Veikko Kytölä & Juuso Valtanen

Algorithmic Rationality vs. Human Bias

A Comparative Review of Robo-Advisors in Portfolio Management

The School of Accounting and Finance
Bachelor’s Thesis Finance
Accounting and Finance

Vaasa 2025

UNIVERSITY OF VAASA
The School of Accounting and Finance

Author:	Veikko Kytölä & Juuso Valtanen
Title of the thesis:	Algorithmic Rationality vs. Human Bias: A Comparative Review of Robo-Advisors in Portfolio Management
Degree:	Bachelor of Science in Economics and Business Administration
Degree Programme:	Finance
Supervisor:	Niranjan Sapkota
Year:	2025 **Pages:** 35

ABSTRACT:

This thesis addresses a major conflict in modern finance. It is the gap between the ideas of Modern Portfolio Theory and the real-world investor decision-making. Cognitive biases are often shown by investors and therefore traditional human advisors were compared with algorithmic models. These models range from basic robo-advisors to advanced Agentic AI systems. The goal was to examine which option performs better and reduces investor biases more effectively.

The results of the algorithm models show a clear advantage. They are used as commitment devices and to remove biases such as the Disposition Effect and overconfidence. This advantage is deepened by the emergence of Generative AI and Large Language Models. Algorithms can use these technologies to interpret market news free from emotional bias. They also predict returns more accurately than humans. Agentic AI introduces autonomous cost-efficiency by automating complex financial workflows without human participation.

Relying only on technology is not enough. Algorithms bring in new risks, such as the "black box" problem and probable biases in the training data. Therefore, human advisors are still needed. They provide complete financial planning, explain complex model results, and offer psychological support. As a result, the thesis suggests a solution. A Hybrid Model, which combines the efficiency of AI with human decision-making and empathy, is the best approach for future portfolio management.

KEYWORDS: Artificial intelligence, Security portfolios, Machine learning, Investment, Performance

Contents

Tables

1 Introduction

Modern portfolio management has been a key topic in the financial literature since the introduction of Modern Portfolio Theory (MPT). MPT established the basis for rational, optimized investment decisions. The classical view is that investors look to maximize returns for a given level of risk, assuming that decisions are made with perfect market rationality (Markowitz, 1952).

However, the assumptions of MPT are often questioned by evidence from behavioral finance, which shows that people are affected by behavioral biases when making decisions (Kahneman & Tversky, 1979). The traditional human-led model tends to fail to achieve the best results for investors, because of these biases. Therefore, algorithmic solutions have grown quickly in the FinTech sector, while traditional advisory services have been struggling (D'Acunto et al., 2019).

In this thesis we compare two different models. One of these models is a traditional human advisor, described by Gennaioli et al. (2015) as a Money Doctor model. The counterpart for this model is an algorithmic system. It is a digital platform that can give automated and customized recommendations without human interference (D'Acunto et al., 2019). This alternative model to humans, ranges from basic rule-based robo-advisors to more advanced Generative AI and also Agentic AI systems. The main issue is that machines offer technical objectivity, while humans provide the subjective and holistic value that has long justified higher fees.

Our thesis's main goal is to identify the main differences between algorithmic and human advisory models in portfolio management. We will examine how these differences affect the areas of cost-efficiency, performance, and minimized investor biases. With this framework we can examine the proposed hypotheses and understand how algorithms impact the financial field.

1.1 Purpose of the Study

The purpose of this thesis is to examine whether the main theoretical benefits of algorithmic logic, like data processing and generating alpha, give a real advantage in cost-effectiveness and overall performance. We start the thesis by introducing the theory and then move on to look at different financial results. The first hypothesis is:

H_1: Algorithmic portfolio management achieves superior cost-efficiency and performance compared to human-led methods.

Second, we study human behavioral issues and ask if algorithms that use consistent logic can help reduce the behavioral mistakes that human investors tend to make. If the algorithm can keep investors disciplined, it should remove errors that people usually find hard to avoid due to their human nature. This leads to the second hypothesis:

H_2: The consistency of algorithmic advisory services significantly reduces common investor cognitive biases.

Many studies show that people exhibit systematic cognitive biases (Kahneman & Tversky, 1979) and demonstrate that advanced Generative AI and machine learning models are capable of making strong predictions (Gu et al., 2020). This research targets to compare how these two different models perform and how they approach risk and decision processes.

1.2 Structure of the Thesis

This thesis is structured into five main sections, each aimed at answering a research question. First, Chapter 1 presents the topic, motivation, and hypothesis. Chapter 2 follows and establishes the underlying theory for the work by discussing key models in behavioral economics and showing how cognitive biases make the human investor imperfect. Chapter 3 introduces the technological solution of robo-advice. It defines automated investment advice and explains in depth the role of artificial intelligence and machine learning as the engine of the portfolio management process. Chapter 4 contrasts human and algorithmic models across three dimensions: rationality versus cognitive biases, performance and cost-efficiency, and algorithmic challenges. Finally, Chapter 5 summarizes the paper's key conclusions.

This thesis was written by Veikko Kytölä and Juuso Valtanen as a pair and together, we developed the research topic, theory, and hypotheses. We both contributed equally to the literature search and analysis. We worked on the thesis at the same time, drafting and editing together to ensure a consistent style.

Since Agentic AI, Generative AI and Large Language Models are new technologies, high-rated academic sources are limited. Therefore, this thesis uses recent working papers to ensure the analysis is up to date. Also for greater clarity and academic language, we used a language model, Gemini, developed by Google, for grammar checks and proofreading. The model was not used to generate content, conduct research, or perform analysis.

2 Human-Led Portfolio Management

This chapter examines the behavioral side of human-led portfolio management to understand the patterns and biases that affect results. It lays the foundation for comparative analysis by first defining traditional portfolio management. Next, it examines how financial theory has evolved to a more Behavioral Finance approach, demonstrating investor irrationality. Finally, the chapter presents empirical evidence for specific behavioral biases to show that the human investor is systematically irrational.

