deployment of input layers of validation and sanitisation that may be able to detect and filter possible anomalous and evil input prior to the inaccessibility of the central model.
10. **Regularly validate and update models to adapt to emerging attack vectors**

AI security cannot be a once and done thing, it should be a set it and forget it job. It needs constant vigilance and adaptation. Financial institutions are urged by this solution to put in place a security validation process that is continuous. This consists in:

- **Continuous Monitoring:** This is active monitoring of academic and security communities to keep up to date on the new emerging attack vectors.
- **Regular Validation:** Occasional re-testing of deployed models against these latest attack techniques to find out any new vulnerability which may have arisen.
- **Model Updating:** In case of identification of vulnerability, models should be updated in time. This could include retraining them using new varieties of adversarial examples or using new defensive layers to ward against the danger.

Exploring these research topics will help get the most out of Agentic and Multimodal AI in quantitative finance. At the same time, work to minimize the risks and use these technologies ethically, fairly and safely. They play a positive role in the financial system and in society.

Getting industry, universities and government bodies to team up is necessary to achieve significant improvements regarding these several problems.

6.4 Overview of Proposed Agentic and Multimodal AI Models

The proposed framework is a combination of Agentic AI focused on autonomy, planning, tools, and reflection with Multimodal AI that deals with various data types such as numerical, textual, and visual data. Such combination uses data manipulation tools such as Qlib and feature engineering as noise-filtering techniques, including wavelet transform and pre-curated alpha factors (examples: Alpha158 and Alpha360) to improve signal quality. Traditionally, proprietary rules-based trading systems do not have such adaptive, multi-source integration, based instead on pre-defined rules, statistical models and historical patterns.

General Performance Differences

Multimodal and Agentic AI systems would deliver superior performance over conventional systems because they can make dynamic decisions and respond to complex market situations via their dynamic nature. The most important dissimilarities are:

Autonomy and Adaptability: Agentic AI takes independent decisions and can make runtime strategy changes using reinforcement learning (RL) agents compared to the strict rule-based logic of the traditional systems which would necessitate human interaction to change the logic. Multimodal AI builds on this by combining different sources of data and provides stronger predictions than the single-modal (e.g., numeric) information used in traditional models. Alpha discovery and risk management: The proposed models use RL and multimodal analysis to find new alpha sources that enhance risk-adjusted returns, but traditional systems cannot keep up with the non-stationarity of markets and tend to ignore unstructured data in news, or sentiment.

Scalability and Efficiency: Democratization of advanced functionality via open-source libraries such as Qlib makes these AI models less steep in their learning curve and allows smaller organizations to quickly reach reliable performance as compared to traditional systems which require plenty of manual tuning and which are not as scalable. Among the representatives of early adopters, the quantitative advances are exhibited in an ability to consume research up to 60 percent faster and accomplish the tasks with 30 percent less time to display the competitive advantage over the traditional methods.

Performance during wide Fluctuations in the market:

During extreme circumstances, standard algorithmic trading systems tend to fail because of their inability to manage quick regime switches, resulting in increased risks such as herding phenomenon or flash crashes. It navigates mechanism scores well in the proposed Agentic and Multimodal AI models as being pro desperately reconfigurable and resilient. The comparison is shown in detail as shown below:

Table 6.1: Comparison of Traditional Algorithmic Trading Systems and Proposed Agentic and Multimodal AI Models in High-Volatility Market Conditions

Aspect	Traditional Algorithmic Trading Systems	Proposed Agentic and Multimodal AI Models
Adaptation to Volatility	Reactive, founded on set regulations, resulting in the slow responses and the ensuing possible losses in the black swan events	Proactive; trains on RL continuously to learn and identify regimes can subsequently automatically re-train on new data to minimize the effects of volatility.
Data Integration	Could only use structured historical data and does not work with real time multimodal inputs such as news opinion.	Agilely combines numerical, textual and visual data, to make better crisis (e.g., geopolitical event) predictions

Risk Mitigation	Tendentious to a correlated failure when subjected to stress conditions, and simple backtesting poses inadequately react to high volatility	Incorporates stress testing and adversarial simulations for robustness, reducing systemic risks like algorithmic herding
Decision-Making Speed	Human-dependent adjustments will occur more slowly hence expose them to market fluctuations that are rapid.	The stability is guaranteed by autonomous, real-time actions (e.g. liquidity provision or counter-cyclical trading).
Alignment with Preferences	Does not have frameworks of subjective variables such as risk tolerance or morals.	Uses RLHF to incorporate decisions into our human values so that ethical answers can be given to volatile times

These models have better resilience in the high volatility environment through simulating extreme conditions and adjusting policies through feedback loops, which may save the disaster that conventional systems cause more damage. An example that illustrates this is that Agentic AI can identify market shocks early and take protective actions whereas Multimodal Fusion contextualizes information that can be easily missed by some of the traditional models.

Trading Strategies Implications

The proposed models are more flexible to support advanced approaches such as event driven trading and High frequency trading (HFT) that would result in better results in a dynamic environment. With the set pattern in the Traditional systems, the method will lack performance against volatility since they are non-changing systems. In general, such integration is expected to result in market efficiency, but the challenges of interpretability of models and ethical management should be solved.

6.5 Conclusion

This DBA research covers the progress in quantitative finance and pinpoints the notable difference that Agentic AI and Multimodal AI integration may make.

