WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

BACHELOR'S THESIS

FINWIKI

a financial search engine

submitted by

Raphael Keisler

intended degree

Bachelor of Science (WU), BSc (WU)

Student ID number:	**12309392**
Degree programme:	Business, Economics and Social Science
Supervisor:	Univ.Prof. Mag.Dr.rer.soc.oec. Rony Flatscher

Vienna, March 2026

WIRTSCHAFTS UNIVERSITÄT WIEN VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS

Declaration of Originality

University of Economics and Business

I hereby declare that:

1. I have written this Bachelor's Thesis myself, independently and without the aid of unfair or unauthorized resources. Whenever content has been taken directly or indirectly from other sources, this has been indicated and the source referenced.

2. This Bachelor's Thesis has not been previously presented as an examination paper in this or any other form in Austria or abroad.

3. This Bachelor's Thesis is identical with the thesis assessed by the examiner.

.

Vienna, Jan. 2026



Signature of Author

Usage of Artificial Intelligence

During the preparation of this bachelor's thesis, AI-based language models were used in a supportive and auxiliary manner. Their primary purpose was to assist in the rephrasing, restructuring, and linguistic refinement of text passages in order to improve clarity, coherence, and academic style. In particular, AI tools were used to enhance readability, adjust phrasing to an appropriate scientific tone, and reduce redundancy, without introducing new content or altering the intended meaning of the text.

All conceptual development, academic reasoning, formulation of research questions, selection and interpretation of sources, and critical evaluation of results were carried out independently by the author. AI-based tools were not used to generate original scientific arguments, derive conclusions, or make methodological decisions. All content produced with the assistance of AI was carefully reviewed, validated, and, where necessary, manually revised by the author to ensure correctness and alignment with academic standards.

Furthermore, the practical component of this thesis explicitly involves working with modern AI technologies, particularly Large Language Models (LLMs). As part of this work, AI-based systems were used during the programming process, for example to support code generation, explain technical concepts, and assist in optimization tasks. In this context, AI tools were treated as development aids rather than autonomous decision-makers.

Consequently, selected functional and content-related elements were generated or adapted with the assistance of AI-based tools. Full responsibility for the selection, application, validation, and evaluation of all AI-generated outputs remains with the author. The use of AI in both the writing and implementation process was conducted transparently and in accordance with principles of academic integrity.

Acknowledgment

A big thank you goes to my supervisor for the trust placed in me and for granting a high degree of academic freedom throughout this thesis. The opportunity to work independently and explore my own research ideas was highly valuable.

I would also like to acknowledge WU Vienna University of Economics and Business for providing an excellent learning environment. The bachelor's program offered a strong foundation and made this thesis a rewarding and enriching learning experience.

Abstract

This thesis examines the applicability of artificial intelligence, particularly Large Language Models (LLMs), in the context of financial analysis and investment decision support. A prototype system was developed that combines financial market data, news analysis and AI-based interpretation to support users in evaluating financial assets.

The results show that AI can effectively assist in processing and summarizing large volumes of financial information. However, the evaluation also reveals significant limitations, including inconsistent outputs, limited reproducibility and challenges in generating structured and reliable results. These factors prevent the use of AI as a fully autonomous financial advisor.

Overall, the findings indicate that while AI is well suited as a decision-support tool, human oversight remains essential for responsible financial decision-making. The thesis concludes by outlining potential directions for future research, including the use of domain-specific models and hybrid decision-support systems.


finwiki

Table of Content

1 INTRODUCTION 1

1.1 Problem, Statement and Motivation 1
1.2 Objectives of the Thesis 1
1.3 Research Question and Methodology 2

2 FINANCIAL DECISION-MAKING IN DIGITAL AGE 3

2.1 Structural Changes in Modern Financial Markets 3
2.2 Challenges Faced by Private Investors 3
2.2.1 Limited Financial Literacy 3
2.2.2 Time Constraints and Cognitive Overload 3
2.3 Limitations of Traditional Financial Advisory Models 4
2.3.1 Cost and Accessibility 4
2.3.2 Standardization and Lack of Individualization 4
2.3.3 Limited Adaptability and Transparency 4
2.4 Digital Transformation of Financial Services 4
2.4.1 Rise of Online Platforms, Robo-Advisors and Trading Apps 4
2.4.2 Increasing Importance of Data-Driven Decision Making 5
2.4.3 Role of Financial News and Market Sentiment 5
2.5 Role of Artificial Intelligence in Financial Analysis 6
2.5.1 Analytical Capabilities of Large Language Models 6
2.5.2 Opportunities and Limitations 6
2.6 Motivation for a Hybrid Decision-Support Approach 6

3 RESEARCH METHODOLOGY AND EMPIRICAL BASIS 7

3.1 Purpose of the Empirical Study 7
3.2 Survey Design and Scope 7
3.3 Sample Characteristics 7
3.4 Investment Experience and Current Behaviour 8
3.5 Risk Perception and Financial Constraints 8
3.6 Expectations toward Portfolio Recommendations 9
3.7 Transparency, Trust and User Control 10
3.8 Implications for System Design 10
3.9 Survey Summary 10

4 ARTIFICIAL INTELLIGENCE AND LARGE LANGUAGE MODELS 11

4.1 Artificial Intelligence as Technical Paradigm 11
4.2 Machine Learning as the Basis of Modern AI Systems 11
4.2.1 Learning Paradigms and Their Characteristics 11
4.2.2 Role of Data in Financial Machine Learning 12
4.3 Natural Language Processing in Finance 12
4.3.1 Importance of Textual Information 12
4.4 Large Language Models: Architecture and Functionality 12
4.4.1 Transformer-Based Architecture 12
4.4.2 Training and Knowledge Representation 13
4.5 Applications of LLMs in Financial Contexts 13
4.5.1 Opportunities and Use Cases 13
4.5.2 Challenges and Risks 13
4.6 Implications for System Design 14

5 CONCEPTUAL DESIGN AND SYSTEM ARCHITECTURE 15

5.1 Design Objectives and System Requirements 15
5.1.1 Functional Requirements 15
5.1.2 Non-Functional Requirements 15
5.2 High-Level Architectural Overview 16

5.3 Frontend Development and UI Design 16
5.3.1 Design Philosophy 16
5.3.2 User Interaction Model 16
5.3.3 Dashboard-Based Interface Design 17
5.4 Backend Architecture and Application Logic 17
5.4.1 Role of the Backend 17
5.4.2 Internal API Design 17
5.5 Data Acquisition and Processing 18
5.5.1 Financial Market Data 18
5.5.2 Asset Identification and Normalization 18
5.5.3 News Acquisition and Preprocessing 18
5.6 Integration of Artificial Intelligence 18
5.6.1 Initial Role of AI in the System 18
5.6.2 Prompt Design and Control Mechanisms 19
5.7 Scoring and Evaluation Model 19
5.8 Design Constraints and Architectural Trade-Offs 19

6 TECHNICAL IMPLEMENTATION 20

6.1 Purpose and Scope of the Implementation 20
6.2 Development Environment and Technology Stack 20
6.2.1 Programming Languages and Tools 20
6.2.2 System Execution Model 21
6.3 Implementation of Asset Search and Identification 21
6.3.1 Asset Reference Dataset 21
6.3.2 Symbol-ISIN Mapping Strategy 21
6.4 Financial Data Retrieval and Processing 22
6.4.1 Market Data Acquisition 22
6.4.2 Data Normalization and Validation 22
6.5 News Retrieval and Processing Pipeline 22
6.5.1 News Acquisition 22
6.5.2 Text Preprocessing 23
6.5.3 Initial Approach: AI-Based News Interpretation 23
6.5.4 Actual Approach: Deterministic Multi-Dimensional News Score 23
6.6 Scoring Architecture and Integration of AI 24
6.6.1 Intended AI-Supported Scoring Framework 24
6.6.2 Transition to Deterministic Multi-Dimensional Scoring 24
6.7 Quantitative Evaluation and Final Score Aggregation 25
6.7.1 Quantitative Scoring Component 25
6.7.2 Final Score Aggregation 25
6.8 Performance and Architectural Implications 25

7 SYSTEM DEMONSTRATION AND UI OVERVIEW 27

7.1 Purpose of the Demonstration 27
7.2 Implemented Financial Search Engine 27
7.2.1 Landing Page 27
7.2.2 Search Function 28
7.2.3 Loading State Visualisation 28
7.2.4 Asset Overview Page - General Information 29
7.2.5 Asset Overview Page - Assessment 30
7.2.6 Asset Overview Page - News 31
7.3 Implemented Components of Advisory Dashboard 31
7.3.1 Planned Sections 31
7.3.2 Survey and Question Types 31
7.3.3 Calculation Waiting Screen 32
7.3.4 Dashboard - Overview Page 33
7.3.5 Dashboard - Risk Assessment Page 34
7.4 User Interaction Flow 34

8 EVALUATION AND ANALYSIS OF RESULTS 35

8.1 Evaluation Objectives and Methodological Approach 35
8.2 Evaluation of the Financial Search and Analysis Engine 35
8.2.1 Asset Identification and Data Integrity 35
8.2.2 Market Data Retrieval and Preprocessing 36
8.2.3 Assessment of LLM-Based News Interpretation 36
8.2.4 Evaluation of the Final Scoring Architecture 36
8.3 Evaluation of the Advisory-Oriented Implementation 37
8.3.1 Output-Consistency and Determinism 37
8.3.2 Reliability of Structured Output Generation 37
8.3.3 Implications for Advisory Automation 38
8.4 Performance Characteristics and Practical Feasibility 38
8.5 Summary of Evaluation Results 38

9 DISCUSSION AND CRITICAL REFLECTION 39

9.1 Aim and Position of Discussion 39
9.2 Interpretation of the Results 39
9.3 Implications for AI-Based Financial-Systems 39
9.4 Reflection on the Abandoned Advisory Component 40
9.5 Scientific and Practical Contribution 40
9.6 Summary of Discussion 41

10 CONCLUSION AND OUTLOOK 42

10.1 Summary of the Thesis 42
10.2 Answering the Research Question 42
10.3 Key Findings and Contributions 43
10.4 Limitations of the Work 43
10.5 Outlook and Future Research Directions 44
10.6 Final Remarks 44

LIST OF DIAGRAMS 45

List of Tables 45
List of Figures 45

REFERENCES 46

1 Introduction

This chapter introduces the motivation and objectives of the thesis. It outlines the central problem, formulates the research question and briefly presents the methodological approach.

1.1 Problem, Statement and Motivation

Managing and growing personal wealth has become increasingly challenging in recent years. While previous generations often relied on traditional savings accounts or personal financial advisors, today's financial environment is characterized by low interest rates, rising inflation and a rapidly expanding range of investment opportunities. As a result, making well-informed financial decisions has become more complex than ever before. (Jung et al., 2018).

At the same time, many private individuals lack the necessary financial knowledge or the time required to analyze financial markets in depth. This often leads to suboptimal financial behaviour, such as leaving capital unused in low-interest accounts or making investment decisions based on incomplete or unreliable information. Although professional financial advisors can partially address this issue, their services are often costly, not fully transparent or based on standardized models that do not sufficiently account for individual circumstances.

In addition, recent advances in artificial intelligence have significantly expanded the possibilities of automated data analysis and decision support. In particular, Large Language Models (LLMs) have demonstrated remarkable capabilities in processing large volumes of information, identifying patterns and generating structured outputs. These developments raise the question of whether such technologies can be applied to the domain of personal finance in order to support users in making better-informed investment decisions. (Bommasani et al., 2022).

This thesis is motivated by the idea that combining financial data analysis with modern AI techniques could provide a more accessible, efficient and scalable alternative to traditional financial advisory models.