2.1 Traditional Portfolio Management

The traditional academic definition of portfolio management is based on Markowitz's Modern Portfolio Theory (MPT). His model suggests that a portfolio manager's task is a mathematical optimization problem. In which they build the client an Efficient Frontier by diversifying the portfolio at a given risk to achieve the optimal return. The portfolio manager's value is therefore based on their technical ability to analyze returns, variances, and correlations, and thus does not take psychological factors into account. (Markowitz, 1952).

Markowitz's purely technical and rational view has been challenged because it fails to explain why clients pay high fees to human advisors, who often underperform the markets after fees. The alternative psychological model defines the portfolio manager as a Money Doctor whose primary task is not technical optimization but building trust with the client. The Money Doctor's service aims to ease anxiety and provide clients with the self-confidence to take on risks they would normally avoid. (Gennaioli et al., 2015).

However, Gennaioli et al. (2015) Money Doctors model's core is maintaining client trust, not profit maximization. This creates a conflict where managers could have a strong financial incentive to cater to clients' cognitive biases. For example they could sell the client a desired but overpriced product, when they really should correct the client's irrational view. After challenging the client's view they could potentially break the trust relationship, which is the foundation of the portfolio manager's profits.

Financial theory is increasingly moving away from the EMH, which assumed investors acted as rational profit-maximizers, to more psychological focus (Fama, 1970). Behavioral Finance and Prospect Theory challenged this idea by showing that investors are loss-averse, feeling losses more strongly than gains, and evaluate risk based on a specific reference point (Kahneman & Tversky, 1979). Thaler (1985) later introduced Mental Accounting, showing that individuals irrationally earmark fungible money for different purposes. Together, these theories help explain the cognitive biases that make the psychological role of the human advisor important.

2.2 Investor Behavioral Biases

Section 2.2 examined how Behavioral Finance challenged traditional, investor-based financial theory. Theories like Prospect Theory and Mental Accounting help explain why investors are imperfect decision-makers. In this section, we will examine specific cognitive biases and the behaviors they cause. We will focus on three key biases for this study: overconfidence, the Disposition Effect, and herding.

2.2.1 Overconfidence

In finance, overconfidence is one of the most studied cognitive biases. Barber and Odean (2001) found that investors often overestimate how accurate their knowledge is and how well they can predict future security performance. This makes them trade too much. The study shows that frequent trading reduces net returns due to transaction costs.

In their essential study on this topic, "Boys will be boys", Barber and Odean (2001) faced the challenge of measuring overconfidence directly. They used gender as a proxy for overconfidence. This was because psychological research shows that men are usually more overconfident than women. The researchers thought that men would trade too much and therefore undermine their returns more than women. The empirical result shows that they were right. Men traded 45% more than women and gained lower net returns because of that. Barber and Odean (2001) note that the result did not only concern the risk. Usually rational models suggest that avoiding risk leads to lower returns but in this study, both groups actually did worse than their benchmarks.

Grinblatt and Keloharju (2009) deepened the analysis related to overconfidence. Instead of relying on the proxy, they used a unique dataset to measure overconfidence directly. They also looked at sensation seeking, which was measured by the number of speeding tickets. Their findings confirm that both overconfidence and sensation seeking correlated with more frequent trading and negative returns. However, they were able to differentiate the two. For example, a one-unit increase in the overconfidence measure raised trading volume by 4%, while one additional speeding ticket, measuring sensation seeking, raised it by 7%. The results show that sensation seeking was a better predictor of high portfolio turnover. This suggests that excessive trading is not only driven by overconfidence but also by the excitement of trading itself.

Although the measurement methods differ – Barber and Odean (2001) used gender as an indirect measure while Grinblatt and Keloharju (2009) used direct psychological behavioral measures – both studies come to the same conclusion. Overconfidence and related psychological traits led to excessive trading, which reduces returns due to transaction costs.

2.2.2 Loss Aversion & Disposition Effect

Besides overconfidence, the Disposition Effect is one of the most significant cognitive biases in finance. Shefrin and Statman (1985, p. 777) called this phenomenon the tendency to “sell winners too early and ride losers too long”. Their theory is based directly on Kahneman and Tversky's (1979) Prospect Theory. Because investors are risk-averse when it comes to gains, they realize even small profits quickly. When facing losses, investors tend to take more risks. This is why they often keep their losing investments, hoping to "break even". Shefrin and Statman (1985) explained that this investment bias is strengthened by two factors: Mental Accounting, which is the reluctance to accept a loss, and Regret Aversion, which is the fear of admitting the original investment was a mistake.

This theoretical model received extensive empirical confirmation by Odean's (1998) classic study, which analyzed 10,000 investors' trading data. Odean (1998) demonstrated that investors sold their winning stocks 50% more likely than their losing stocks (Proportion of Gains Realized = 0.148 vs. Proportion of Losses Realized = 0.098). The main finding of the study was that rational explanations did not account for the phenomenon. It was not caused by portfolio rebalancing, higher transaction costs, or belief in mean reversion. In fact, this behavior was financially suboptimal, since the winning stocks that were sold gained, on average, 3.4% more than the losing stocks held (Odean, 1998).

The most substantial evidence for psychological factors is the so-called December effect, where tax planning is involved. Odean (1998) explains that rational tax planning requires steady selling of losses throughout the year to maximize tax benefits (Constantinides, 1984). However, Odean's (1998) data showed that loss realizations were dramatically concentrated in December. This finding supports Shefrin and Statman's (1985) theory of self-control. Instead of being a logical time to optimize, the end of the tax year acts as an outside deadline that finally pushes investors to make decisions they have put off all year because of the Disposition Effect (Odean, 1998; Shefrin & Statman, 1985).

2.2.3 Herding

While the Disposition Effect focuses on the investor's internal psychology, herding explains how social interactions between investors lead to irrational decisions. Herding is the tendency to ignore one's own private information and to follow other investors' decisions. This can be explained by informational cascades and conversation (Shiller, 1995). He states that investing is rarely done independently. Investors talk to each other and try to figure out what they know. As Shiller (1995) explains, citing the classic experiments of Asch (1952), if someone sees an entire group acting differently from them, they may doubt their own judgment and follow the herd, even if the herd is wrong.