The transition from traditional statistics to AI agents that can plan, search for solutions, learn from various types of input and exchange information with humans has an important impact in how to examine and deal with financial markets.

The framework I have devised, using various sources, sees an AI system working like a partner in algorithmic trading and controlling personal wealth and portfolios and not as only a predictive model. LLMs are used to coordinate functions, access data from a range of channels and adjust learning as part of reinforcement learning.

It is important to focus on RLHF for value alignment—systems with these methods to promise better outcomes, adjust to different financial markets and can address diverse personalized financial goals. The advantages can range from better risk-adjusted results, to discovering original investment ideas, better running the operation and offering more investment strategies to a wider audience.

On the other hand, these powerful changes can only happen with major challenges and significant moral questions. Problems related to how data is organized, addressing real robustness and generalization, managing systemic risk and ensuring AI governance is trustworthy, fair and transparent remain very challenging.

They can only be deployed successfully if we incorporate new techniques and, most importantly, create strict ethical rules and good risk practices and laws that are flexible and update with changes in technology.

The journey ahead requires a concerted, collaborative effort to achieve our goals. Those working in AI should set new standards and be attentive to matters related to safety and transparency. The financial organizations should focus on being ethical and reliable as they seek to succeed in today’s market. Proper regulations should be put in place to support new developments that benefit the market and investors.

Essentially, merging Agentic AI and Multimodal AI with quantitative finance transforms the way people work together in finance and machine capability in one of society's most critical sectors. If the financial industry is aware of possible dangers and remains dedicated to following the right rules, they can use these tools effectively to work toward improving the stability, efficiency and justice found in our financial system.

This study helps continue the discussion by looking at the current situation and designing an outline of future systems helps by setting out the course of action clearly.

REFERENCES

Ahadzadeh, A.S., Ong, F.S. and Veeraiah, C. (2024) 'The influence of competitiveness trait on attitudes toward memory implants: exploring the mediating role of perfectionism discrepancy', *Current Psychology*, 43(24), pp.1-12.

Albahri, A.S., Duhaim, A.M., Fadhel, M.A., Alnoor, A., Baqer, N.S., Alzubaidi, L., Albahri, O.S., Alamoodi, A.H., Bai, J., Salhi, A. and Santamaría, J. (2023) 'A systematic review of trustworthy and explainable artificial intelligence in healthcare: Assessment of quality, bias risk, and data fusion', *Information Fusion*, 96, pp.156-191.

Albert, S. and Hamann, M. (2021) 'Putting wake words to bed: We speak wake words with systematically varied prosody, but CUIs don't listen', in *Proceedings of the 3rd Conference on Conversational User Interfaces*, pp. 1-5.

Alidousti, M., Khakzar Bafruei, M. and Afshar Sedigh, A.H. (no date) '*Pdqn: A Novel Data-Efficient Double Deep Q-Network Framework for Intelligent Financial Portfolio Management*'. Available at: SSRN 5288029.

Anwar, M., Wang, C., De Nijs, F. and Wang, H. (2022) 'Proximal policy optimization based reinforcement learning for joint bidding in energy and frequency regulation markets', in *2022 IEEE Power & Energy Society General Meeting (PESGM)*, Piscataway, NJ: IEEE, pp. 1-5.

Avancha, K.M. (2024) '*The Post-Lockdown Thought Process: Analyzing Shifts in Social Behavior, Personal Finance, and Ethical Values*'.

Azzutti, A. (2024) '*AI Governance in Algorithmic Trading: Some Regulatory Insights from the EU AI Act*'. Available at: SSRN.

Bai, S., Kolter, J.Z. and Koltun, V. (2018) 'An empirical evaluation of generic convolutional and recurrent networks for sequence modeling', *arXiv preprint arXiv:1803.01271*.

Bai, T., Lang, Q., Song, S., Fang, Y. and Liu, X. (2022) 'Feature Fusion Deep Reinforcement Learning Approach for Stock Trading', in *2022 41st Chinese Control Conference (CCC)*, Piscataway, NJ: IEEE, pp. 7240-7245.

Baimuratov, I. (2024) 'Diagrammatic Reasoning for ALC Visualization with Logic Graphs', in *Proceedings of the ACM Web Conference 2024*, pp. 4295-4303.

Baldacci, B. (2021) '*Quantitative finance at the microstructure scale: algorithmic trading and regulation*'. Unpublished doctoral dissertation, Institut Polytechnique de Paris.

Bhoyar, S.R. and Ingole, Y.B. (2025) 'Cosmic evolution of the Kantowski-Sachs universe in the context of a bulk viscous string in teleparallel gravity', *Indian Journal of Physics*, 99(5), pp.1589-1601.

Boyd, S., Johansson, K., Kahn, R., Schiele, P. and Schmelzer, T. (2024) 'Markowitz portfolio construction at seventy', *arXiv preprint arXiv:2401.05080*.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A. and Agarwal, S. (2020) 'Language models are few-shot learners', *Advances in neural information processing systems*, 33, pp.1877-1901.

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y.T., Li, Y., Lundberg, S. and Nori, H. (2023) 'Sparks of artificial general intelligence: Early experiments with gpt-4', *arXiv preprint arXiv:2303.12712*.

Budish, E., Cramton, P. and Shim, J. (2015) 'The high-frequency trading arms race: Frequent batch auctions as a market design response', *The Quarterly Journal of Economics*, 130(4), pp.1547-1621.