1.2 Objectives of the Thesis

The primary objective of this bachelor's thesis is to explore the application of artificial intelligence in the context of financial decision support through the conceptualization and prototypical implementation of two complementary system components.

The first component focuses on the development of an AI-supported financial search and analysis engine that aggregates, processes and evaluates information about financial assets in a structured and transparent manner.

The second objective is to investigate the feasibility of an AI-supported advisory approach aimed at generating personalized financial insights based on user-specific information. In this context, the thesis examines the suitability of Large Language Models for analyzing both quantitative market data and qualitative information, such as financial news, while critically assessing their limitations when applied to structured and high-stakes decision-making scenarios.

Rather than developing a market-ready product, this work follows an exploratory and analytical approach. Emphasis is placed on evaluating technical feasibility, interpretability and risk, with the goal of identifying realistic application boundaries for AI-supported systems in financial contexts.

1.3 Research Question and Methodology

Based on the outlined problem statement, the central research question of this thesis is:

To what extent are modern AI models suitable for supporting or automating financial analysis and investment-related decision-making and what technical and conceptual limitations arise in this context?

To address this question, a prototype system was developed consisting of the two already mentioned main components:

1. A financial search engine for collecting and analyzing asset-related data and news content.
2. A conceptual framework for an AI-driven digital investment advisor visualized in a dashboard.

The methodology follows a practice-oriented approach. After establishing the theoretical foundation, the system is implemented step by step and subsequently evaluated based on its performance, reliability and practical limitations. Particular attention is paid to challenges encountered during development and to the reasons that ultimately led to a reduction in the system's functional scope.

2 Financial Decision-Making in Digital Age

This chapter examines the structural developments and challenges influencing modern financial decision-making. It outlines changes in financial markets, investor behaviour and technological transformation, providing the foundation for understanding the relevance of digital and AI-supported financial systems.

2.1 Structural Changes in Modern Financial Markets

Over the past decades, financial markets have undergone profound structural changes driven by globalization, digitalization and technological innovation. Financial products that were once accessible only to institutional investors are now widely available to private individuals. Online brokerage platforms, real-time market data and low transaction costs have significantly lowered the barriers to market participation.

At the same time, this increased accessibility has also led to greater complexity. Investors are no longer limited to a small number of savings products but can choose from thousands of stocks, exchange-traded funds (ETFs), bonds, derivatives and cryptocurrencies. Each of these asset classes follows different market mechanisms and risk-return profiles, which makes informed decision-making increasingly difficult for non-professional investors.

In addition, financial markets have become more volatile due to geopolitical uncertainties, macroeconomic developments and rapid information dissemination through digital media. As a result, investment decisions today require not only financial knowledge but also the ability to process and interpret large volumes of dynamic information in real time.

2.2 Challenges Faced by Private Investors

Private investors operate in an environment characterized by uncertainty, complexity and increasing informational pressure. Although access to financial markets has improved significantly, the ability to navigate them effectively remains unevenly distributed. Behavioural biases, knowledge gaps and cognitive constraints shape investment decisions and often result in suboptimal financial outcomes.

2.2.1 Limited Financial Literacy

One of the central challenges in personal finance is the limited financial literacy among large parts of the population. Many individuals lack a fundamental understanding of concepts such as diversification, risk-adjusted return, inflation or compound interest. This knowledge gap often leads to suboptimal financial behaviour, such as excessive risk aversion, overconfidence or reliance on anecdotal information.

Furthermore, financial education is rarely a formal part of school curricula, which means that many people are forced to acquire financial knowledge independently. This process is often fragmented and influenced by non-scientific sources such as social media, online forums or anecdotal success stories.

2.2.2 Time Constraints and Cognitive Overload

In addition to knowledge gaps, time constraints play a significant role. Proper financial planning requires continuous monitoring of markets, evaluation of new information and regular portfolio adjustments. For individuals with full-time employment or other responsibilities, dedicating sufficient time to financial analysis is often unrealistic.

The sheer volume of available information further exacerbates this issue. News portals, analyst reports, financial indicators and social media generate an overwhelming amount of data, making it increasingly difficult to distinguish relevant signals from noise. This phenomenon, often referred to as information overload, can lead to decision paralysis or poorly informed choices.

2.3 Limitations of Traditional Financial Advisory Models

Conventional financial advisory services were originally designed to provide structured guidance and professional expertise. However, structural characteristics of these models increasingly reveal practical limitations in accessibility, flexibility and transparency. As financial markets evolve and user expectations shift, traditional advisory approaches face growing pressure to adapt.

2.3.1 Cost and Accessibility

Traditional financial advisory services are often associated with high costs, making them less accessible to individuals with lower investment volumes. In many cases, professional advice is only economically viable for clients with significant capital, which excludes a large portion of the population from personalized financial guidance. (D'Acunto et al., 2019).

2.3.2 Standardization and Lack of Individualization

Many advisory services rely on standardized risk profiles and predefined portfolio models. While this approach allows for scalability, it often fails to account for individual circumstances, personal goals or changing life situations. As a result, recommendations may be overly generic and insufficiently tailored to the user's needs. (Jung et al., 2018).

2.3.3 Limited Adaptability and Transparency

Traditional advisory processes typically operate on periodic reviews rather than continuous analysis. This makes it difficult to react quickly to market developments or newly available information. Moreover, clients often have limited insight into how recommendations are derived, which reduces transparency and trust.

2.4 Digital Transformation of Financial Services

Technological innovation has fundamentally reshaped how individuals interact with financial markets. Digital platforms, algorithmic systems and mobile applications have transformed investment from a specialized activity into an accessible, everyday process. This transformation has expanded opportunities while simultaneously introducing new behavioural and structural challenges.

2.4.1 Rise of Online Platforms, Robo-Advisors and Trading Apps

The digital transformation of the financial sector has led to the widespread adoption of online brokers, fintech platforms and robo-advisors. These systems aim to lower entry barriers by automating parts of the investment process and offering user-friendly access to financial markets. Portfolio construction and rebalancing are often based on predefined risk profiles and algorithmic rules, allowing for cost-efficient and scalable solutions. (Jung et al., 2018).

Robo-advisors and digital trading platforms provide several advantages, including reduced fees, high accessibility and intuitive user interfaces. As a result, they have gained significant

popularity, particularly among younger investors and individuals with limited prior financial experience.

However, many of these platforms rely on business models that generate revenue through transaction-based fees, spreads or payment-for-order-flow mechanisms. This can create indirect incentives to promote frequent trading activity rather than long-term investment strategies. User interfaces are often optimized for simplicity and engagement, which may encourage rapid decision-making and repeated transactions. (D'Acunto et al., 2019).

In addition, the decision-making logic of many digital platforms remains largely rule-based and focuses primarily on structured quantitative data. The ability to interpret unstructured information, such as financial news, macroeconomic developments or market sentiment, is often limited. As a consequence, automated recommendations or trading suggestions may overlook important contextual factors, potentially leading to suboptimal decisions.

2.4.2 Increasing Importance of Data-Driven Decision Making

Modern financial decision-making is increasingly driven by large volumes of data. Market prices, corporate reports, economic indicators and real-time news all contribute to shaping investment strategies. While access to data has improved significantly, the primary challenge lies in extracting meaningful insights from heterogeneous and rapidly changing information sources.

Although many digital platforms provide extensive data and analytical tools, the interpretation of this information is often left to the user. In addition, the presentation of data is frequently optimized for usability and engagement rather than for critical reflection. This can make it difficult for users to assess information objectively, particularly when platforms are not fully neutral in their incentives.

These developments have led to growing interest in advanced analytical approaches, especially those based on artificial intelligence and machine learning. Such methods promise to support users by aggregating, contextualizing and interpreting large datasets more systematically, potentially reducing behavioural biases and information overload.

2.4.3 Role of Financial News and Market Sentiment

Financial markets are influenced not only by numerical indicators but also by information flow and collective expectations. News articles, corporate announcements and macroeconomic reports shape how both private and professional investors perceive assets and assess future developments. These perceptions are commonly reflected in market sentiment, which plays a significant role in investment behaviour. (Tetlock, 2007).

Market sentiment is largely formed through the continuous consumption and interpretation of financial news. Positive or negative narratives can influence investor decisions before their effects are fully reflected in market prices. As a result, relying exclusively on historical prices or quantitative indicators provides an incomplete picture of market dynamics, particularly in situations driven by expectations or external events. (Tetlock, 2007).

Incorporating news analysis into financial systems therefore enables the contextualization of market data by capturing prevailing sentiment and expectations. By analysing multiple articles related to a specific asset, it becomes possible to approximate the general market mood and to better understand how other market participants may be positioning themselves.

Additionally, the inclusion of up-to-date news helps reduce the risk of outdated or misleading guidance. Financial markets can react rapidly to new information, such as product announcements or strategic decisions. Integrating real-time news analysis allows analytical

systems to reflect such developments more promptly and to support more informed, context-aware financial decision-making.

2.5 Role of Artificial Intelligence in Financial Analysis

Advances in artificial intelligence have opened new possibilities for processing complex financial information. In environments characterized by large volumes of structured and unstructured data, AI-based systems offer the potential to enhance analytical depth, reduce cognitive burden and support more informed decision-making.

2.5.1 Analytical Capabilities of Large Language Models

Large Language Models (LLMs) have demonstrated remarkable capabilities in processing unstructured textual data. They can summarize documents, extract key information, recognize patterns and generate human-like text. These properties make them particularly suitable for analyzing financial news, earnings reports and market commentary. (Bommasani et al., 2022).

Unlike traditional algorithms, LLMs can interpret contextual meaning and combine information from multiple sources. This enables a more holistic view of financial events compared to purely quantitative models.

2.5.2 Opportunities and Limitations

While LLMs offer significant potential, their application in finance is associated with critical limitations. Their outputs are probabilistic and not deterministic, which can lead to inconsistencies. Furthermore, LLMs may generate plausible but incorrect information, a phenomenon commonly referred to as hallucination. (Ji et al., 2023).

In financial contexts, where incorrect information can lead to substantial monetary losses, such limitations are particularly problematic. Additionally, LLMs lack inherent domain awareness and must rely on prompt design and external validation to produce reliable results.

2.6 Motivation for a Hybrid Decision-Support Approach

Given the strengths and weaknesses of existing systems, this thesis adopts a hybrid approach that combines algorithmic data processing with AI-based analysis while maintaining human oversight.

The goal is not to replace human decision-making but to enhance it by:

- aggregating and structuring financial data,
- analyzing news and market sentiment,
- providing transparent and interpretable indicators,
- supporting, rather than automating, investment decisions.

This approach acknowledges both the potential and the limitations of artificial intelligence and forms the conceptual foundation for the system developed in this thesis.

3 Research Methodology and Empirical Basis

This chapter presents the empirical study conducted to complement the theoretical and technical components of this thesis. It outlines the study design, describes the data collection process and discusses how the results contributed to the conceptual development of the system.

3.1 Purpose of the Empirical Study

To complement the theoretical and technical components of this thesis, an exploratory empirical study was conducted at the beginning of the project. The purpose of this study was to gain insight into financial behaviour, investment readiness and perceived user needs regarding personal investing and financial decision support. The survey was designed as an exploratory research instrument to identify user perceptions, expectations and challenges related to digital investment tools, providing indicative insights rather than statistically generalizable results.

Rather than validating a finished system, the survey served as an input-oriented research instrument. Its results were used to identify typical user characteristics, common pain points and expectations toward digital investment tools, directly informing the conceptual design of the system and the structure of the advisory questionnaire developed later in the project.