Although this social interaction explains part of herd behavior, Scharfstein and Stein (1990) introduce a model based on reputational concerns. It explains why even professional portfolio managers follow the herd. According to their theory, the manager's compensation is based on the labor market's assessment of their ability (smart vs. dumb manager). In this model, it is often safer for the portfolio manager to follow other professionals, even though they would act differently as individuals. If they make the same mistake as the other professionals, they can share the blame. The model suggests that if a manager makes the same mistake as everyone else, they can share the blame. In contrast, if they make a mistake unconventionally, they will be seen more likely as dumb. Keynes (Scharfstein & Stein, 1990, p. 465) summarized the issue: "Worldly wisdom teaches that it is better for reputation to fail conventionally than to succeed unconventionally."

While Shiller (1995) and Scharfstein and Stein (1990) approach herd behavior from different angles, their conclusions are pretty similar. Shiller focuses on the social interactions and information cascades. He examines how they drive individual investors toward collective movement and market volatility. In contrast, Scharfstein and Stein reveal the structural incentives that lead even professionals to mimic one another. In their model, this is caused by rational reputational pressure: failing along with the herd is safer for the manager's career than failing alone by swimming against the tide. Scharfstein and Stein (1990) explain that although herd behavior can be rational, it is socially inefficient because it causes investment decisions to lose touch with real fundamental factors.

2.3 The Human Investor

Recent research on behavioral biases provides strong support for this idea. A primary driver for human behavior is overconfidence. It causes people to trade too much and that hurts their net returns (Barber & Odean, 2001; Grinblatt & Keloharju, 2009). Also, a poor financial strategy is the Disposition Effect which makes investors to "sell winners too early and ride losers too long," (Odean, 1998, p. 777). It is also linked to psychological patterns like Mental Accounting and Regret Aversion (Shefrin & Statman, 1985). Beyond individual biases is herding behavior. It is influenced by social and reputational pressures and gets investors to follow the crowd (Shiller, 1995; Scharfstein & Stein, 1990). This disconnects investment choices from economic fundamentals. Ultimately these biases and emotional responses show why a solely human judgement in the financial environment should not be relied on. A need rises for a better solution and therefore the following chapter will introduce an algorithmic approach designed to tackle these human limitations.

3 Algorithmic Portfolio Management

As a response to tackle the human irrationalities and biases, this chapter introduces the basics of algorithmic portfolio management. We start by defining what a robo-advisor is and what role it plays in automating portfolio management. Then we take a closer look into a detailed description of the mechanical multi-step process. In the end we jump to the current day and most modern versions of robo-advisors to discuss the function of artificial intelligence and machine learning behind algorithms. We examine how different learning models and artificial intelligence can outshine classic rule-based models in both forecast precision and alpha generation. This chapter follows the evolution of traditional human portfolio management by illustrating the technological evolutions toward more a robust way of portfolio management.

3.1 Definition of Algorithmic Services

Algorithmic portfolio management’s core is artificial intelligence (AI). It is dedicated to emulating human mental processes like complex human processes, problem-solving, and pattern recognition. The term AI has many subfields. Most notable of these subfields are machine learning (ML) and deep learning (DL). These technologies receive actionable outputs by using immense computing power to go through big datasets, learning from external data sources. However, Kaplan and Haenlein (2019) state that it is important to recognize that AI and these systems are always dependent on the training data that is given to them. That is why AI functions as an amazing processor of current data and is effective at pattern recognition, but it is unable to produce its own information - at least not yet.

The swift technological progress has also found its footing in the financial sector with the emergence of robo-advisors. They are digital services that use algorithms to automate tailored financial advice and portfolio management and they do all this without constant human interaction (D’Acunto et al., 2019; Jung et al., 2018). They provide a modern alternative to the traditional person-to-person service model (Jung et al., 2018).

Robo-advisors replicate the essential functions of traditional portfolio management, including client profiling, portfolio construction, and periodic rebalancing. The difference is that they do it mechanically. These platforms operate with a level of consistency that is very difficult for an individual human to maintain. By strictly following a predefined set of rules while managing a portfolio robo-advisors can eliminate the feelings and thought-based biases, that commonly weaken human decision-making. (D’Acunto et al., 2019).

3.2 The Algorithm's Operating Process

Robo-advisors make objective investment decisions using a consistent and automated process. Depending on whether the portfolio is managed passively or actively, robo-advisors are split into two different types. The two strategies have a main difference in the degree of algorithm’s involvement. Passive robo-advisors favor products that require less active management. Examples of such products are Exchange Traded Funds (ETFs). Passive strategies usually remain static after the account setup and focus on instrument and allocation selection rather than ongoing trading. (Jung et al., 2018).

Compared to passive robo-advisors, active robo-advisors use more complex models like AI (D’Acunto et al., 2019). Their paper shows that active robo-advisors make decisions about security selection or market timing more actively than passively. This active strategy’s goal is generating excess returns and alpha. The process of both strategies follows the same three-stage sequence of profiling, construction, and maintenance. That means that the underlying process remains almost the same in both strategies (Jung et al., 2018).

As mentioned before, the process starts with the customer’s risk profiling (Jung et al., 2018). The algorithm does this by using a digital survey to collect information about the client's risk tolerance, investment goals, and time horizon (D’Acunto et al., 2019; Jung et al., 2018). While this digital assessment primarily maps the client’s risks, the most highly developed algorithms can directly infer the client's risk appetite when constructing the portfolio. After this assessment is completed, the robo-advisor automatically builds an investment portfolio. Most robo-advisors use portfolio strategies based on the MPT, as described in Chapter 2 (D'Acunto et al., 2019). MPT was the first major mathematical model for robo-advisors, but methods and processing power have advanced a lot over the past decades (Salo et al., 2024). Currently, the portfolio building includes more advanced techniques, such as deep reinforcement learning, to handle ongoing investment decisions (Hambly et al., 2023).