Buening, T.K., Gan, J., Mandal, D. and Kwiatkowska, M. (2025) 'Strategyproof reinforcement learning from human feedback', *arXiv preprint arXiv:2503.09561*.

Cai, Y. (2024) 'Overcoming data limitations in credit risk assessment with fingpt-generated synthetic data', in *2024 International Conference on Electronics and Devices, Computational Science (ICEDCS)*, Piscataway, NJ: IEEE, pp. 137-141.

Campbell, J.Y., Lo, A.W., MacKinlay, A.C. and Whitelaw, R.F. (1998) 'The econometrics of financial markets', *Macroeconomic Dynamics*, 2(4), pp.559-562.

Cao, Z., Xu, J., Dong, C., Yu, P. and Bai, T. (2024) 'Matcc: A novel approach for robust stock price prediction incorporating market trends and cross-time correlations', in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 187-196.

Chavan, S., Kumar, P. and Gianelle, T. (2021) 'Intelligent Investment Portfolio Management using Time-Series Analytics and Deep Reinforcement Learning', *SMU Data Science Review*, 5(2), p.7.

Chen, T. and Guestrin, C. (2016) 'Xgboost: A scalable tree boosting system', in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785-794.

Chen, Z., Chen, J., Chen, J. and Sra, M. (2025) 'Position: Standard Benchmarks Fail--LLM Agents Present Overlooked Risks for Financial Applications', *arXiv preprint arXiv:2502.15865*.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H. and Bengio, Y. (2014) 'Learning phrase representations using RNN encoder-decoder for statistical machine translation', *arXiv preprint arXiv:1406.1078*.

Costa, C.D.S.B. and Costa, A.H.R. (2023) 'POE: a general portfolio optimization environment for FinRL', in *Brazilian Workshop on Artificial Intelligence in Finance (BWAIF)*, Porto Alegre: SBC, pp. 132-143.

Costa, C.D.S.B. and Costa, A.H.R. (2024) 'RLPortfolio: Reinforcement Learning for Financial Portfolio Optimization', in *Brazilian Conference on Intelligent Systems*, Cham: Springer Nature Switzerland, pp. 412-426.

Croom, J. (2025) '*Interactivity and Illusions of Ability: How Using Generative AI Affects Investor Judgments*'. Unpublished doctoral dissertation, University of Washington.

Deng, L., Wang, T., Zhao, Y. and Zheng, K. (2024) 'MILLION: A general multi-objective framework with controllable risk for portfolio management', *arXiv preprint arXiv:2412.03038*.

Deng, S., Peng, H., Xu, J., Liu, C., Giurcuaneanu, C.D. and Liu, J. (2025) 'Understanding Financial Reasoning in AI: A Multimodal Benchmark and Error Learning Approach', *arXiv preprint arXiv:2506.06282*.

Devlin, J., Chang, M.W., Lee, K. and Toutanova, K. (2019) 'Bert: Pre-training of deep bidirectional transformers for language understanding', in *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pp. 4171-4186.

Dhar, V. and Sedoc, J. (2025) 'DBOT: Artificial Intelligence for Systematic Long-Term Investing', *arXiv preprint arXiv:2504.05639*.

Ding, H., Li, Y., Wang, J. and Chen, H. (2024) 'Large language model agent in financial trading: A survey', *arXiv preprint arXiv:2408.06361*.

Ding, Q., Shi, H., Guo, J. and Liu, B. (2024) 'Tradexpert: Revolutionizing trading with mixture of expert llms', *arXiv preprint arXiv:2411.00782*.

Ezeilo, C. and Nimo, C. (no date) '*Automated Stock Trading using Trust Region Policy Optimization*'.

Fama, E.F. (1970) 'Efficient capital markets: A review of theory and empirical work', *The journal of Finance*, 25(2), pp.383-417.

Fatemi, S. and Hu, Y. (2024) 'FinVision: A multi-agent framework for stock market prediction', in *Proceedings of the 5th ACM International Conference on AI in Finance*, pp. 582-590.

Fu, H. (2023) 'A graph-based network with attention for stock price prediction', in *2023 International Conference on Communications, Computing, Cybersecurity, and Informatics (CCCI)*, Piscataway, NJ: IEEE, pp. 1-8.

Gao, Y., Chen, H., Wang, X., Wang, Z., Wang, X., Gao, J. and Ding, B. (2024) 'Diffsformer: A diffusion transformer on stock factor augmentation', *arXiv preprint arXiv:2402.06656*.

Gong, Z. (2023) '*Deep learning for trading and hedging in financial markets*'. Unpublished doctoral dissertation, University of Essex.

Greenblatt, R., Denison, C., Wright, B., Roger, F., MacDiarmid, M., Marks, S., Treutlein, J., Belonax, T., Chen, J., Duvenaud, D. and Khan, A. (2024) 'Alignment faking in large language models', *arXiv preprint arXiv:2412.14093*.

Gu, J., Ye, J., Wang, G. and Yin, W. (2024) 'Adaptive and explainable margin trading via large language models on portfolio management', in *Proceedings of the 5th ACM International Conference on AI in Finance*, pp. 248-256.

Guarino, A., Grilli, L., Santoro, D., Messina, F. and Zaccagnino, R. (2022) 'To learn or not to learn? Evaluating autonomous, adaptive, automated traders in cryptocurrencies financial bubbles', *Neural Computing and Applications*, 34(23), pp.20715-20756.