3.2 Survey Design and Scope

The survey was conducted as an **anonymous online questionnaire** and distributed digitally. It was intentionally designed to cover a broad range of topics related to personal finance, including financial situation, investment behaviour, risk tolerance and expectations toward automated portfolio recommendations.

The questionnaire was conducted in German, as the primary target group consisted of German-speaking participants. All results were subsequently analyzed and interpreted for inclusion in this thesis, which is written in English.

The introductory text of the survey emphasized anonymity and transparency and informed participants that the results would be used to design a program capable of generating personalized portfolio recommendations. This framing helped position the survey as a research-oriented initiative rather than a commercial product

3.3 Sample Characteristics

In total, **60 participants** completed the survey. The sample predominantly consisted of young adults, with a strong representation of students and early-career professionals. Participants covered a broad range of employment situations, including full-time employment, self-employment, and mixed statuses. Income levels and financial obligations varied considerably across the sample, indicating heterogeneous financial circumstances rather than a uniform target group.

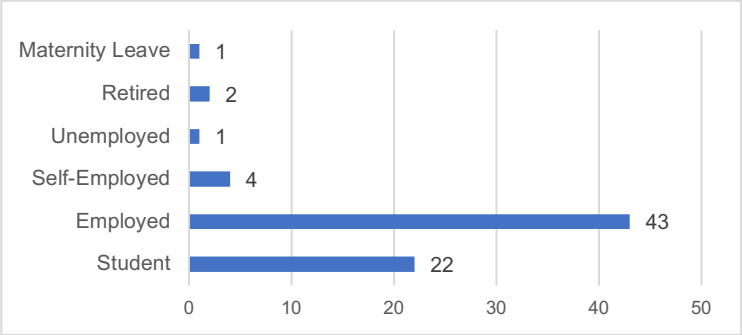

Diagram 1: Distribution of Employment Status (Multiselection possible)

3.4 Investment Experience and Current Behaviour

A key focus of the survey was to assess participants existing investment experience. The results indicate that while a substantial portion of respondents had already engaged in some form of investing, a notable share either had very limited experience or had not invested at all.

Among those who were already investing, commonly used instruments included stocks, ETFs and cryptocurrencies. However, self-reported confidence levels varied widely. Many participants indicated uncertainty regarding asset selection, portfolio composition and long-term strategy.

These findings suggest that access to investment platforms alone does not equate to financial confidence or understanding. Even users who actively invest often lack structured guidance and rely on fragmented information sources.

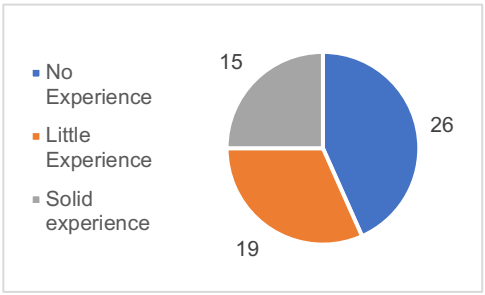

Diagram 2: Experience with Investing

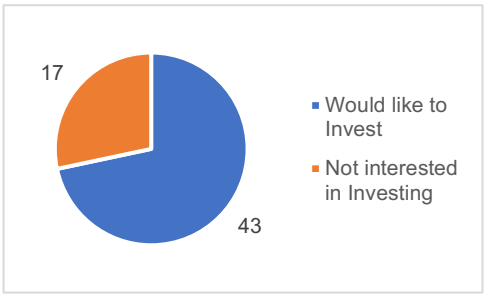

Diagram 3: Interest in Personal Investing

3.5 Risk Perception and Financial Constraints

The survey further examined participants risk tolerance and financial constraints. Results show that most respondents classified themselves as having a **medium risk tolerance**, with fewer participants indicating either very low or very high risk preferences.

At the same time, many participants reported financial limitations, such as limited disposable income, existing obligations or uncertainty regarding future income stability. This combination highlights a central tension: users are interested in investing but are simultaneously cautious and risk-aware.

These insights underline the importance of **risk-aware and context-sensitive financial guidance**, particularly for individuals who cannot afford significant losses.

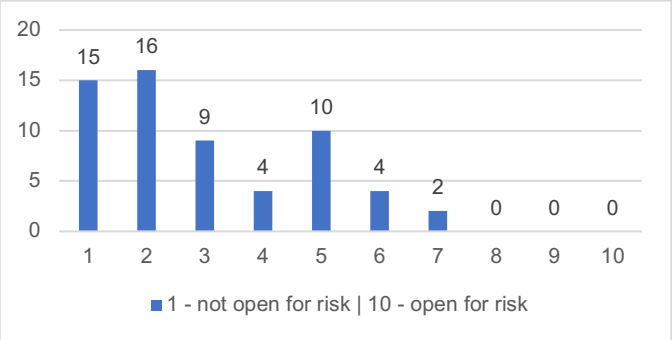

Diagram 4: Risk Awareness (1-10)

3.6 Expectations toward Portfolio Recommendations

Participants were explicitly asked what they would expect from a personalized portfolio recommendation system. Qualitative responses indicate a strong demand for transparency, explanation and adaptability.

Common themes included:

- understandable reasoning behind recommendations,
- consideration of individual financial situations,
- explanation of risks and assumptions,
- avoidance of overly aggressive strategies.

Interestingly, several participants emphasized that they did not want a system that merely "tells them what to buy," but rather one that helps them **understand why certain decisions might be reasonable**.
This feedback strongly supports the decision to design the system as a **decision-support tool rather than an automated advisor.**

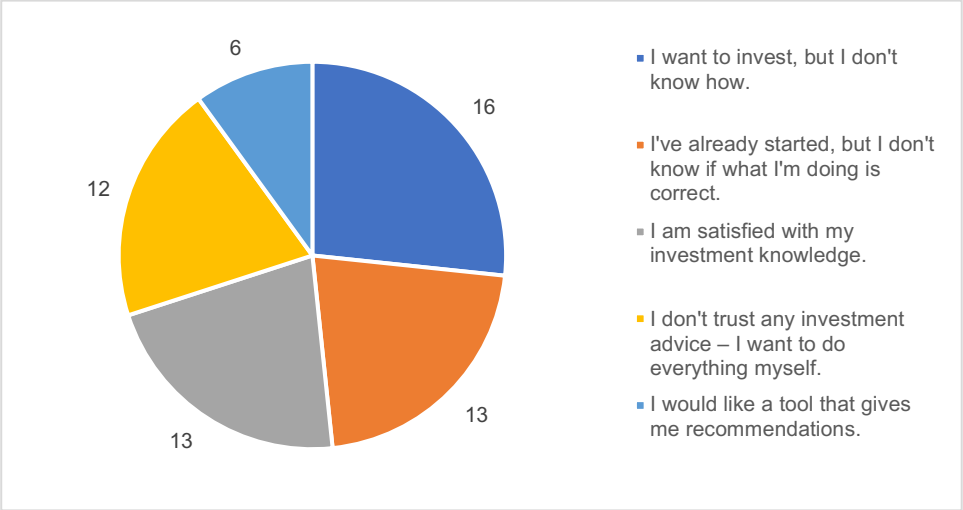

Diagram 5: Personal Assessment

3.7 Transparency, Trust and User Control

A central outcome of the survey is the strong demand for **transparent and comprehensible** investment recommendations. Participants consistently expressed that they want to understand why a specific asset is recommended and which factors influence a particular decision. Rather than receiving prescriptive instructions, users emphasized the importance of explanations that allow them to form their own judgement based on clearly presented information.

Closely related to this expectation is the desire to be informed before any financial commitment is made. The survey results indicate a clear reluctance toward investing in assets that are perceived as unfamiliar or insufficiently explained. This highlights the importance of informational clarity and gradual engagement, enabling users to maintain a sense of **self-control** and avoid investing in the unknown. Participants therefore value systems that support informed decision-making without removing personal responsibility.

These findings also point to a broader issue of trust in existing digital investment platforms. Many users perceive **current recommendation systems as non-transparent**, with decision logic that resembles a “black box.” When users cannot trace how recommendations are generated, confidence in such systems is reduced. The survey therefore suggests that transparency, explainability and user control are essential prerequisites for the acceptance of AI-supported financial systems.

3.8 Implications for System Design

The findings of the survey had direct implications for the subsequent system design. In particular, they informed:

- the structure of the advisory questionnaire,
- the emphasis on transparency and explanation,
- the integration of risk-aware evaluation,
- the avoidance of fully automated decision-making.

The survey results confirmed that the intended target group values guidance and structure but remains sceptical of transparency or overly aggressive financial automation. These insights provided a strong empirical foundation for the conceptual design decisions presented in the following chapter.

3.9 Survey Summary

In summary, the empirical study provided valuable insight into user behaviour, financial constraints and expectations regarding digital investment tools. The results demonstrate a clear demand for systems that reduce complexity, support informed decision-making and respect individual risk profiles.

Rather than serving as validation of a finished product, the survey functioned as an empirical anchor for the design and scope of the proposed system. This approach ensures that the technical and conceptual components of the thesis are grounded in observed user needs rather than purely theoretical assumptions.

4 Artificial Intelligence and Large Language Models

This chapter introduces the theoretical foundations of artificial intelligence and large language models. It explains the underlying technological principles, relevant machine learning approaches and their potential applications and limitations in financial contexts.

4.1 Artificial Intelligence as Technical Paradigm

Artificial Intelligence (AI) represents a broad field of computer science concerned with the development of systems capable of performing tasks that typically require human cognitive abilities. These tasks include learning from experience, recognizing patterns, understanding natural language and making decisions under uncertainty. The concept of AI has evolved significantly since its early theoretical foundations in the mid-20th century, shifting from symbolic rule-based systems to data-driven learning approaches.

Modern AI systems are largely based on statistical methods and machine learning algorithms that extract patterns from large datasets. Rather than being explicitly programmed with fixed rules, these systems adapt their behaviour through training processes that optimize performance based on predefined objectives. This paradigm shift has enabled AI to be applied successfully in complex and data-intensive domains such as image recognition, speech processing, medical diagnostics and financial analysis.

The increasing relevance of AI in economic and financial contexts is driven by two main factors: **the exponential growth of available data** and **the rising computational power that enables the processing of this data in real time**. Together, these developments form the technological foundation for advanced decision-support systems.

4.2 Machine Learning as the Basis of Modern AI Systems

Modern artificial intelligence is fundamentally driven by machine learning, which enables systems to derive patterns and improve performance through data rather than predefined rules. This data-centric approach forms the technological foundation of most contemporary AI applications, including those in finance.

4.2.1 Learning Paradigms and Their Characteristics

Machine learning can be categorized into three primary paradigms, each of which plays a distinct role in modern AI applications:

- **Supervised learning** involves training a model using labeled data, where both input and desired output are known. This approach is commonly used in prediction and classification tasks, such as credit risk assessment or price forecasting. (Fischer & Krauss, 2017)

- **Unsupervised learning** focuses on identifying patterns or structures in unlabeled data. Typical applications include clustering, anomaly detection and dimensionality reduction, which are particularly useful for exploratory financial analysis.

- **Reinforcement learning** is based on interaction with an environment, where an agent learns to optimize its behaviour through trial and error. While powerful, this approach is less common in financial advisory systems due to its complexity and high data requirements.

In financial applications, supervised and unsupervised learning dominate due to their interpretability and relatively stable behaviour.

4.2.2 Role of Data in Financial Machine Learning

Financial machine learning systems rely heavily on large volumes of structured and unstructured data. Structured data includes numerical information such as prices, trading volumes, financial ratios and balance sheet figures. Unstructured data, on the other hand, includes textual information such as news articles, earnings reports, social media content and analyst commentary. (Heaton et al., 2016).