After the portfolio is set up, the robo-advisor is responsible for its ongoing maintenance, also known as the maintenance phase (Jung et al., 2018; Salo et al., 2024). As portfolio values shift and allocation changes from its target weights, automatic rebalancing is performed according to pre-set rules. According to Jung et al. (2018) rebalancing is usually grouped into two groups: passive or active. Passive rebalancing is fully automatic, while active rebalancing occurs when the investor receives suggestions and then decides whether to act. This process also assures that the client's portfolio stays steady with their original risk profile. Because algorithms can handle very short holding periods, they need to rebalance portfolios automatically (Krauss et al., 2017). The core of automated maintenance is therefore tasks where assets are reallocated at every point in time (Jung et al., 2018). Robo-advisors eliminate human cognitive biases that often weaken investment decisions, by using these maintenance strategies (D'Acunto et al., 2019).

3.3 AI and Machine Learning Behind the Algorithms

The necessary framework and theory for portfolio diversification and risk management is based on the MPT, which is a fixed system based on set rules (Markowitz, 1952). MPT focuses on optimizing variance using historical data, but it does not consider daily new information and complicated market conditions (Salo et al., 2024). This is why algorithmic portfolio management has benefited most from progress of artificial intelligence. Therefore, we will discuss AI and especially ML and DL in more detail in this chapter.

3.3.1 The Role of ML/DL

ML has quickly developed as an important tool in the financial sector to help with decision-making and strategies. ML covers the term DL as well (Aziz et al., 2022). Their strengths lie in their ability to manage huge amounts of data and find patterns that are not visible to human analysts or simpler models (Gu et al., 2020). The difference between simpler models and learning models is the ability to learn from external data sources (Hambly et al., 2023). Learning models are adaptive and dynamic, but not as flexible as neural networks. Neural networks mimic the biological brain but are actually mathematical systems that consist of neurons or nodes. They are often described as universal approximators. These models have theoretical power to map out virtually any complex relationship allowing them to find non-linear market trends (Gu et al., 2020; Hambly et al., 2023). This enables portfolio management to move from fixed rules toward a dynamic model that can constantly respond to new market information when it becomes available.

3.3.2 ML's Predictive Power

Gu et al. (2020) demonstrated that ML models, such as deep neural networks significantly surpass traditional econometric models. Their advantage comes from flexibility. They can process immense amounts of heterogeneous information. These models are also able to capture asset returns far more effectively than linear models. This is because they can include correlations between predictors that typical regression methods miss (Gu et al., 2020; Krauss et al., 2017). Simpler models usually do not consider these important interactions.

The advanced models' ability to predict is confirmed by many statistics. Gu et al. (2020) demonstrate that ML models can achieve monthly stock-level predictive R^2 values as high as 0.40%, which is significantly higher than those achieved by classic linear models. For instance, a strong linear benchmark model obtains an R^2 of only 0.16% per month. An OLS model is a basic statistical technique that finds the straight line that best fits historical data. It uses the full 900+ input features without regularization, but still fails, producing a negative R^2 of -3.46% per month. This means it does worse than a simple zero forecast. This clear statistical edge is the foundation of the measurable financial gains ML models produce.

3.3.3 Alpha Generation

The financial benefits from ML predictions are strong. Gu et al. (2020) demonstrate that a strategy using neural network forecasts to time the S&P 500 index achieved an annualized out-of-sample Sharpe ratio of 0.77. They compared it to a buy-and-hold investor, who achieved a Sharpe ratio of 0.51. The third strategy compared was the value-weighted long-short strategy that also used neural network predictions to select stocks. It earned an annualized out-of-sample Sharpe ratio of 1.35. Based on these results algorithms could have a competitive edge beyond classic strategies.

More evidence is presented by Gu et al. (2020), who note that ML portfolios usually produce significant alphas compared to factor pricing models. This alpha is not just a reward for taking on known systemic risk, according to the study. For example, complex factor models explain only 10% to 30% of the variation in returns generated using neural network-based portfolios (Gu et al., 2020). The six-factor model is an example of such a model. A substantial 70% to 90% of these returns still remain unexplained by established market risks. This suggests that traditional factor models offer only a partial explanation for the better performance of ML portfolios.

The progressive advancement of portfolio optimization has led to the reinforcement learning (RL) being integrated into algorithmic trading (Hambly et al., 2023). The study explains that RL uses a persistent learning approach based on environment to maximize long-term portfolio returns. This approach shifts from fixed optimization to dynamic and adaptive decision-making and is one of the most advanced approaches for dynamic portfolio optimization. Because these algorithms hold assets for only short periods, the continuous and automatic rebalancing discussed in Chapter 3.2 is needed.

3.4 Evolution of Algorithmic Services

Traditional robo-advisors rely on simple risk surveys to assign investors to static portfolios (Jung et al., 2018). Linear operations are transformed into dynamic systems with the help of advanced AI technologies. New models use more Natural Language Processing (NLP), enabling computers to read unstructured text such as news and reports (Sheng et al., 2025). NLP is also the power behind Large Language Models (LLMs). These systems are trained on massive datasets to understand complex patterns in human language (Lopez-Lira & Tang, 2025). Liu and Shi (2025) also add to this. With these models, market risks and volatility are easier to understand. This allows capturing market sentiment more effectively than traditional methods.

Generative AI is defined by Sheng et al. (2025) as a technology that creates new analysis and content. Processing numbers is not its only job. By analyzing firm-specific information, Generative AI can generate abnormal returns. It is also able to explain the reasons behind investment decisions.

An autonomous system that can make independent decisions to work towards specific goals is called Agentic AI. It represents an important evolution in algorithmic services. For example, without human control, Agentic AI can execute portfolio adjustments. Complex financial workflows can also be automated by these agents. This shifts the algorithm's role. From a simple tool, it becomes an independent operator. (Luqman et al., 2025).

4 Comparative Analysis: Algorithm vs. Human

Comparing algorithmic and human advisory services is an important topic in financial research today. The main question is, can the technical strengths of advanced algorithms exceed the emotional benefits of human advisors. In this chapter, we aim to examine the hypotheses. We look at two main areas to do this. These are rationality versus cognitive biases, and performance and cost-efficiency.