Guo, J., Peng, J., Yuan, H. and Ni, L.M.S. (2023) 'HXPY: A high-performance data processing package for financial time-series data', *Journal of Computer Science and Technology*, 38(1), pp.3-24.

Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N.V., Wiest, O. and Zhang, X. (2024) 'Large language model based multi-agents: A survey of progress and challenges', *arXiv preprint arXiv:2402.01680*.

Han, W., Xu, J., Cheng, Q., Zhong, Y. and Qin, L. (2024) 'Robo-Advisors: Revolutionizing Wealth Management Through the Integration of Big Data and Artificial Intelligence in Algorithmic Trading Strategies', *Journal of Knowledge Learning and Science Technology*, 3(3), pp.33-45.

Hochreiter, S. and Schmidhuber, J. (1997) 'Long short-term memory', *Neural computation*, 9(8), pp.1735-1780.

Hornuf, L., Streich, D.J. and Töllich, N. (2025) '*Making GenAI Smarter: Evidence from a Portfolio Allocation Experiment (No. 11862)*'. CESifo Working Paper.

Huang, J., Wang, J., Li, Q. and Jin, X. (2024) 'Social Media Development and Multi-Modal Input for Stock Market Prediction: A Review', in *2024 International Conference on Computing, Networking and Communications (ICNC)*, Piscataway, NJ: IEEE Computer Society, pp. 198-202.

Jenkinson, T. and Jones, H. (2009) 'IPO pricing and allocation: a survey of the views of institutional investors', *The review of financial studies*, 22(4), pp.1477-1504.

Johnson, B. (2024) '*Reinforcement Learning and Online Learning for Dynamic Trading Strategies*'. Unpublished master's thesis, Northeastern University.

Joshi, S. (2025) 'Advancing innovation in financial stability: A comprehensive review of ai agent frameworks, challenges and applications', *World Journal of Advanced Engineering Technology and Sciences*, 14(2), pp.117-126.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q. and Liu, T.Y. (2017) 'Lightgbm: A highly efficient gradient boosting decision tree', *Advances in neural information processing systems*, 30.

Kokare, K., Pimpale, S., Kapadnis, N., Navghare, P., Raigar, D. and Urane, K.R. (2024) 'Trading Insight AI: An Intelligent Research Tool Revolutionizing Financial Analysis with AI and NLP', in *3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024)*.

Kolesar, J., Ali, S., Antonopoulos, T. and Piskac, R. (2025) 'Coinductive Proofs of Regular Expression Equivalence in Zero Knowledge', *arXiv preprint arXiv:2504.01198*.

Kollintzas, T. and Vassilatos, V. (2023) '*Implications of market and political power interactions for growth and the business cycle II: Politico-economic equilibrium*'.

Kong, Y., Nie, Y., Dong, X., Mulvey, J.M., Poor, H.V., Wen, Q. and Zohren, S. (2024) 'Large Language Models for Financial and Investment Management: Applications and Benchmarks', *Journal of Portfolio Management*, 51(2).

Konstantinidis, T., Iacovides, G., Xu, M., Constantinides, T.G. and Mandic, D. (2024) 'Finllama: Financial sentiment classification for algorithmic trading applications', *arXiv preprint arXiv:2403.12285*.

Kou, Z., Yu, H., Luo, J., Peng, J., Li, X., Liu, C., Dai, J., Chen, L., Han, S. and Guo, Y. (2024) 'Automate strategy finding with llm in quant investment', *arXiv preprint arXiv:2409.06289*.

Krishnan, N. (2025) 'Ai agents: Evolution, architecture, and real-world applications', *arXiv preprint arXiv:2503.12687*.

Lee, N. and Moon, J. (2023) 'Offline reinforcement learning for automated stock trading', *IEEE Access*, 11, pp.112577-112589.

Lettau, M. and Ludvigson, S. (2001) 'Consumption, aggregate wealth, and expected stock returns', *the Journal of Finance*, 56(3), pp.815-849.

Lewington, A., Vittalam, A., Singh, A., Uppuluri, A., Ashok, A., Athmaram, A.M., Milt, A., Smith, B., Weinberger, C., Sarin, C. and Bergmeir, C. (2024) 'Creating a Cooperative AI Policymaking Platform through Open Source Collaboration', *arXiv preprint arXiv:2412.06936*.

Li, C., Shi, Y., Luo, Y. and Tang, N. (2025) 'DeepFund: Will LLM be Professional at Fund Investment? A Live Arena Perspective', *arXiv preprint arXiv:2503.18313*.

Li, Y., Luo, B., Wang, Q., Chen, N., Liu, X. and He, B. (2024) 'A reflective llm-based agent to guide zero-shot cryptocurrency trading', *arXiv preprint arXiv:2407.09546*.

Li, Y., Yu, Y., Li, H., Chen, Z. and Khashanah, K. (2023) 'Tradinggpt: Multi-agent system with layered memory and distinct characters for enhanced financial trading performance', *arXiv preprint arXiv:2309.03736*.

Li, Z., Liu, X.Y., Zheng, J., Wang, Z., Walid, A. and Guo, J. (2021) 'Finrl-podracer: high performance and scalable deep reinforcement learning for quantitative finance', in *Proceedings of the second ACM international conference on AI in finance*, pp. 1-9.

Lim, B., Arık, S.Ö., Loeff, N. and Pfister, T. (2021) 'Temporal fusion transformers for interpretable multi-horizon time series forecasting', *International journal of forecasting*, 37(4), pp.1748-1764.