The ability to integrate both data types is essential for comprehensive financial analysis. While structured data provides quantitative insights, unstructured data often contains contextual information that can influence market behaviour. This duality makes financial analysis a particularly suitable domain for advanced AI techniques.

4.3 Natural Language Processing in Finance

A large share of financially relevant information is communicated through text, such as news articles, reports and corporate disclosures. Processing this unstructured information requires specialized methods, making Natural Language Processing a key component of modern financial analysis.

4.3.1 Importance of Textual Information

A significant portion of financially relevant information is published in textual form. Corporate announcements, macroeconomic news, regulatory statements and market analyses all influence investor sentiment and asset prices. Traditional quantitative models are limited in their ability to incorporate such information effectively. (Kearney & Liu, 2013).

Natural Language Processing (NLP) addresses this limitation by enabling machines to analyze and interpret human language. NLP techniques allow systems to extract meaning, identify sentiment and detect relationships within textual data, making them highly valuable for financial applications.

4.4 Large Language Models: Architecture and Functionality

Large Language Models represent a major advancement in AI-based text processing. Their architecture allows them to capture contextual relationships within language and generate coherent outputs, making them particularly relevant for complex analytical tasks.

4.4.1 Transformer-Based Architecture

Large Language Models are typically built on transformer architectures, which consist of multiple layers of self-attention mechanisms and feed-forward neural networks. Unlike recurrent models, transformers process entire sequences simultaneously, allowing them to capture long-range dependencies efficiently.

The core idea behind transformers is the attention mechanism, which assigns different weights to parts of the input text depending on their relevance. This enables the model to focus on meaningful relationships and contextual cues within the data.

4.4.2 Training and Knowledge Representation

LLMs are trained using unsupervised or self-supervised learning objectives, most commonly next-token prediction. During training, the model learns statistical relationships between words and phrases by processing billions of text samples.

It is important to note that **LLMs do not possess factual understanding** or reasoning in a human sense. Instead, they generate outputs based on learned probability distributions. Their apparent intelligence emerges from scale rather than comprehension.

4.5 Applications of LLMs in Financial Contexts

The use of Large Language Models in finance enables automated interpretation of textual information and supports data-intensive analytical workflows. At the same time, their probabilistic nature introduces challenges that must be carefully considered in high-stakes environments.

4.5.1 Opportunities and Use Cases

Within the financial domain, LLMs offer a range of potential applications, including:

- automated analysis of financial news,
- sentiment classification of market-related texts,
- summarization of corporate disclosures,
- extraction of key financial indicators,
- support for investment research and reporting.

These capabilities make LLMs particularly attractive for systems that aim to reduce information overload and enhance decision transparency.

4.5.2 Challenges and Risks

Despite their potential, LLMs present significant challenges when applied to financial decision-making:

- **Lack of Determinism**
  LLMs generate probabilistic outputs, meaning identical inputs may produce different results. This lack of determinism conflicts with the need for reproducibility in financial analysis.

- **Hallucination Risk**
  LLMs may generate plausible but factually incorrect statements. In finance, such errors can lead to misleading conclusions and financial losses.

- **Limited Explainability**
  The internal reasoning processes of LLMs are largely non-transparent, making it difficult to explain or justify specific outputs - a critical requirement in regulated financial environments.

- **Formatting and Validation Issues**

While LLMs can generate structured outputs, ensuring strict compliance with predefined formats remains challenging. This poses a significant obstacle for automated processing pipelines.

4.6 Implications for System Design

The characteristics of LLMs necessitate careful system design choices.

In particular:

- LLMs should not be used as sole decision-makers.
- Outputs must be validated and contextualized.
- Human oversight remains essential.
- AI should function as an analytical assistant rather than an autonomous agent.

These considerations directly influenced the architectural and conceptual decisions made in this thesis. Instead of pursuing full automation, the system was designed as a decision-support tool that combines data-driven analysis with AI-based interpretation.

5 Conceptual Design and System Architecture

The following chapter outlines the conceptual framework and architectural structure of the developed system. It introduces the guiding design principles, system requirements and technical building blocks that form the foundation for the subsequent implementation.

5.1 Design Objectives and System Requirements

The conceptual design of the developed system is guided by the overarching objective of investigating how artificial intelligence can be applied to support financial decision-making in a transparent, reliable and user-controlled manner. Rather than pursuing a single inflexible solution, the system was intentionally designed around two complementary application components that address different user needs and risk profiles.

The first component is a financial search and analysis engine that enables users to explore individual assets through structured market data, contextualized news analysis and interpretable evaluation metrics. This component emphasizes exploratory analysis, information transparency and low-risk user interaction.

The second component explores an advisory-oriented approach in which user-specific financial information is processed to generate structured insights and portfolio-related feedback. Due to the higher stakes associated with personalized financial guidance, this component was treated more cautiously and evaluated primarily with respect to feasibility, interpretability and risk rather than full automation.

Across both components, the system follows a **decision-support paradigm rather than decision replacement.** Artificial intelligence is used to augment information processing and interpretation, while final judgment and control remain explicitly with the user. Based on these objectives, a set of functional and non-functional requirements was derived.

5.1.1 Functional Requirements

From the defined design objectives, the following functional requirements were identified. The system should be able to:

- provide structured access to financial assets for exploratory analysis,
- retrieve and display relevant market data and historical indicators,
- integrate and preprocess financial news from multiple sources,
- support user-driven analysis without enforcing decisions,
- generate interpretable evaluation and scoring metrics,
- process user input for personalized analytical context,
- integrate AI-based interpretation in a controlled and non-autonomous manner.

These requirements apply across both system components, with varying emphasis depending on the interaction context.

5.1.2 Non-Functional Requirements

In addition to functional capabilities, several non-functional requirements significantly influenced the system design and architectural decisions.

Transparency: All calculations and evaluations must be comprehensible.
Modularity: Components should be loosely coupled and independently extendable.

Scalability: The system should allow future expansion.
Reproducibility: Results should be traceable and logically explainable.
Reliability: Errors must be detectable and isolated.
Ethical compliance: No autonomous financial decisions are permitted.

These non-functional requirements directly shaped the architectural structure and constrained the use of artificial intelligence within the system.

5.2 High-Level Architectural Overview

The system follows a layered architecture that separates concerns and ensures a clean division of responsibilities. This approach improves maintainability, testability and extensibility.

At a high level, the system consists of three interacting layers:

1. **Presentation Layer (Frontend)**
2. **Application & Logic Layer (Backend)**
3. **Data & Integration Layer**

Each layer fulfills a distinct role and communicates with the others via clearly defined interfaces.

5.3 Frontend Development and UI Design

User interaction represents a central aspect of the system, as analytical results must be presented in a comprehensible and intuitive manner. The frontend therefore focuses on usability, clarity and structured visualization of financial information.

5.3.1 Design Philosophy

The frontend serves as the primary interaction point between user and system. Its design is guided by the following principles:

- simplicity over complexity,
- visual clarity over information density,
- intuitive navigation,
- immediate feedback for user actions.

Given the target audience - users without deep financial or technical expertise - the interface was designed to minimize cognitive load while still presenting detailed analytical information (especially relevant for graphs and scorings).

5.3.2 User Interaction Model

The frontend supports two primary interaction modes:

1. **Exploratory Mode**
   Users search for financial assets, review data and analyze trends.

2. **Advisory Mode (Conceptual)**
   Users input personal financial information and receive structured feedback.

5.3.3 Dashboard-Based Interface Design

For the advisory mode, a **dashboard-based user interface** was selected due to its suitability for data-intensive applications. Dashboards provide an effective way to present heterogeneous information in a structured and comprehensible manner while maintaining usability and clarity.

By aggregating different types of data within a single view, dashboards allow users to observe multiple indicators simultaneously and to understand relationships between them without navigating across multiple pages. The modular nature of dashboard components further enables flexible layout design and facilitates future extensibility through additional widgets or views.

The implemented frontend includes asset overview panels, time-series visualizations, summaries of key financial indicators, as well as sentiment and scoring visualizations derived from data analysis. In addition, structured input forms are integrated to support user interaction within the advisory context. The use of predefined UI components ensures visual consistency across the application and significantly reduces implementation complexity, allowing development efforts to focus on functionality rather than purely aesthetic concerns.

5.4 Backend Architecture and Application Logic

Behind the user interface, the system relies on a structured backend responsible for coordinating data processing, business logic and external integrations. This part of the architecture ensures consistent data handling and enables communication between analytical components and user-facing interfaces.

5.4.1 Role of the Backend

The backend functions as the central control unit of the system and is responsible for coordinating all core processing tasks. Its primary role is to manage data flow between system components, execute business logic, perform data transformations and interface with both external services and internal AI components.

In addition, the backend manages communication with external APIs, such as financial data providers and news services and acts as an intermediary between the frontend and the AI models. A **Python-based backend** was chosen due to the availability of robust data processing libraries, native compatibility with AI frameworks, flexibility in API development and the ease with which complex logic can be prototyped, tested and debugged.

5.4.2 Internal API Design

To ensure a clean separation between frontend and backend logic, an **internal API was implemented**. This architectural decision provides abstraction from underlying data sources and allows centralized validation and transformation of all incoming and outgoing data.

The internal API enforces consistent data formats and improves maintainability by encapsulating functionality within clearly defined endpoints. Each endpoint serves a specific purpose, such as retrieving asset-related data, executing news analysis, calculating evaluation scores or triggering AI-based assessments. Communication follows a request-response-

model, ensuring that all data returned to the frontend conforms to predefined schemas and validation rules.

5.5 Data Acquisition and Processing

Reliable financial analysis depends on consistent access to high-quality data. The system therefore integrates structured market information and unstructured news content, both of which require dedicated acquisition and preprocessing mechanisms to ensure stability and analytical relevance.

5.5.1 Financial Market Data

Financial market data forms the quantitative foundation of the system. This includes asset prices, historical price movements, trading volumes and basic financial indicators. Due to licensing restrictions and the high cost of professional data providers, the **system relies on publicly accessible data sources**.

To ensure that users receive current information, data retrieval is performed dynamically. While this approach improves timeliness, it also introduces dependencies on external services, which must be considered when evaluating system reliability and scalability.

5.5.2 Asset Identification and Normalization

A major challenge in financial data processing is the unambiguous identification of assets. Stock symbols alone are insufficient, as they are not globally unique, differ across exchanges and may change over time. To address this issue, the system employs a hybrid identification strategy.

Internally, assets are primarily identified using the International Securities Identification Number (ISIN), while ticker symbols are used as secondary identifiers for user interaction and compatibility with external data sources. This dual-reference approach ensures consistent asset mapping and reduces ambiguity without sacrificing usability.

5.5.3 News Acquisition and Preprocessing

Financial news is retrieved via specialized news APIs that aggregate content from a wide range of publishers. Retrieved articles undergo multiple preprocessing steps before analysis. These steps include relevance filtering, duplicate removal, keyword matching and text normalization.

Only articles that meet predefined relevance criteria are forwarded to the analysis stage. This preprocessing pipeline reduces noise, improves data quality and ensures that subsequent analytical components operate on meaningful and contextually relevant information.

5.6 Integration of Artificial Intelligence

Artificial intelligence was originally intended to enhance analytical capabilities within the system, particularly in the interpretation of unstructured information. The following section describes the conceptual role of AI and the mechanisms designed to control and evaluate its integration.

5.6.1 Initial Role of AI in the System

The conceptual system design initially assumed artificial intelligence as a central analytical component. The implementation process and resulting architectural refinements are discussed

in later chapters. The initial approach involved integrating artificial intelligence as a complementary analytical layer intended to extend conventional data processing through the interpretation of unstructured textual information rather than replacing deterministic logic.