4.1 Rationality vs. Cognitive Biases

As presented in Chapter 2, the main problem of the human-led model is its vulnerability to cognitive biases. Human financial advisors can fall into these biases. According to D'Acunto et al. (2019), they may even transfer their own biases to their clients. The Money Doctor model shows that portfolio managers may have a strong financial incentive to take advantage of their clients' existing irrationality rather than correct them (Gennaioli et al., 2015). In contrast, algorithms used in robo-advising are designed to avoid these human mistakes.

As we explained in Chapter 3, algorithmic portfolio management gives us a mechanical solution to this psychological problem. The algorithm serves as a commitment device. It helps settle the conflict among emotional impulses and rational decision-making (Shefrin & Statman, 1985). This is especially clear in how it counters the Disposition Effect. While a human investor might avoid closing a losing Mental Account due to regret aversion (Shefrin & Statman, 1985), the algorithm sticks to its set rebalancing rules (Jung et al., 2018). It sells the losing asset once its share falls below the optimal allocation, an action the human bias may resist. The algorithm is also able to avoid excessive trading caused by overconfidence (Barber & Odean, 2001). In addition, it ignores the social pressures of herding (Scharfstein & Stein, 1990). Because it follows only its programmed data and a predefined allocation plan, it is incapable of feeling the emotional impulses that lead to these costly human errors (D'Acunto et al., 2019).

New empirical evidence confirms that Generative AI removes the emotional optimism bias that is typical for humans. Hsu et al. (2025) compared investment recommendations that GPT-4 models made against human analysts. They found a clear difference in objectivity. Human analysts issued "Strong Buy" recommendations for 19.23% of stocks, while GPT-4 models were much more careful, issuing "Strong Buy" ratings for only 1.67% to 5% of cases. Furthermore, while humans rarely recommended selling (only 0.30% "Sell" ratings), AI models issued "Sell" recommendations in 10% of cases, especially for stocks that were performing poorly. The human reluctance to sell is because investors are prone to the Disposition Effect (Shefrin & Statman, 1985). This data proves that algorithmic logic acts as a stricter screening tool that filters out the "optimism bias" found in human advisory (Hsu et al., 2025).

This rationality also applies to how market news is analyzed. Lopez-Lira and Tang (2025) found that LLMs can predict stock price movements from news headlines with a 93.3% success rate for initial market reactions. While human traders may panic or overreact to news, the LLM strategy achieved returns that were only based on economic facts. This effectively separates decision-making from human emotional noise (Lopez-Lira & Tang, 2025).

The idea that algorithms can reduce bias is supported by several studies. Rossi and Utkus (2020) studied investors who adopted a robo-advisor. They found that investors who use robo-advising showed a significant decline in the Disposition Effect and rank-order effect bias. The analysis shows that the benefits are best for investors whose behavior has been most biased before. This finding is supported by the D'Hondt et al. (2020) study, which utilized "AI alter ego" models. They found that algorithmic advice particularly benefits low-education and low-income investor groups, which are more likely to make cognitive errors. R. Chen and Ren (2022) also confirm this by studying AI-powered funds. According to them, AI can reduce the most common cognitive biases, such as the Disposition Effect, while also lowering transaction costs. Overall, these studies strongly support the second hypothesis and show that algorithms can fix the main systematic errors in the human-led model discussed in Chapter 2.

4.2 Performance and Cost-Efficiency

To examine the first hypothesis, we need to compare algorithmic and human advisory services across two key areas: cost-efficiency and performance. The analysis begins with the fundamental pricing difference. The cost difference between algorithmic and human advisory services is a substantial competitive factor.

The study by R. Chen and Ren (2022) presents comparative evidence on the cost and the efficiency of algorithmic services versus traditional portfolio management. They systematically studied AI-powered mutual funds that use ML for active stock selection and compared them to human-managed peer funds. The study analyzes the differences between the two types of funds over the sample period from 2017 to 2019. The table below provides three key metrics: expense percentage, turnover percentage, and number of stocks held.

Table 1. Differences in structure and cost-efficiency between AI-powered funds and human-managed peers (R. Chen & Ren, 2022, p. 6).

Fund Characteristic	AI-Powered Funds	Human-Managed Peer Funds	Difference (AI-Peer)
Expense (%)	0.37	0.40	-0.03
Turnover (%)	31.06	72.38	-41.32
Number of Stocks Held	149.05	197.47	-48.41

AI-powered funds have lower expenses than human-managed peer funds. The table shows an average difference of 0.03 percentage points in AI's favor. However, the turnover percentage holds the biggest difference. AI-powered funds average 31.06%, while human-managed peers average 72.38%. Trading fewer stocks and doing so less frequently can mean lower turnover. This leads to lower costs. Table also shows that AI-powered funds hold significantly fewer stocks (149.05 compared to 197.47). This suggests a strategy that is more selective and cost-conscious. Compounding decreases the final investment value, meaning that the differences between these two strategies are important. (R. Chen & Ren, 2022).

The cost advantage is growing. The industry is moving from automation to total autonomy with the help of Agentic AI. According to Luqman et al. (2025), these systems can run multi-step financial workflows on their own. They can for example, run compliance audits and real-time risk modeling. These tasks normally would require expensive human work. With the help and evolution of this Agentic AI, financial institutions can cut running costs significantly. M. Chen (2025) supports this, noting that Agentic AI enables a "human-plus-AI" model. In this model, AI agents handle most routine investment decisions. This allows firms to manage more assets with fewer employees, creating a structural cost advantage that traditional firms cannot match.

The data also show that algorithmic services are less expensive, but it also points out that human advisory costs are higher, partly because of high turnover and greater expenses. The reason for the higher cost of human advisors must therefore also be explained by the Money Doctor concept from Chapter 2. The higher fee does not come solely from technical portfolio management, but also from the value of the psychological support that advisors provide. The cost mainly covers the client's peace of mind and the maintenance of a trust. This approach is called comprehensive wealth management, which goes beyond the technical portfolio management offered by algorithms. Therefore, the higher cost or premium fees are justified. (Gennaioli et al., 2015).