Liu, J. (2024) 'A survey of financial ai: Architectures, advances and open challenges', *arXiv preprint arXiv:2411.12747*.

Liu, W., Ge, Y. and Gu, Y. (2024) 'Multi-factor stock price prediction based on GAN-TrellisNet', *Knowledge and Information Systems*, 66(7), pp.4177-4198.

Liu, X.Y., Rui, J., Gao, J., Yang, L., Yang, H., Wang, Z., Wang, C.D. and Guo, J. (2021) 'FinRL-meta: A universe of near-real market environments for data-driven deep reinforcement learning in quantitative finance', *arXiv preprint arXiv:2112.06753*.

Liu, X.Y., Wang, G., Yang, H. and Zha, D. (2023) 'Fingpt: Democratizing internet-scale data for financial large language models', *arXiv preprint arXiv:2307.10485*.

Liu, X.Y., Xia, Z., Rui, J., Gao, J., Yang, H., Zhu, M., Wang, C., Wang, Z. and Guo, J. (2022) 'Finrl-meta: Market environments and benchmarks for data-driven financial reinforcement learning', *Advances in Neural Information Processing Systems*, 35, pp.1835-1849.

Liu, X.Y., Xia, Z., Yang, H., Gao, J., Zha, D., Zhu, M., Wang, C.D., Wang, Z. and Guo, J. (2024) 'Dynamic datasets and market environments for financial reinforcement learning', *Machine Learning*, 113(5), pp.2795-2839.

Liu, X.Y., Yang, H., Chen, Q., Zhang, R., Yang, L., Xiao, B. and Wang, C.D. (2020) 'FinRL: A deep reinforcement learning library for automated stock trading in quantitative finance', *arXiv preprint arXiv:2011.09607*.

Liu, X.Y., Yang, H., Gao, J. and Wang, C.D. (2021) 'FinRL: Deep reinforcement learning framework to automate trading in quantitative finance', in *Proceedings of the second ACM international conference on AI in finance*, pp. 1-9.

Lo, A. (2017) *Adaptive markets: Financial evolution at the speed of thought*. Princeton, NJ: Princeton University Press.

Markowitz, H.M. (2008) *Portfolio selection: efficient diversification of investments*. New Haven, CT: Yale University Press.

Miehling, E., Ramamurthy, K.N., Varshney, K.R., Riemer, M., Bouneffouf, D., Richards, J.T., Dhurandhar, A., Daly, E.M., Hind, M., Sattigeri, P. and Wei, D. (2025) 'Agentic ai needs a systems theory', *arXiv preprint arXiv:2503.00237*.

Montazeri, S., Mirzaeinia, A., Jumakhan, H. and Mirzaeinia, A. (2023) 'Cnn-drl for scalable actions in finance', in *2023 International Conference on Computational Science and Computational Intelligence (CSCI)*, Piscataway, NJ: IEEE, pp. 302-306.

Mun, Y. and Kim, N. (2025) 'Leveraging Large Language Models for Sentiment Analysis and Investment Strategy Development in Financial Markets', *Journal of Theoretical and Applied Electronic Commerce Research*, 20(2), p.77.

Murphy, J.J. (1999) *Technical analysis of the financial markets: A comprehensive guide to trading methods and applications*. New York: Penguin.

Nasser, M., Green, C.C. and Kalmoun, E.M. (2025) 'A computational method for multiple steady Hele-Shaw bubbles in planar domains', *arXiv preprint arXiv:2503.12422*.

Ngo, C.N., Friolo, D., Massacci, F., Venturi, D. and Battaiola, E. (2020) 'Vision: What If They All Die? Crypto Requirements For Key People', in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, Piscataway, NJ: IEEE, pp. 178-183.

Obeid, J. and Hoque, E. (2020) 'Chart-to-text: Generating natural language descriptions for charts by adapting the transformer model', *arXiv preprint arXiv:2010.09142*.

Okpala, I., Golgoon, A. and Kannan, A.R. (2025) 'Agentic AI systems applied to tasks in financial services: modeling and model risk management crews', *arXiv preprint arXiv:2502.05439*.

Olujimi, P.A., Owolawi, P.A., Mogase, R.C. and Wyk, E.V. (2025) 'Agentic AI frameworks in SMMEs: A systematic literature review of ecosystemic interconnected agents', *AI*, 6(6), p.123.

Pal, A., Gopi, S. and Lee, K.M. (2023) 'Fintech agents: Technologies and theories', *Electronics*, 12(15), p.3301.

Pástor, Ľ. (2000) 'Portfolio selection and asset pricing models', *The Journal of Finance*, 55(1), pp.179-223.

Pástor, Ľ. and Stambaugh, R.F. (2000) 'Comparing asset pricing models: an investment perspective', *Journal of Financial Economics*, 56(3), pp.335-381.

Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A.V. and Gulin, A. (2018) 'CatBoost: unbiased boosting with categorical features', *Advances in neural information processing systems*, 31.

Qi, Y. (2025) 'MultiModalGraphSearch: Intelligent Massive-Scale SubGraph Discovery for Multi-Category Financial Pattern Mining', *IEEE Access*.

Qin, Y., Song, D., Chen, H., Cheng, W., Jiang, G. and Cottrell, G. (2017) 'A dual-stage attention-based recurrent neural network for time series prediction', *arXiv preprint arXiv:1704.02971*.