The primary tasks assigned to AI included sentiment analysis of financial news, context-related interpretation of events, text summarization and contribution to qualitative evaluation metrics. To ensure data protection and maintain full control over execution and data processing, the AI component was accessed via a locally hosted large language model interface.

5.6.2 Prompt Design and Control Mechanisms

To achieve consistent and interpretable AI outputs, structured prompts were developed. These prompts explicitly define the expected output format, evaluation criteria, scope limitations and response constraints. The goal of this approach was to reduce ambiguity and improve output reliability.

Despite these measures, inconsistencies in AI-generated outputs remained a significant challenge. This observation highlights fundamental limitations of current LLM technology, particularly in applications that require deterministic behaviour and strict adherence to predefined structures. Due to these existing problems, **any influence of LLMs and AI was removed from the final application**. The exact reasons and replacement solutions are explained in the following chapters.

5.7 Scoring and Evaluation Model

The system's scoring mechanism combines multiple dimensions:

1. **Quantitative metrics** (price movement, volume, market capitalization)
2. **Qualitative analysis** (news sentiment, event relevance)
3. **Contextual weighting** (market environment, volatility)

Each component contributes to a normalized score. The final value represents an aggregated assessment rather than a definitive recommendation.
This design ensures transparency while allowing users to interpret results independently.

5.8 Design Constraints and Architectural Trade-Offs

Several trade-offs influenced the final design:

- **Accuracy vs. Complexity:**
  More complex models increased error potential and reduced interpretability.
- **Automation vs. Control:**
  Full automation was avoided to reduce financial risk.
- **Scalability vs. Feasibility:**
  Continuous data crawling was abandoned due to resource limitations.
- **Innovation vs. Reliability:**
  Experimental AI features were limited to non-critical components.

These trade-offs were necessary to ensure a stable and defensible system design.

6 Technical Implementation

The following chapter translates the previously defined conceptual architecture into a concrete technical realization. It describes the system's development paths, implementation decisions and practical constraints encountered during the construction of the prototype.

6.1 Purpose and Scope of the Implementation

The implementation followed two distinct but closely related development paths, each addressing different objectives and technical challenges within the broader context of AI-supported financial decision support.

The first implementation path focused on the financial search and analysis engine. Its objective was to develop a functional and stable prototype capable of aggregating financial market data, processing and analyzing news content, initially integrating AI-based interpretation and presenting interpretable results through a user-oriented interface. This component represents the primary realized outcome of the implementation and serves as the main basis for subsequent evaluation.

The second implementation path explored the technical feasibility of an AI-supported advisory dashboard based on user-specific input. This path was implemented selectively to investigate limitations related to output consistency, determinism, interpretability and risk. Rather than aiming for full functional completeness, the advisory implementation served as an experimental evaluation of the practical boundaries of current AI technology in high-stakes financial contexts.

Accordingly, this chapter presents both successfully implemented functionality and intentionally constrained developments. The focus lies on feasibility, transparency and critical assessment under realistic technical constraints, rather than on production readiness or commercial scalability.

6.2 Development Environment and Technology Stack

The technical realization of the system required a carefully selected set of development tools and execution models. The chosen technologies were intended to ensure flexibility, local controllability and compatibility with AI-driven components.

6.2.1 Programming Languages and Tools

The system was implemented using the following technologies:

- **Python** - backend logic, data processing, AI integration
- **HTML / CSS / JavaScript** - frontend implementation
- **REST-based APIs** - communication between frontend and backend
- **Ollama** - execution of a locally hosted Large Language Model
- **CSV-based datasets** - asset reference storage
- **Third-party APIs** - financial data and news retrieval

Python was chosen due to its extensive ecosystem for data processing, AI integration and rapid prototyping. Its compatibility with external APIs and LLM frameworks made it particularly suitable for this project.

6.2.2 System Execution Model

The application follows a locally executed client-server model:

- The frontend runs in a web browser.
- The backend runs as a local Python service.
- Communication occurs via HTTP requests.
- The LLM is executed locally and accessed via an internal API.

This approach allows full control over data flow and avoids reliance on external cloud-based AI services.

6.3 Implementation of Asset Search and Identification

Accurate asset identification forms the technical foundation of the search engine. Since financial instruments cannot be reliably distinguished by ticker symbols alone, a structured identification strategy was required to ensure consistent data retrieval and downstream processing.

6.3.1 Asset Reference Dataset

One of the first technical challenges was the construction of a reliable asset reference dataset. Financial assets cannot be uniquely identified by ticker symbols alone, as these are exchange-specific and often reused.

To address this, a reference dataset was created containing:

- Company name
- Ticker symbol
- ISIN

This dataset was stored in CSV format and loaded into memory at application startup.

6.3.2 Symbol-ISIN Mapping Strategy

During user interaction, financial assets are searched using ticker symbols or company names, as these identifiers are familiar and intuitive for users. Internally, however, all further processing relies on the International Securities Identification Number (ISIN) as the primary reference. This separation between user-facing identifiers and internal identifiers was introduced to ensure technical correctness and consistency across data sources.

To achieve this, the system implements a structured resolution process. User input is first matched against a symbol-based reference list to identify potential assets. Once a match is found, the corresponding ISIN is resolved and used as the internal reference for all subsequent data retrieval and processing steps. All asset-related operations are therefore performed using the ISIN rather than the ticker symbol.

1. User input → symbol-based lookup
2. Symbol → ISIN resolution
3. ISIN → internal data reference

This approach ensures consistent data retrieval across multiple sources, avoids ambiguities caused by non-unique ticker symbols and maintains compatibility with external APIs that rely on different identification schemes. As a result, asset handling remains both user-friendly and technically robust.

6.4 Financial Data Retrieval and Processing

Reliable quantitative evaluation depends on consistent, validated and comparable market data. Integrating heterogeneous financial information from publicly available sources therefore required dedicated mechanisms for acquisition, normalization and error handling to ensure analytical stability.

6.4.1 Market Data Acquisition

Market data such as historical prices, volume and basic indicators were retrieved using publicly accessible data sources. Due to the absence of free institutional-grade APIs, a hybrid approach was implemented:

- API-based retrieval where possible
- Structured extraction from public sources
- Fallback mechanisms for missing values

All retrieved data is processed in real time to ensure current information.

6.4.2 Data Normalization and Validation

Incoming financial data varies significantly in format, precision and completeness depending on its source. To ensure consistency and comparability, a dedicated normalization and validation layer was implemented.

This preprocessing stage converts all numerical values into uniform formats, aligns timestamps across datasets and handles missing or null values in a controlled manner. Invalid or implausible entries are filtered out to prevent the propagation of errors into later processing stages.

These normalization steps are essential for meaningful comparison between assets and play a critical role in ensuring that the subsequent scoring logic operates on reliable and consistent data.

6.5 News Retrieval and Processing Pipeline

In addition to structured market data, unstructured textual information plays a central role in the system's evaluation logic. The news processing pipeline was therefore designed to ensure relevance, consistency and analytical suitability of retrieved articles before further assessment.

6.5.1 News Acquisition

Financial news articles are retrieved via a news aggregation API. Search queries are dynamically generated based on asset names and relevant keywords.

To improve relevance, articles are filtered based on:

- publication date,
- source credibility,

- keyword relevance,
- duplication detection.

Only articles passing these filters are forwarded to the analysis pipeline.

6.5.2 Text Preprocessing

Before being passed to the evaluation pipeline, news articles undergo several preprocessing steps. These include text cleaning and normalization, removal of non-informative elements such as boilerplate content, keyword extraction and sentence segmentation. These steps ensure that subsequent analytical components operate on structured, comparable and contextually relevant input data.

6.5.3 Initial Approach: AI-Based News Interpretation

During the initial system design phase, the preprocessing pipeline was primarily developed to prepare textual input for analysis by a Large Language Model (LLM). The objective of this approach was to leverage the contextual interpretation capabilities of modern language models to evaluate financial news articles. By reducing linguistic noise and standardizing input structure, the system aimed to improve the reliability and consistency of AI-generated sentiment assessments and contextual interpretations.

The LLM-based approach was intended to automatically extract sentiment, identify relevant financial events and contribute qualitative insights to the asset evaluation process. However, experimental implementation revealed significant limitations. Despite extensive prompt engineering and structured input preparation, the model exhibited **inconsistent output formatting, probabilistic variability and limited reliability in producing reproducible evaluation** results. These issues posed a substantial challenge for integration into a deterministic scoring pipeline and introduced potential risks regarding interpretability and result stability.

6.5.4 Actual Approach: Deterministic Multi-Dimensional News Score

Due to the identified limitations of LLM-based evaluation (analyzation follows in later chapters), the final implementation replaced the AI-driven interpretation with a deterministic, rule-based scoring model. This model evaluates financial news articles using multiple predefined assessment dimensions to determine the relevance and impact of each article for a specific company.

First, sentiment polarity is assessed using curated keyword lists that identify positive and negative linguistic indicators. The scoring mechanism also considers linguistic intensity by differentiating between strong and uncertain expressions, allowing sentiment weights to be amplified or dampened accordingly.

Second, article relevance is determined by analyzing the frequency and contextual placement of company references within the text. Higher mention density indicates stronger direct relevance of the article to the evaluated asset.

Third, event-based classification is applied by detecting predefined financial topics such as mergers, earnings growth, strategic expansions or corporate scandals. Each recognized event category contributes a weighted positive or negative adjustment based on its expected financial significance.

To ensure temporal relevance, the aggregated base score is multiplied by a time-decay factor, giving greater influence to recent news articles. Additionally, a source reliability factor adjusts the score based on the perceived credibility and authority of the publishing media outlet.

Finally, the resulting value is normalized and constrained to a bounded evaluation range, representing the overall sentiment and informational relevance of the analyzed news article.

This deterministic scoring approach improves reproducibility, interpretability and integration stability while still capturing key qualitative aspects of financial news analysis. The transition from probabilistic AI interpretation to rule-based evaluation represents a central design decision aimed at balancing analytical sophistication with technical reliability.

6.6 Scoring Architecture and Integration of AI

The scoring architecture represents the central evaluation mechanism of the application and combines multiple quantitative and qualitative data sources to generate a consolidated asset assessment. During the initial system design, the qualitative interpretation component was intended to be primarily supported by a Large Language Model (LLM), thereby extending traditional data-driven evaluation methods with contextual analysis of unstructured information.

While quantitative indicators were always designed to provide the structural foundation of the scoring system, the LLM was originally planned to deliver dynamic qualitative interpretation of financial news and events. Consequently, the overall scoring framework was initially conceptualized as a hybrid system combining deterministic market data evaluation with AI-based contextual analysis.

During implementation, however, several technical constraints became apparent, requiring an architectural refinement to ensure system stability, reproducibility and operational reliability.

6.6.1 Intended AI-Supported Scoring Framework

The original system design integrated a locally hosted Large Language Model using the Ollama framework. The objective was to enable automated interpretation of financial news, extraction of sentiment and event relevance, and structured qualitative contribution to the overall asset evaluation.

The model was accessed via a local REST interface and invoked as part of the qualitative analysis pipeline. Prompt structures were designed to define analytical tasks, enforce structured output formats and ensure compatibility with automated scoring aggregation.

While the LLM demonstrated strong contextual understanding capabilities, several limitations emerged during implementation. Output variability, inconsistent formatting and probabilistic deviations created integration challenges within a deterministic scoring environment. Despite prompt optimization and output validation mechanisms, reproducibility could not be consistently ensured. These observations correspond with known limitations of foundation models. (Bommasani et al., 2022).