While the Money Doctor concept explains the higher fee for a human advisor, it does not address financial performance. Barras et al. (2010) studied actively managed funds, analyzing 2,076 mutual funds over the sample period from 1975 to 2006. The study shows that cost overrides skill. They found that 9.6% of fund managers could generate alpha before fees, but this dropped to only 0.6% after accounting for management fees and expenses. This is strong evidence that actively managed funds fail to deliver positive returns after costs, as high fees wipe out any extra profit managers may generate. (Barras et al., 2010).

The evidence above highlights the inefficiency that algorithmic services are designed to solve. When comparing algorithmic services to human advisory services, one must separate the two main strategies. The analysis should focus on the typically passive nature of robo-advisors and the performance potential of ML models. Most passive robo-advisors have a principal purpose. It is cost-efficient portfolio maintenance. Instead of doing active security selection, they frequently use passive instruments and fixed strategies (Jung et al., 2018). The ML model's potential to generate alpha is not used by this passive strategy (D’Acunto et al., 2019).

A clear performance edge is shown by recent studies on Generative AI. Hedge funds that used this Generative AI were analyzed by Sheng et al. (2025). They found that annualized abnormal returns earned by these funds were 3% to 5% higher compared to those that did not use Generative AI. AI processes firm-specific data more effectively than humans, which drove the better performance. Lopez-Lira and Tang (2025) also built a trading strategy. However, theirs was based on LLM news analysis. This strategy generated an annualized Sharpe ratio of 2.97 before transaction costs, which is much higher than traditional market benchmarks. Liu and Shi (2025) also confirm this technical superiority. They show that LLMs reduce the error in estimating volatility by about 80% in comparison to traditional methods. This proves that advanced algorithms understand market risks with better precision.

Also, as discussed before, Gu et al. (2020) demonstrate that a long-short strategy based on neural network forecasts reached a Sharpe ratio of 1.35, significantly exceeding the 0.51 Sharpe ratio for a buy-and-hold investor. They also found that a complex six-factor models explain only 10% to 30% of the variation in returns generated by ML models. These results strongly support the first hypothesis, confirming that the algorithmic services offer a proven technical advantage in areas where human advisors struggle.

4.3 Human Strengths and Algorithmic Challenges

Sections 4.1 and 4.2 show that algorithms outperform in rational decision-making and performance, but this advantage is mainly limited to technical, quantitative tasks like portfolio optimization. Both basic robo-advisors and advanced ML models focus on solving practical portfolio management challenges by using systematic, data-driven methods to improve risk and returns (Rossi & Utkus, 2020).

4.3.1 Human Strengths

The technical approach often comes from the Markowitz (1952) model. However, it misses the wider range of services that traditional wealth management offers. Human advisors add lasting value in two important ways. Algorithms cannot match them in holistic financial planning and psychological support.

Human advisors offer a holistic service. It goes further than portfolio allocation (Lightbourne, 2017). This approach includes complex services. For example, it includes tax planning, retirement structuring, and inheritance advice. A thorough grasp of a person's individual personal and financial situation is required for these services. But for robo advisors achieving this is difficult, because their processes are standardized and algorithm-driven (Jung et al., 2018).

Just as important is the psychological support that algorithms cannot offer. The Money Doctor model helps explain why. It addresses the main question in delegated management. Why do investors pay high fees for human managers who usually underperform the market after fees (Gennaioli et al., 2015)? Investors are often "too nervous or anxious" to make risky investments on their own. The model points this out.

Clients are really looking for trust. They focus on it sometimes more than higher returns. Trust in a human manager helps lower investors sense of risk (Gennaioli et al., 2015). Personal relationships, familiarity, and communication build this trust. The higher fee covers both professional skills and the peace of mind that comes from empathy and trust. These qualities give clients the confidence to handle market risks and avoid panic-selling, which a data-driven algorithm cannot.

4.3.2 Algorithmic Challenges

While algorithmic models can help reduce the cognitive biases shown in Chapter 2, they introduce several technological and ethical issues. One of the issues is algorithmic bias. Algorithms are not completely objective but rather reflect biases that originate from their developers and the data used to train them (Phoon & Koh, 2018). If the training data used is outdated, biased, or has systematic biases, the algorithm will learn and re-make these errors (D'Acunto et al., 2019). This leads to new types of systematic errors in investment decisions. Luqman et al. (2025) also warn that Agentic AI systems operate autonomously. If their training data has historical biases, they can accidentally increase discriminatory practices. This can lead to systematic errors in loan approvals or risk assessments without human oversight.

The efficiency of algorithms creates a black-box effect (Lightbourne, 2017). The study explains that the most effective models are often so complex that their decision-making processes are impossible to fully explain to a human. Similarly, another study provides evidence for this problem, noting that ML models often fail to offer much to readers (Rossi & Utkus, 2020). M. Chen (2025) highlights that this lack of transparency creates a serious governance problem. Asset managers have a legal duty to act in the client's best interest, but still they cannot fully explain why an autonomous AI agent made a specific trade. This lack of transparent reasoning transforms the technical problem into a profound liability issue, which makes it nearly impossible for the robo-advising service to defend the prudence of its investment choices to regulators or clients. This means that the main question is how robo-advisors can meet their legal responsibilities if their algorithms' decisions cannot be explained to investors or regulators.

Trusting the machine creates a new behavioral risk known as the trust trap. When a robo-advisor uses Generative AI to become conversational, it builds affective trust in the user. This human-like trust in the robot creates a dilemma. It makes the investor more prone to influence and makes them more likely to accept objectively incorrect or more expensive portfolio recommendations, even when clear warnings are visible. The study found that this over-reliance could not be reduced by clear disclaimers, suggesting that affective trust makes the investor highly sensitive to influence. (Hildebrand & Bergner, 2021).