Rane, N., Choudhary, S. and Rane, J. (2023) '*Leading-edge artificial intelligence (AI)-powered financial forecasting for shaping the future of investment strategies*'. Available at: SSRN 4640828.

Ren, T., Zhou, R., Jiang, J., Liang, J., Wang, Q. and Peng, Y. (2024) 'RiskMiner: Discovering Formulaic Alphas via Risk Seeking Monte Carlo Tree Search', in *Proceedings of the 5th ACM International Conference on AI in Finance*, pp. 752-760.

Sanghi, N. (2024) *Deep Reinforcement Learning with Python: RLHF for Chatbots and Large Language Models*. Berkeley, CA: Apress.

Santos, G.C., Garruti, D., Barboza, F., de Souza, K.G., Domingos, J.C. and Veiga, A. (2023) 'Management of investment portfolios employing reinforcement learning', *PeerJ Computer Science*, 9, p.e1695.

Savidge, D. (2023) '*Exploring Frequent Batch Auctions*'. Available at: SSRN 4444229.

Scaletta, G. (2024) '*Deep reinforcement learning for portfolio optimization*'. Unpublished doctoral dissertation, Politecnico di Torino.

Shaji, N., Andrade, T., Ribeiro, R.P. and Gama, J. (2022) 'Study on correlation between vehicle emissions and air quality in Porto', in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Cham: Springer Nature Switzerland, pp. 181-196.

Shi, H., Song, W., Zhang, X., Shi, J., Luo, C., Ao, X., Arian, H. and Seco, L.A. (2025) 'Alphaforge: A framework to mine and dynamically combine formulaic alpha factors', in *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39, No. 12, pp. 12524-12532.

Son, M.G. (2022) '*High-Frequency Trading in the Options Market*'. Available at: SSRN 4199462.

Sun, H., Liu, X., Bian, Y., Zhu, P., Cheng, D. and Liang, Y. (2024) 'Asymmetric Graph-Based Deep Reinforcement Learning for Portfolio Optimization', in *Joint European*

Conference on Machine Learning and Knowledge Discovery in Databases, Cham: Springer Nature Switzerland, pp. 174-189.

Sun, R., Stefanidis, A., Jiang, Z. and Su, J. (2025) 'Factor-MCLS: Multi-agent learning system with reward factor matrix and multi-critic framework for dynamic portfolio optimization', *arXiv preprint arXiv:2504.11874*.

Sun, R., Xi, Y., Stefanidis, A., Jiang, Z. and Su, J. (2025) 'A novel multi-agent dynamic portfolio optimization learning system based on hierarchical deep reinforcement learning', *Complex & Intelligent Systems*, 11(7), p.311.

Sun, S., Qin, M., Zhang, W., Xia, H., Zong, C., Ying, J., Xie, Y., Zhao, L., Wang, X. and An, B. (2023) 'Trademaster: A holistic quantitative trading platform empowered by reinforcement learning', *Advances in Neural Information Processing Systems*, 36, pp.59047-59061.

Sun, S., Wang, R. and An, B. (2022) 'Quantitative stock investment by routing uncertainty-aware trading experts: A multi-task learning approach', *arXiv preprint arXiv:2207.07578*.

Tang, Z., Chen, Z., Yang, J., Mai, J., Zheng, Y., Wang, K., Chen, J. and Lin, L. (2025) 'AlphaAgent: LLM-Driven Alpha Mining with Regularized Exploration to Counteract Alpha Decay', *arXiv preprint arXiv:2502.16789*.

Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F. and Rodriguez, A. (2023) 'Llama: Open and efficient foundation language models', *arXiv preprint arXiv:2302.13971*.

Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S. and Bikel, D. (2023) 'Llama 2: Open foundation and fine-tuned chat models', *arXiv preprint arXiv:2307.09288*.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I. (2017) 'Attention is all you need', *Advances in neural information processing systems*, 30.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P. and Bengio, Y. (2017) 'Graph attention networks', *arXiv preprint arXiv:1710.10903*.

Vishal, M., Satija, Y. and Babu, B.S. (2021) 'Trading agent for the Indian stock market scenario using actor-critic based reinforcement learning', in *2021 IEEE international conference on computation system and information technology for sustainable solutions (CSITSS)*, Piscataway, NJ: IEEE, pp. 1-5.

Wan, R., Liu, Y., McQueen, J., Hains, D. and Song, R. (2023) 'Experimentation platforms meet reinforcement learning: Bayesian sequential decision-making for continuous monitoring', in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 5016-5027.

Wang, M., Tian, J., Zhang, M., Guo, J. and Jia, W. (2025) 'Mtmd: multi-scale temporal memory learning and efficient debiasing framework for stock trend forecasting', *IEEE Transactions on Emerging Topics in Computational Intelligence*.

Wang, Z., Li, D. and Li, P. (2022) 'Latent coreset sampling based data-free continual learning', in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 2077-2087.

Wu, J., Shi, W., Cao, X., Chen, J., Lei, W., Zhang, F., Wu, W. and He, X. (2021) 'DisenKGAT: knowledge graph embedding with disentangled graph attention network', in *Proceedings of the 30th ACM international conference on information & knowledge management*, pp. 2140-2149.

Wu, Z., Xiao, C. and Sun, J. (2023) 'Medlink: De-identified patient health record linkage', in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 2672-2682.

Xiang, Q., Chen, Z., Sun, Q. and Jiang, R. (2024) 'Rsap-dfm: Regime-shifting adaptive posterior dynamic factor model for stock returns prediction', in *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*.