Because qualitative evaluation directly influences the overall asset score, these limitations affected the reliability of the entire scoring architecture rather than a single subsystem.

6.6.2 Transition to Deterministic Multi-Dimensional Scoring

To ensure consistent and reproducible scoring results, the qualitative evaluation component was redesigned as a deterministic multi-dimensional scoring model. This redesign preserves the conceptual goal of integrating qualitative information while replacing probabilistic interpretation with rule-based evaluation mechanisms.

The qualitative component evaluates financial news using multiple assessment dimensions, including sentiment polarity, company relevance, event-based classification, temporal

weighting and source credibility adjustment. The resulting evaluation is normalized to a score within a bounded range.

This deterministic qualitative evaluation is combined with quantitative market indicators to form a hybrid scoring architecture that balances interpretability, stability and analytical coverage.

6.7 Quantitative Evaluation and Final Score Aggregation

Once qualitative and quantitative inputs are defined, their consolidation requires a coherent aggregation framework. Combining heterogeneous indicators into a single interpretable metric demands standardization, weighting and clear structural logic to ensure consistency and transparency of the resulting asset evaluation.

6.7.1 Quantitative Scoring Component

Quantitative indicators form the structural basis of the asset evaluation. These metrics are normalized and transformed into comparable scoring dimensions. The main indicators include:

- price development
- trading volume
- relative market size

Each metric contributes to a normalized sub-score within predefined bounds, ensuring comparability across assets and market segments.

6.7.2 Final Score Aggregation

The final asset score is computed as a weighted combination of quantitative and qualitative evaluation components. The weighting scheme is intentionally conservative to prevent excessive influence of qualitative interpretation and to maintain numerical stability of the overall scoring system.

The aggregated score is designed as an informational indicator supporting comparative analysis and exploration rather than providing automated investment recommendations. The system therefore operates as a decision-support tool rather than an autonomous advisory mechanism.

6.8 Performance and Architectural Implications

The refinement of the scoring architecture was primarily driven by implementation constraints identified during development. Key limitations included:

- LLM response latency, which accumulated significantly when processing multiple information sources
- inconsistent output formatting and probabilistic evaluation results
- inability to guarantee deterministic scoring behaviour
- limited reliability of publicly available financial data sources
- scalability constraints affecting real-time system performance

Because the qualitative evaluation component directly influences the overall scoring results, these limitations affected the reliability of the entire application. The transition to a deterministic

scoring framework therefore represents an architectural trade-off between advanced AI interpretation and system-level stability.

The final system design prioritizes **reproducibility, transparency and scalability** while still incorporating structured qualitative information into the asset evaluation process.

7 System Demonstration and UI Overview

The implemented prototype is best understood through its observable behaviour and user interaction patterns. The following chapter therefore presents selected interface views and illustrates how system functionality is made accessible to users.

7.1 Purpose of the Demonstration

The objective is to illustrate how the developed application behaves from a user perspective and how the implemented functionality is exposed through the user interface.

The focus relies on visualizing real system components and their interaction rather than introducing new technical concepts. By documenting the observable behaviour of the prototype, this section supports transparency and facilitates understanding of how the system can be used in practice. Conceptual elements and planned extensions are addressed separately and are clearly distinguished from implemented functionality.

7.2 Implemented Financial Search Engine

The financial search engine represents the primary functional component of the application and serves as the main entry point for exploratory asset analysis. The following subsections describe the key interface elements and their role in supporting user-driven financial research.

7.2.1 Landing Page

The landing page represents the primary entry point to the application and follows the design principles of well-known search engines. As shown in *Figure 1*, the interface is intentionally minimalistic, focusing the user's attention on the central search functionality.

Users can directly search for financial assets using either the asset name or its ticker symbol. In addition, a prominently placed button enables access to the portfolio generation process.



Figure 1: Landing Page of Search Engine

This dual entry concept reflects the system's two main use cases: exploratory asset research and portfolio-oriented analysis.

At the bottom of the page, transparency is further supported by explicitly referencing the partner API used for data acquisition and by displaying the number of accessible asset symbols. This design choice communicates data provenance and scope to the user and aligns with the overall emphasis on traceability.

7.2.2 Search Function

The search function supports dynamic and incremental asset lookup, as illustrated in Figure 2. As users begin typing, the system provides real-time suggestions based on the available asset reference dataset. This behaviour reduces input errors and improves usability, particularly for users unfamiliar with exact ticker symbols.

At the current stage of development, the search engine supports stock assets. However, the underlying design was intentionally conceived to allow extension to additional asset classes such as ETFs, commodities or cryptocurrencies. The suggestion-based interaction model therefore represents a scalable foundation rather than a stock-specific solution.

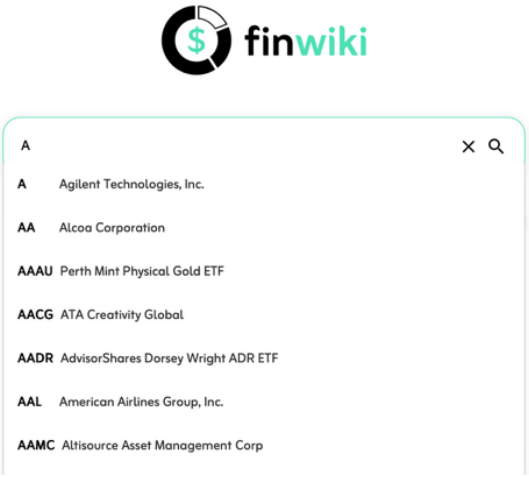

Figure 2: Search Function on Landing Page

7.2.3 Loading State Visualisation

To ensure transparency during data-intensive operations, the system employs explicit loading state visualizations. These states inform the user about the current processing stage and reduce uncertainty during longer backend operations.

As shown in Figures 3-5, the loading process is divided into multiple semantically meaningful steps, such as asset lookup, news retrieval and result preparation. Internally, these states correspond to backend requests for financial data and news analysis, including aggregation, ranking and scoring.

Once all processing steps are completed, the user is automatically forwarded to the updated asset overview page. Although the waiting time is static rather than dynamically calculated, the staged visualization improves perceived responsiveness and user trust.

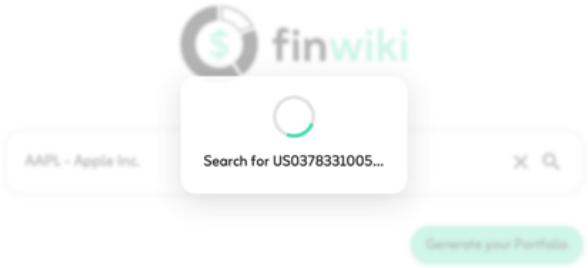

Figure 3: Loading State - Search

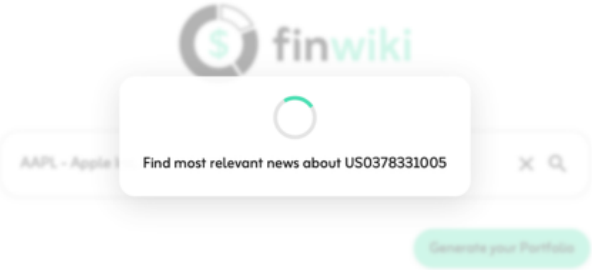

Figure 4: Loading State - News

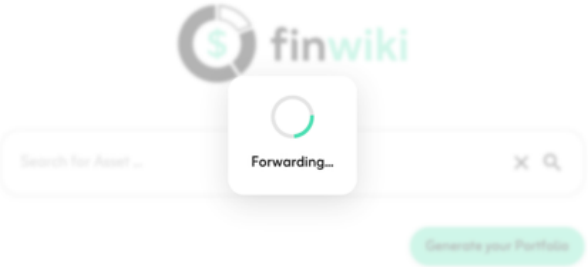

Figure 5: Loading State - Forwarding

7.2.4 Asset Overview Page - General Information

The asset overview page provides a consolidated view of the selected asset's core information. As illustrated in Figure 6, this section presents general company data, key financial statistics and a historical price chart.

The intent of this layout is to offer users an immediate high-level understanding of the asset before engaging with more detailed analytical components. By combining textual descriptions,

numerical indicators and visual time-series data, the interface supports different information preferences and levels of financial literacy.

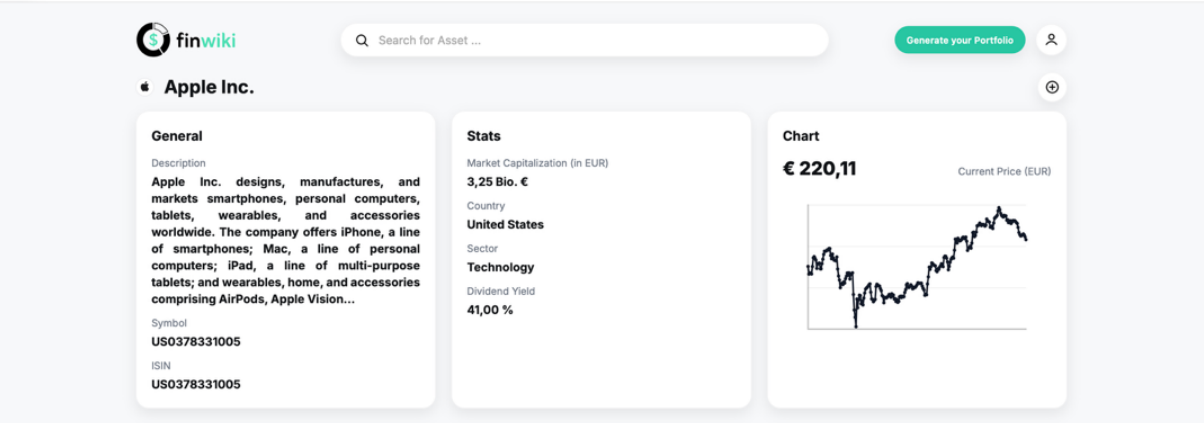

Figure 6: Asset Overview Page - General Information

7.2.5 Asset Overview Page - Assessment

The assessment section translates quantitative and qualitative indicators into an interpretable evaluation framework. Individual metrics such as valuation price indicators, money flow index, relative volume, on-balance volume, short interest and earnings revisions are presented in a structured manner.

Each indicator represents a distinct analytical perspective. Money-flow-based metrics reflect trading activity and capital movement, while earnings revisions capture changes in market expectations. In addition to technical indicators, basic qualitative criteria such as market positioning, scalability, cost structure and return profile are incorporated.
News-related criteria, including overall market sentiment, current price drivers, unexpected events and volatility, further contribute to the assessment. The final score represents an aggregated average of these dimensions and is explicitly presented as an informational indicator rather than a definitive investment recommendation.