4.4 The Hybrid Model

The analysis above shows that the best approach to portfolio management is not found in competition between models, but in their combined strengths. A comparative analysis shows that neither pure human advice nor pure algorithm is perfect on its own. While algorithms are better than most human analysts in forecasting stock returns, humans maintain an advantage when institutional knowledge matters. Therefore, humans and algorithms are perfect complements, unlocking the highest potential. (Cao et al., 2024).

The hybrid model is built on behavioral analysis (Phoon & Koh, 2018). Human analysts often show predictable biases and irrational behavior due to psychological traits (Cao et al., 2024). In comparison, Cao et al. (2024) demonstrate that the algorithm is rational and follows a disciplined plan, using its ability to analyze huge amounts of data while being immune to human biases. Their study found that removing these biases explains about 22% of the performance difference between humans and machines. Phoon and Koh (2018) also found that the algorithm is more cost-efficient than humans. They state that robo-advisors can handle quantitative routine tasks at lower costs than humans and automate the simpler tasks of portfolio management.

Automation is becoming more advanced because of new technologies. M. Chen (2025) argues that the hybrid model is being transformed by Agentic AI. It moves from simple automation to a "human-plus-AI" partnership. Autonomous AI agents do not just follow rules in this setup: they analyze data independently, execute trades, and handle compliance workflows. Because of this the human investor can focus on strategy instead of micromanaging. By delegating execution to autonomous agents, they can train employees on higher-level decision-making (M. Chen, 2025). This also means firms can reorganize their teams.

Performance improves because of this operational advantage. The “Man + Machine” approach makes more precise forecasts than humans or machines working alone (Cao et al., 2024). Reducing forecasting mistakes is the model’s key advantage and it can avoid nearly 90% of the extreme errors that humans would usually make. Sheng et al. (2025) note that human ability to identify sophisticated market signals is improved by adding Generative AI into this workflow. Vast amounts of raw information are processed by Generative AI. For example, it can process earnings calls and news. Human advisors get useful insights from this because finding them manually would be almost impossible.

Tasks are assigned based on strengths in the best hybrid model. Mechanical, data-driven portfolio optimization is handled by the algorithm, while humans provide psychological support and handle holistic goals (Phoon & Koh, 2018). M. Chen (2025) agrees with this. He notes that AI should be in charge of executing the trades. Humans should then define the strategy and handle moral limits and the market patterns also support this (Cao et al., 2024). In the study they say that about 20% of robo-advisors now combine computer recommendations with human touch. A good example of this hybrid model is Vanguard Personal Advisor Services. It combines a personal advisor and a robo-advisor (Phoon & Koh, 2018).

5 Conclusion

This thesis started by identifying a conflict in modern finance. It is the difference between the rational, optimized approach of Modern Portfolio Theory (MPT) and the human investor, who often shows cognitive biases. In this thesis we compared two models. One is the traditional, human-led advisory model. The other is the newer algorithmic model, which ranges from basic robo-advisors to advanced Agentic AI. Our main hypothesis was that algorithms achieve superior cost-efficiency and performance. We also hypothesized that they are more effective towards lessening cognitive biases. The literature review supports both hypotheses. However, it also shows that human advisors still have an important role.

Defining the "human problem" was important for validating the second hypothesis. Chapter 2 showed that investors often act irrationally. This is seen in behavioral biases like overconfidence, the Disposition Effect, and herding. Investment decisions detach from fundamental analysis because of these biases and Chapter 4.1 answered this. It showed that the algorithm functions as a commitment device that is able to automatically counter these biases. Regret aversion or social pressures are not faced by it. Studies also show that using robo-advisors greatly reduces the Disposition Effect. This advantage is strengthened by new technologies like Large Language Models (LLMs) which allow the system to interpret market news and risks objectively. This happens without the emotional reactions that often mislead human investors.

A clear technical and financial advantage for the algorithmic model was found when examining the first hypothesis. The traditional model has a main weakness. It is its high costs, and these high costs remove any possible alpha. In contrast the algorithmic models are more cost-efficient having much lower expenses and portfolio turnover. With Agentic AI this advantage is growing. It lowers operating costs by handling complex financial tasks autonomously. Stronger predictive power was also shown by advanced machine learning (ML) models. They are able to produce a high adjusted alpha, which cannot be explained by traditional factor models. Generative AI supports this by through evaluation of large amounts of unstructured data. Better investment opportunities are found this way and this also confirms that the algorithmic model has a clear edge in terms of cost-efficiency and performance.

This thesis also finds that relying only on quantitative comparison is not enough. The analysis revealed that human advisors offer unique, non-technical value that algorithms cannot match. As stated in the Money Doctor model, human advisors are important for psychological support, empathy, and the trust that gives clients the confidence to take on market risk. They are also better at holistic wealth management when dealing with complex personal situations such as tax planning, inheritance, and retirement structuring that fall outside what algorithms are designed to handle.

Both models being imperfect lead to the conclusion of this thesis. Human advisors can be affected by cognitive bias, and the algorithmic model is limited by its purely technical focus, the black box problem, and its own potential for algorithmic bias. The best approach is not to replace the human advisors, but to combine them with the algorithm, and therefore, the future of portfolio management lies in the Hybrid Model. As explained in Chapter 4.4, this model leverages both. The algorithm and Agentic AI handle mechanical, data-driven optimization and bias correction. At the same time, the human advisor can focus on holistic planning, setting goals, and managing the one thing algorithms cannot: the client's unique psychology.

This study points to two main directions for future research. First, since this thesis is a literature review, more empirical studies are needed to directly compare the long-term performance and cost-efficiency of traditional human-led, advanced algorithmic, and hybrid models in a controlled environment. Second, more work is required to develop algorithms that are more transparent and explainable to address the black-box problem, and the ethical risks of autonomous Agentic AI identified in this thesis.

References

Aziz, S., Dowling, M., Hammami, H., & Piepenbrink, A. (2022). Machine learning in finance: A topic modeling approach. *European Financial Management*, *28*(3), 744–770. https://doi.org/10.1111/eufm.12326

Barber, B. M., & Odean, T. (2001). BOYS WILL BE BOYS: GENDER, OVERCONFIDENCE, AND COMMON STOCK INVESTMENT. *The Quarterly Journal of Economics*, *116*(1), 261–292. https://doi.org/10.1162/003355301556400

Barras, L., Scaillet, O., & Wermers, R. (2010). False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas. *The Journal of Finance*, *65*(1), 179–216.