Xiao, Y., Sun, E., Luo, D. and Wang, W. (2024) 'TradingAgents: Multi-agents LLM financial trading framework', *arXiv preprint arXiv:2412.20138*.

Xiong, G., Deng, Z., Wang, K., Cao, Y., Li, H., Yu, Y., Peng, X., Lin, M., Smith, K.E., Liu, X.Y. and Huang, J. (2025) 'FLAG-Trader: Fusion LLM-Agent with Gradient-based Reinforcement Learning for Financial Trading', *arXiv preprint arXiv:2502.11433*.

Xu, Y., Hao, J., Tang, K., Chen, J., Liu, A., Liu, P. and Zhang, G. (2025) 'FinRipple: Aligning Large Language Models with Financial Market for Event Ripple Effect Awareness', *arXiv preprint arXiv:2505.23826*.

Zhang, Q., Guo, D., Zhao, X. and Luo, L. (2022) 'Handling RDF streams: Harmonizing subgraph matching, adaptive incremental maintenance, and matching-free updates together', in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 2580-2589.

Zhao, C., Chen, F. and Thuraisingham, B. (2021) 'Fairness-aware online meta-learning', in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 2294-2304.

APPENDIX A SURVEY COVER LETTER

[Your University/Institution Letterhead]

[Date]

[Participant's Name/Salutation, e.g., Dear Finance Professional,]

Subject: Invitation to Participate in Doctoral Research on Agentic AI and Multimodal AI in Quantitative Finance

My name is Banka Tagur Ranga Satish Babu, and I am a Doctorate in Business Administration (DBA) candidate at the Swiss School of Business and Management Geneva. I am conducting research under the supervision of **"Dr. Hemant Palivela"** for my dissertation titled: "Agentic AI and Multimodal AI to Quantitative Finance using Algorithmic Trading for Wealth & Portfolio Management."

The purpose of this research is to explore the current applications, potential benefits, inherent challenges, and future trajectory of integrating advanced Artificial Intelligence paradigms—specifically Agentic AI (systems capable of autonomous goal-oriented action) and Multimodal AI (systems processing diverse data types like text, numbers, and visuals)—within quantitative finance, algorithmic trading, and wealth/portfolio management.

You have been selected to participate in this survey because of your expertise and experience in the fields of finance, technology, artificial intelligence, or investment management. Your insights would be invaluable in understanding the practical implications, opportunities, and concerns associated with these emerging technologies.

This survey is expected to take approximately [e.g., 15-20 minutes] to complete. Your participation is entirely voluntary, and you may choose to stop at any time or skip any questions you are not comfortable answering. All responses will be kept confidential and anonymous. The data collected will be aggregated and used solely for the purpose of this academic research, and any publications or presentations resulting from this study will not contain any information that could personally identify you.

There are no direct risks anticipated from participating in this survey beyond the time taken to complete it. Your honest and thoughtful responses will contribute significantly to a deeper academic understanding of a rapidly evolving and impactful area of financial technology.

If you have any questions about the research or this survey, please do not hesitate to contact me at satishbanka.ai@gmail.com or my dissertation chair, Hemant Palivela, at Hemant.datascience@gmail.com.

To proceed with the survey, please click on the following link: **[Insert Survey Link Here - if applicable, otherwise describe distribution method]**

Thank you for considering this invitation and for your valuable contribution to this doctoral research.

Sincerely,

Banka Tagur Ranga Satish Babu DBA Candidate at Swiss School of Business and Management Geneva [satishbanka.ai@gmail.com] [+917798846835]

APPENDIX B INFORMED CONSENT

Informed Consent Form for Participation in Doctoral Research

Research Project Title: Agentic AI and Multimodal AI to Quantitative Finance using Algorithmic Trading for Wealth & Portfolio Management

Principal Investigator: Banka Tagur Ranga Satish Babu, DBA Candidate, Swiss School of Business and Management Geneva **Faculty Supervisor:** Dr. Hemant Palvela, Dr., Swiss School of Business and Management Geneva

1. Purpose of the Research: You are invited to participate in a research study investigating the applications, benefits, challenges, and future directions of Agentic AI and Multimodal AI in quantitative finance, algorithmic trading, and wealth/portfolio management. The aim is to understand the transformative potential and practical considerations of these advanced AI technologies in the financial industry.

2. Procedures: If you agree to participate, you will be asked to complete an anonymous survey that will take approximately [e.g., 15-20 minutes]. The survey will include questions about your perspectives on AI in finance, its current uses, perceived benefits, potential risks, and expectations for the future. *[If interviews are also part of the methodology, a separate section for interview consent would be needed, or this section would be expanded to describe interview procedures, estimated time, audio recording, etc.]*

3. Voluntary Participation: Your participation in this study is entirely voluntary. You may refuse to participate or withdraw from the study at any time without any penalty or loss of benefits to which you are otherwise entitled. You may also choose to skip any specific questions in the survey that you do not wish to answer.

4. Confidentiality and Anonymity: All information you provide will be kept confidential. Survey responses will be collected anonymously. No personally identifiable information will be linked to your responses in any reports, publications, or presentations resulting from this research. Data will be stored securely on password-protected devices/servers accessible only to the principal investigator and faculty supervisor. Aggregated and anonymized data may be used for academic publications and presentations.