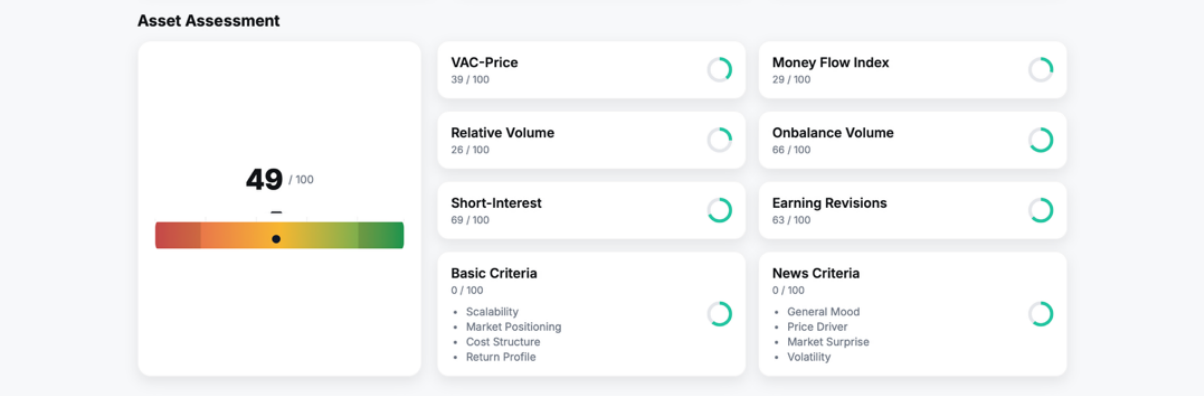

Figure 7: Asset Overview Page - Assessment

7.2.6 Asset Overview Page - News

The news section displays up to fifteen of the most relevant and impactful news articles associated with the selected asset, as shown in Figure 8. Each entry includes the publication date, a sentiment-based score and a direct link to the original source.

This section is included to provide contextual grounding for the assessment results and to allow users to independently verify and explore underlying information. Articles are ranked and scored by the system in the background, primarily based on relevance and estimated market impact.

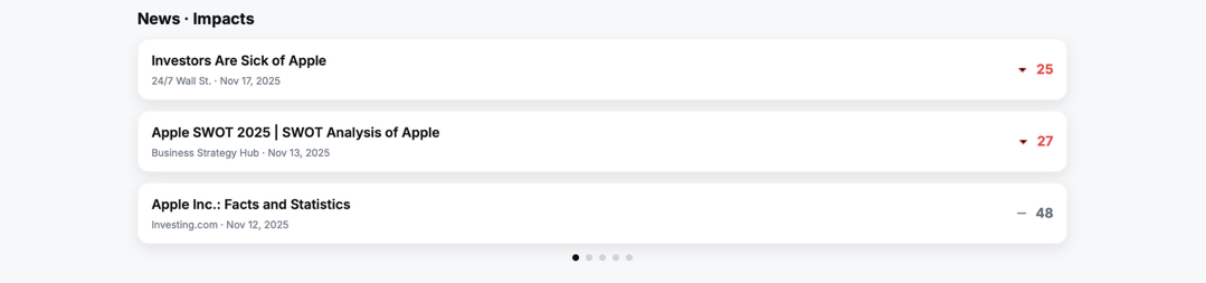

Figure 8: Asset Overview Page - News

7.3 Implemented Components of Advisory Dashboard

Beyond exploratory asset analysis, the system also includes a partially implemented advisory dashboard designed to process user-specific financial information. The following subsections present the realized interface components and outline their intended analytical purpose.

7.3.1 Planned Sections

During the design phase, a comprehensive advisory dashboard was conceptualized. Planned sections included an overview page displaying a suggested portfolio, a risk assessment page linking personal financial data to portfolio risk, an asset control page for manual adjustments, simulation views for historical and projected performance, a portfolio-related news section and a wiki-style information view.

Of these components, only the survey, the overview page and the risk assessment page were realized in the implemented prototype. The remaining components were intentionally excluded due to technical limitations and unresolved reliability concerns, as discussed in previous chapters.

7.3.2 Survey and Question Types

The survey interface supports multiple question types to capture diverse user information. These include multiple-choice, single-choice, binary (yes/no), numeric input and free-text responses.

The questions cover financial situation, investment goals, risk tolerance, asset preferences and experience levels. This structured yet flexible approach enables comprehensive user profiling while maintaining usability.

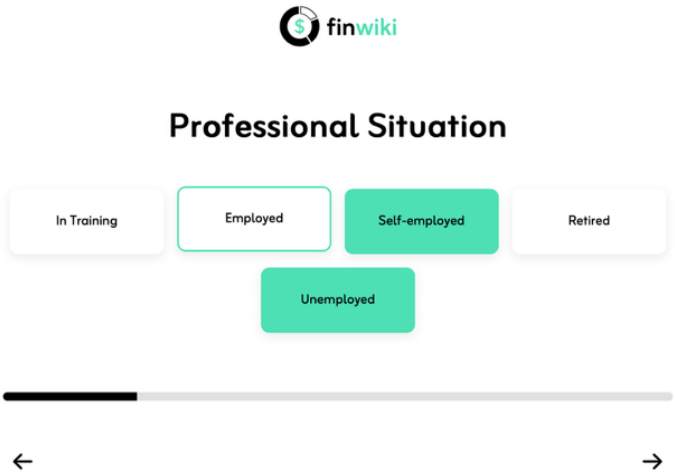

Figure 9: Survey Question (Multiple Choice)

Question	Type	Answer Possibilities
Generation Request	Single-Choice	Portfolio / Savings Plan / Portfolio & Savings Plan
Personal Net-Income	Numeric Input	
Existing Assets	Numeric Input	
Exisiting Debts	Numeric Input	
Professional Situation	Multiple-Choice	In Training / Employed / Self-Employed / Retired / Unemployed
Available Starting Capital	Numeric Input	
Desired Monthly Savings Amount	Numeric Input	
Uninvested Capital	Numeric Input	
Planned Time Horizon	Numeric Input	
Investment Goal	Free-Text Input	
Planned Interim Withdrawals	Binary	Yes / No
Permitted Asset Classes	Multiple-Choice	Stocks / ETFs / Commodities / Crypto / Funds
Willingness to Take Risks	Single-Choice	Low Fluctuations / Medium Fluctuations / High Fluctuations
Sector Restrictions	Free-Text Input	
Region Restrictions	Free-Text Input	
Minimum Number of Positions	Numeric Input	
Maximum Number of Positions	Numeric Input	
Personal Experience with Investing	Binary	Yes / No
Daily Self-Updating of Financial Situation	Binary	Yes / No
Other Comments and Requests	Free-Text Input	

Table 1: Survey Questions Overview

7.3.3 Calculation Waiting Screen

After completing the survey, users are forwarded to a calculation waiting screen, shown in Figure 10. This screen provides the system with sufficient time to process survey data and perform AI-based analysis.

The waiting screen displays a sequence of predefined status messages, such as data upload, analysis execution, visualization preparation and portfolio optimization. Although the progression is time-based rather than dynamically measured, the staged messaging communicates system activity and mitigates perceived delay. Upon completion, users are automatically forwarded to the dashboard.

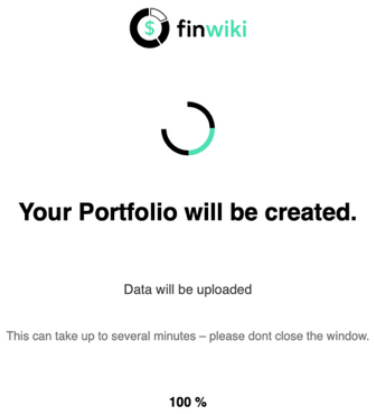

Figure 10: Calculation Waiting Screen

7.3.4 Dashboard - Overview Page

The dashboard overview page presents a high-level summary of the generated portfolio. As illustrated in Figure 11, asset allocation is visualized through a circular chart showing the proportional distribution across asset classes.

Additional panels display extracted goals, inferred risk classification and key investment parameters derived from the survey. A detailed list of selected assets is shown at the bottom, including asset type and individual risk scores. This layout prioritizes clarity and enables users to quickly understand the reason behind portfolio composition.

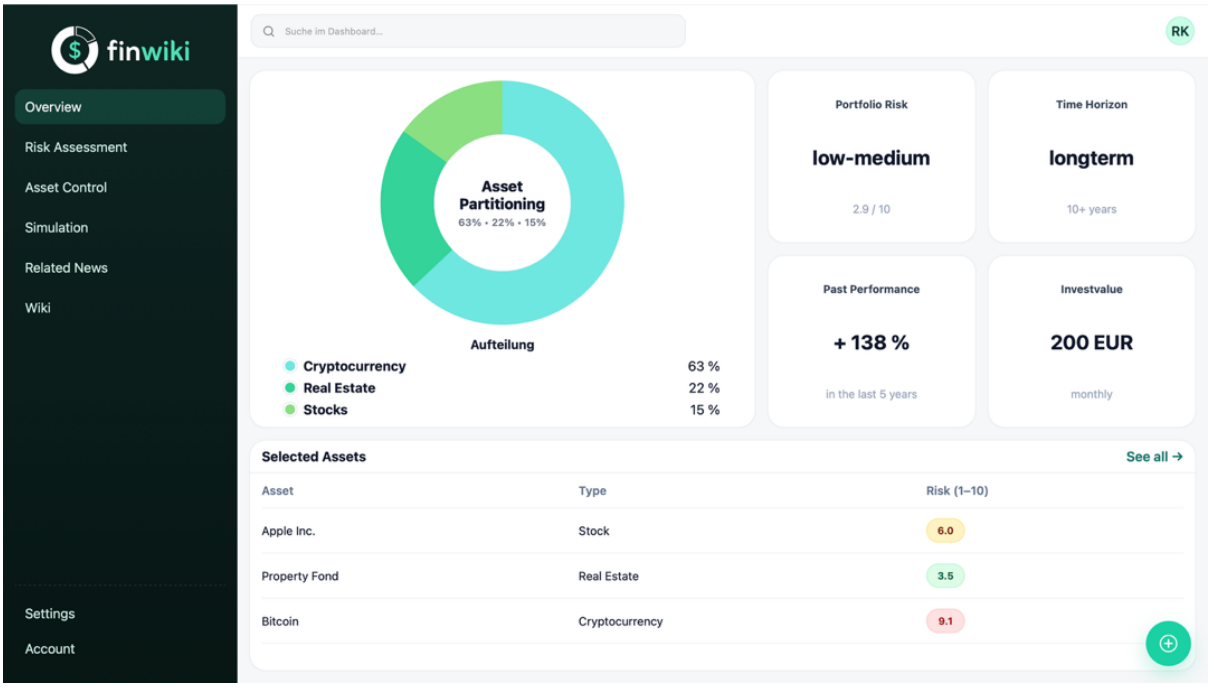

Figure 11: Dashboard - Overview Page

7.3.5 Dashboard - Risk Assessment Page

The risk assessment page focuses on the traceability of risk-related decisions. It visualizes multiple risk dimensions, such as income stability, financial obligations and overall portfolio risk, as shown in Figure 12.

Interactive information icons allow users to inspect the reasoning behind individual assessments. This feature directly addresses concerns raised in the empirical survey regarding transparency and explainability and reinforces the system's decision-support orientation.