Cao, S., Jiang, W., Wang, J., & Yang, B. (2024). From Man vs. Machine to Man + Machine: The art and AI of stock analyses. *Journal of Financial Economics*, *160*, 103910. https://doi.org/10.1016/j.jfineco.2024.103910

Chen, M. (2025). Agentic AI and the Future of Institutional Asset Management. *Journal of Portfolio Management*, *51*(10), 144–152. https://doi.org/10.3905/jpm.2025.1.751

Chen, R., & Ren, J. (2022). Do AI-powered mutual funds perform better? *Finance Research Letters*, *47*, 102616. https://doi.org/10.1016/j.frl.2021.102616

D’Acunto, F., Prabhala, N., & Rossi, A. G. (2019). The Promises and Pitfalls of Robo-Advising. *The Review of Financial Studies*, *32*(5), 1983–2020. https://doi.org/10.1093/rfs/hhz014

D'Hondt, C., De Winne, R., Ghysels, E., & Raymond, S. (2020). Artificial Intelligence Alter Egos: Who might benefit from robo-investing? *Journal of Empirical Finance*, *59*, 278–299. https://doi.org/10.1016/j.jempfin.2020.10.002

Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, *25*(2), 383–417. https://doi.org/10.2307/2325486

Gennaioli, N., Shleifer, A., & Vishny, R. (2015). Money Doctors. *The Journal of Finance*, *70*(1), 91–114.

Grinblatt, M., & Keloharju, M. (2009). Sensation Seeking, Overconfidence, and Trading Activity. *The Journal of Finance*, *64*(2), 549–578.

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies*, *33*(5), 2223–2273.

Hambly, B., Renyuan, X., & Huining, Y. (2023). Recent Advances in Reinforcement Learning in Finance. *Mathematical Finance*, *33*(3), 437–503. https://doi.org/10.1111/mafi.12382

Hildebrand, C., & Bergner, A. (2021). Conversational robo advisors as surrogates of trust: Onboarding experience, firm perception, and consumer financial decision making. *Journal of the Academy of Marketing Science*, *49*(4), 659–676. https://doi.org/10.1007/s11747-020-00753-z

Hsu, W.-C., Wang, M.-C., & Ting, H.-I. (2025). Reducing emotional bias in investment decisions: The role of GPT-4 in financial analysis. *Asia-Pacific Journal of Business Administration*, 1–26. https://doi.org/10.1108/APJBA-03-2025-0181

Jung, D., Dorner, V., Glaser, F., & Morana, S. (2018). Robo-Advisory: Digitalization and Automation of Financial Advisory. *Business & Information Systems Engineering*, *60*(1), 81–86. https://doi.org/10.1007/s12599-018-0521-9

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica (Pre-1986)*, *47*(2), 263–291. https://doi.org/10.2307/1914185

Kaplan, A., & Haenlein, M. (2019). Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, *62*(1), 15–25. https://doi.org/10.1016/j.bushor.2018.08.004

Krauss, C., Do, X. A., & Huck, N. (2017). Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. *European Journal of Operational Research*, *259*(2), 689–702. https://doi.org/10.1016/j.ejor.2016.10.031

Lightbourne, J. (2017). ALGORITHMS & FIDUCIARIES: EXISTING AND PROPOSED REGULATORY APPROACHES TO ARTIFICIALLY INTELLIGENT FINANCIAL PLANNERS. *Duke Law Journal*, *67*(3), 651–679.

Liu, T., & Shi, Y. (2025). News sentiment and investment risk management: Innovative evidence from the large language models. *Economics Letters*, *247*, 112124. https://doi.org/10.1016/j.econlet.2024.112124

Lopez-Lira, A., & Tang, Y. (2025). *Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models* (SSRN Scholarly Paper No. 4412788). Social Science Research Network. https://doi.org/10.2139/ssrn.4412788

Luqman, M., Gauba, H., Kumar, P., Prabhu Desai, A., Yadav, A., Prajapati, R., & Yadav, P. (2025, August 3). *Agentic AI in Finance: A Comprehensive Overview*.

https://www.researchgate.net/publication/394500844_Agentic_AI_in_Finance_A_Comprehensive_Overview

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, *7*(1), 77–91. https://doi.org/10.2307/2975974

Odean, T. (1998). Are Investors Reluctant to Realize Their Losses? *The Journal of Finance*, *53*(5), 1775–1798.

Phoon, K., & Koh, F. (2018). Robo-Advisors and Wealth Management. *The Journal of Alternative Investments*, *20*(3), 79–94. https://doi.org/10.3905/jati.2018.20.3.079

Rossi, A. G., & Utkus, S. P. (2020). Who Benefits from Robo-advising? Evidence from Machine Learning. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3552671

Salo, A., Doumpos, M., Liesiö, J., & Zopounidis, C. (2024). Fifty years of portfolio optimization. *European Journal of Operational Research*, *318*(1), 1–18. https://doi.org/10.1016/j.ejor.2023.12.031

Scharfstein, D. S., & Stein, J. C. (1990). Herd behavior and investment. *American Economic Review*, *80*(3), 465–479.

Shefrin, H., & Statman, M. (1985). The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence. *The Journal of Finance*, *40*(3), 777–790. https://doi.org/10.2307/2327802

Sheng, J., Sun, Z., Yang, B., & Zhang, A. L. (2025). *Generative AI and Asset Management* (SSRN Scholarly Paper No. 4786575). Social Science Research Network. https://papers.ssrn.com/abstract=4786575

Shiller, R. J. (1995). Conversation, information, and herd behavior. *The American Economic Review*, *85*(2), 181–185.

Thaler, R. (1985). MENTAL ACCOUNTING AND CONSUMER CHOICE. *Marketing Science (Pre-1986)*, *4*(3), 199.