5. Risks and Benefits: There are no foreseeable direct risks to you from participating in this study other than the time commitment involved in completing the survey. While you may not receive any direct personal benefit from participating, your responses will contribute to a better academic understanding of the role and impact of advanced AI in the financial sector. This research may help inform industry best practices and future technological developments.

6. Contact Information: If you have any questions about this research study, your rights as a participant, or if you wish to withdraw from the study, please contact: Banka Tagur Ranga Satish Babu at satishbanka.ai@gmail.com or Dr. Hemant Palivela at Hemant.datascience@gmail.com, Swiss School of Business and Management Geneva.

If you have concerns about the ethical conduct of this research, please contact the [Relevant Ethics Board/Review Committee at SSBM, if applicable, along with their contact details].

7. Consent Statement: By proceeding with and completing this survey, you acknowledge that:

- You have read and understood the information provided in this Informed Consent Form.
- You have had the opportunity to ask questions and have had them answered to your satisfaction.
- You understand that your participation is voluntary and that you can withdraw at any time.
- You understand that your responses will be kept confidential and anonymous.
- You are at least 18 years of age.
- You agree to participate in this research study.

[If this were an online survey, ticking a box "I have read and agree to the terms above" would typically follow. For a paper survey, a signature line would be here.]

Thank you for your willingness to contribute to this research.

APPENDIX C: INTERVIEW GUIDE

Interview Guide: Perspectives on Agentic AI and Multimodal AI in Quantitative Finance

Research Project Title: Agentic AI and Multimodal AI to Quantitative Finance using Algorithmic Trading for Wealth & Portfolio Management **Interviewer:** Banka Tagur Ranga Satish Babu, DBA Candidate **Interviewee:** [Name/Role of Expert] **Date:** [Date of Interview] **Location/Method:** [e.g., Zoom, In-Person]

I. Introduction (5 minutes)

- Thank the interviewee for their time and participation.
- Briefly re-introduce myself and the purpose of the DBA research: To explore the applications, benefits, challenges, and future trajectory of Agentic AI and Multimodal AI in quantitative finance, algorithmic trading, and wealth/portfolio management.
- Explain the interview format: Semi-structured discussion, estimated duration [e.g., 45-60 minutes].
- Confirm understanding and consent for audio recording (if applicable) and assure confidentiality and anonymity in reporting findings (referencing the Informed Consent Form).
- Invite any initial questions from the interviewee.

II. Current Landscape and Understanding (10-15 minutes)

1. From your perspective, what are the most significant ways you see AI, particularly more advanced forms like Agentic AI (autonomous, goal-oriented systems) and Multimodal AI (systems processing diverse data like text, numbers, visuals), currently impacting or beginning to impact quantitative finance and wealth/portfolio management?
2. How would you define "Agentic AI" and "Multimodal AI" in the context of practical financial applications? Are these concepts well-understood in the industry?
3. What specific types of financial tasks or problems do you believe are most ripe for disruption or enhancement by these advanced AI paradigms?

III. Benefits and Opportunities (10-15 minutes)

4. What do you perceive as the primary benefits or opportunities that Agentic and Multimodal AI could bring to: * Algorithmic trading strategies (e.g., in terms of alpha generation, risk management, efficiency)? * Wealth and portfolio management (e.g., personalization, client advisory, decision support)?

5. Are there specific types of trading strategies (e.g., factor-based, event-driven, HFT) or portfolio management approaches that stand to gain the most from these technologies? Why?

6. How might these AI systems enhance the ability to adapt to changing market conditions or discover novel investment insights that are difficult for humans or traditional models to find?

IV. Challenges, Risks, and Limitations (10-15 minutes)

7. What are the major challenges or hurdles financial institutions face in developing and deploying Agentic and Multimodal AI systems effectively? (Probes: Data quality/availability, model complexity, integration with legacy systems, talent gap, computational cost).

8. In your view, what are the most significant risks associated with the increasing autonomy and complexity of AI in financial decision-making? (Probes: Model "black boxes"/interpretability, reliability, potential for systemic risk, accountability).

9. How significant are concerns around data bias being amplified by these AI systems in financial contexts? What about data privacy?

10. What are the limitations of current Agentic or Multimodal AI approaches when applied to the nuances of financial markets?

V. Human Role and Alignment (5-10 minutes)

11. How do you see the role of human quantitative analysts, traders, and portfolio managers evolving alongside these increasingly capable AI systems? Will AI augment or replace human roles?

12. The concept of Reinforcement Learning from Human Feedback (RLHF) is proposed to align AI with complex human preferences. How feasible and effective do you think RLHF could be in practical financial settings for goals like aligning with nuanced risk profiles or ethical considerations? What are the challenges?

VI. Future Outlook and Ethical Considerations (5-10 minutes)

13. Looking ahead 5-10 years, what is your vision for the role of Agentic and Multimodal AI in shaping the future of quantitative finance? What breakthroughs or major shifts do you anticipate?

14. What are the most pressing ethical considerations that the industry and regulators need to address as these technologies become more powerful and widespread? (Probes: Fairness, transparency, accountability, market stability, governance).

15. What kind of governance structures or regulatory adaptations do you believe are necessary to ensure the responsible innovation and deployment of AI in finance?

VII. Conclusion (5 minutes)

- Are there any other critical aspects or perspectives on this topic that we haven't discussed but you feel are important?
- Thank the interviewee again for their valuable time and insights.
- Explain how the information will be used (anonymized, aggregated for the dissertation).

Offer to share a summary of the research findings when completed (if appropriate and desired).