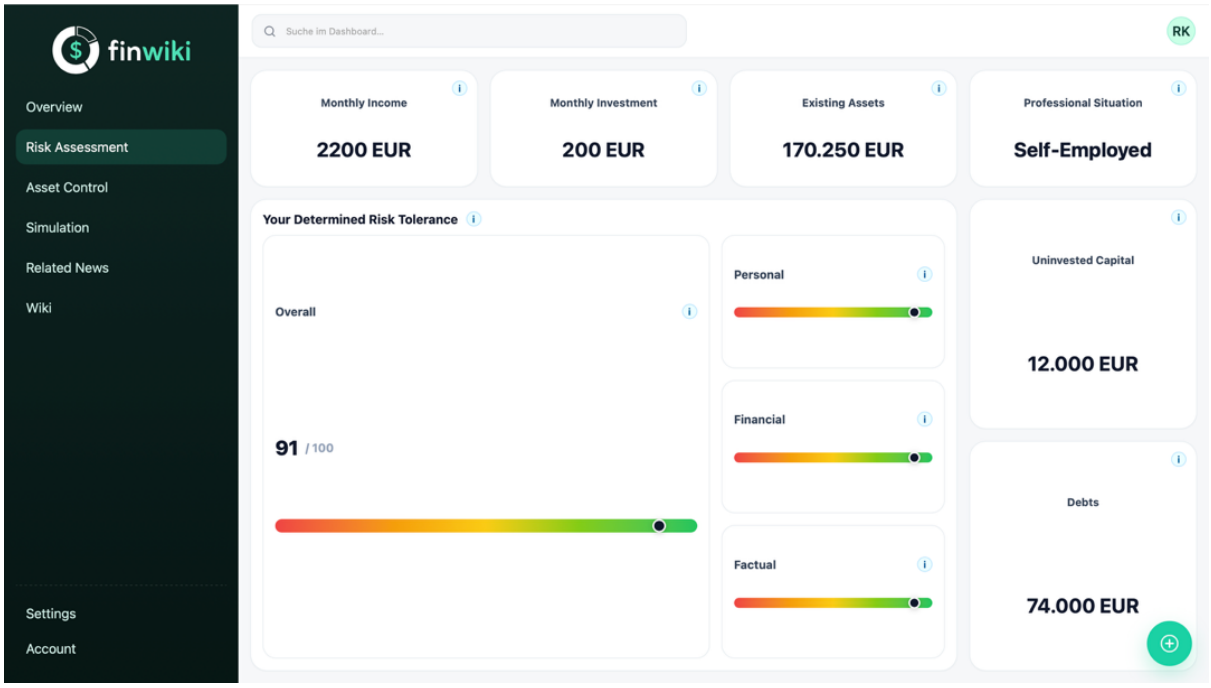

Figure 12: Dashboard - Risk Assessment Page

7.4 User Interaction Flow

The system supports two complementary user interaction paths that reflect different usage intentions. Users may either explore individual financial assets through the search engine or engage in a more structured, portfolio-oriented analysis via the survey and dashboard components.

Search Engine
In the exploratory path, users search for an asset using the landing page and are guided through a short loading sequence while relevant data and news are processed. They are then presented with a consolidated asset overview, including market data, assessment indicators and curated news content. This path allows users to freely navigate between assets and analyze information without further commitment.

Portfolio Dashboard
In the portfolio-oriented path, users complete a structured questionnaire that captures financial context, goals and risk preferences. After a brief processing phase, the system presents the results in a dashboard view, providing an overview of asset allocation and risk assessment. Both interaction paths are explicitly non-binding and designed to support informed interpretation rather than automated decision-making.

8 Evaluation and Analysis of Results

The practical value of the developed prototype becomes evident through a systematic assessment of its functionality and limitations. The following chapter examines how effectively the implemented components fulfil the objectives defined earlier and evaluates their performance under realistic usage conditions.

8.1 Evaluation Objectives and Methodological Approach

The evaluation does not aim to measure performance against industrial benchmarks, as the system was developed as a research prototype rather than a production-ready solution. Instead, the focus relies on assessing practical feasibility, reliability, interpretability and conceptual validity.

Given the exploratory nature of the project, a qualitative evaluation approach was adopted. The system was evaluated through continuous testing during development, targeted test scenarios and systematic observation of system behaviour under realistic usage conditions. This approach enables an in-depth assessment of interactions between system components, particularly the interplay between structured financial data, AI-generated interpretations and user-facing representations.

Importantly, the evaluation differentiates between two distinct implementation paths pursued in this thesis. The first path concerns the financial search and analysis engine, which was realized as a functional prototype. The second path concerns the advisory-oriented implementation, which was explored experimentally to evaluate feasibility and limitations rather than to achieve full automation. Both paths are evaluated separately to reflect their different objectives and risk profiles.

8.2 Evaluation of the Financial Search and Analysis Engine

The financial search and analysis engine represents the core functional outcome of the development process. Its evaluation focuses on data reliability, analytical consistency and the effectiveness of the implemented scoring and interpretation mechanisms.

8.2.1 Asset Identification and Data Integrity

One of the most significant achievements of the implemented search engine is the stability and correctness of asset identification. The decision to rely on ISINs as internal identifiers proved highly effective in preventing ambiguities commonly associated with ticker symbols, particularly in multi-exchange contexts.

During evaluation, asset resolution remained consistent even in cases where ticker symbols overlapped or differed across markets. This stability enabled reliable downstream processing and ensured that market data and news content were correctly associated with the intended asset. However, evaluation also revealed that asset coverage is inherently constrained by the availability and completeness of external data sources.

Widely traded and well-documented assets were consistently supported, while less prominent securities occasionally exhibited incomplete datasets. This limitation reflects structural characteristics of publicly accessible financial data rather than shortcomings of the system design. Within these constraints, asset identification and data integrity can be considered robust and well-suited for exploratory financial analysis.

8.2.2 Market Data Retrieval and Preprocessing

The market data retrieval pipeline performed reliably under most test conditions. The implemented normalization and validation mechanisms successfully addressed inconsistencies in numerical formats, missing values and timestamp alignment, enabling meaningful comparison across assets.

Evaluation revealed, however, that technical correctness does not always imply semantic accuracy. In some cases, data providers returned delayed or contextually incomplete information, which could lead to interpretations that diverge from actual market conditions. This observation highlights an important limitation: even well-designed analytical systems remain fundamentally dependent on the quality of their underlying data.

The preprocessing pipeline proved essential in reducing such issues by stabilizing inputs and preventing error propagation. Nevertheless, the evaluation confirms that data quality represents a critical external constraint that cannot be fully resolved through technical means alone.

8.2.3 Assessment of LLM-Based News Interpretation

The application initially aimed to integrate Large Language Models as a core component for qualitative interpretation of financial news. Evaluation demonstrated that LLMs possess strong capabilities in identifying sentiment trends, extracting contextual relationships and generating structured summaries of complex articles.

When used interactively or in controlled analytical settings, the model frequently produced detailed and contextually accurate interpretations that closely aligned with human assessments. This capability directly addresses the challenge of information overload identified in both the theoretical analysis and the empirical survey.

However, the evaluation revealed that the primary limitation was not analytical quality, but technical reproducibility and operational reliability. Within an automated processing pipeline, LLM outputs exhibited variability in structure, formatting and numerical consistency, even when identical input data and prompt instructions were provided. These characteristics complicate integration into deterministic evaluation frameworks and increase validation overhead.

These findings are consistent with known properties of probabilistic language models (Bommasani et al., 2022). While LLMs demonstrate strong interpretative performance in interactive or semi-structured environments, current technical limitations restrict their suitability for fully automated, reproducible scoring processes.

Within the scope of the application, LLM-based interpretation therefore remains a conceptually valuable analytical approach but was not implemented as the primary driver of the final scoring architecture.

8.2.4 Evaluation of the Final Scoring Architecture

The scoring mechanism was evaluated as the central method for aggregating heterogeneous financial information into a compact and interpretable indicator. The final scoring architecture combines quantitative market metrics with a structured qualitative evaluation model to support comparative asset exploration and user orientation.

Evaluation demonstrated that the scoring system responds sensitively to short-term news developments and market dynamics. This responsiveness improves the system’s ability to reflect current information flows but may introduce short-term fluctuations in asset evaluation.

Additionally, qualitative components inherently introduce weighting assumptions that influence score composition.

The evaluation further confirmed that deterministic qualitative scoring significantly improved reproducibility and system stability compared to the originally intended AI-based evaluation approach. While the LLM-based interpretation demonstrated strong contextual understanding, it introduced variability that conflicted with the reliability requirements of the overall application.

Consequently, the final scoring model prioritizes interpretability, transparency and reproducibility. The resulting score does not represent an objective measure of asset quality but functions as an informational indicator supporting exploratory analysis and comparison. Presenting the score as a decision-support metric rather than an investment recommendation is therefore both technically justified and ethically appropriate.

8.3 Evaluation of the Advisory-Oriented Implementation

In contrast to the fully realized search engine, the advisory-oriented implementation primarily served as an experimental investigation into the feasibility of automated financial guidance. The following evaluation therefore concentrates on technical reliability, output consistency and the limitations encountered during development.

8.3.1 Output-Consistency and Determinism

The experimental evaluation of the advisory-oriented implementation revealed a fundamental limitation in the use of Large Language Models for structured, high-stakes output generation. The core issue was not merely inconsistent results or formatting errors, but the inability of the model to reliably satisfy simultaneous constraints on both content correctness and output structure.

In practical terms, the system repeatedly encountered a trade-off between structural validity and semantic correctness. In many cases, the language model produced outputs that adhered to the requested structure, such as a valid JSON format, but contained incorrect, incomplete or improperly ordered values. In other cases, the generated content was semantically plausible and internally consistent, but violated the required structural constraints through missing elements, incorrect nesting or malformed syntax.

This behaviour persisted despite extensive prompt engineering, explicit formatting instructions and clearly defined output constraints.

8.3.2 Reliability of Structured Output Generation

To address this issue, multiple validation and correction strategies were implemented, including recursive checking loops that repeatedly prompted the model to correct malformed outputs. Structural validation could be enforced to a limited extent by re-invoking the model until syntactic constraints were satisfied.

However, this approach revealed a critical limitation: while structural compliance could sometimes be achieved through repeated correction, semantic correctness could not be preserved simultaneously. In many cases, structurally valid outputs exhibited new semantic errors, such as missing values, incorrect ordering, incomplete datasets or altered content. Conversely, attempts to correct semantic issues often reintroduced structural violations.

As a result, the correction process did not converge toward a stable solution. Instead, it oscillated between structurally valid but semantically incorrect outputs and semantically

plausible but structurally invalid outputs. Although correct outputs occasionally occurred by chance, they could not be reproduced consistently or reliably.

8.3.3 Implications for Advisory Automation

This evaluation demonstrates that current Large Language Models are not suitable for applications that require both deterministic structure and guaranteed semantic correctness. In the context of financial advisory systems, where outputs must be reproducible, complete and verifiable, this limitation represents a critical barrier.

The inability to enforce reliable convergence through validation and correction loops indicates that the issue cannot be resolved through additional engineering effort alone. Instead, it reflects a fundamental mismatch between probabilistic language generation and the deterministic requirements of financial decision support systems.

Consequently, the decision to abandon the implementation of a fully automated advisory component was not based on incomplete development, but on empirical evidence that current AI models cannot meet the necessary reliability and accountability standards for such use cases.

8.4 Performance Characteristics and Practical Feasibility

From a performance perspective, the system operates reliably under experimental conditions. However, evaluation revealed several constraints that limit scalability and real-time applicability. AI inference introduces noticeable latency and real-time analysis of multiple assets significantly increases computational demands.

In addition, external API rate limits restrict requests and continuous data acquisition would require substantial infrastructure resources. These findings indicate that while the system is suitable for research and prototyping, production-level deployment would require significant architectural redesign and investment.

8.5 Summary of Evaluation Results

In summary, the evaluation demonstrates that the financial search and analysis engine successfully achieves its intended purpose as an AI-supported decision-support tool. It effectively aggregates information, integrates contextual analysis and presents interpretable results to users.

In contrast, the advisory-oriented implementation path reveals fundamental limitations of current AI technology when applied to personalized, high-stakes financial decision-making. These limitations are not the result of implementation shortcomings but reflect inherent characteristics of probabilistic language models.

The evaluation confirms that AI is well suited for supporting analysis and interpretation, but remains unsuitable for autonomous financial decision-making. These findings provide a strong empirical foundation for the discussion in the following chapter.

9 Discussion and Critical Reflection

The results obtained during system development and evaluation provide valuable insight into both the opportunities and boundaries of AI-supported financial analysis. Building on the empirical and technical findings presented previously, the following chapter reflects on their broader implications, conceptual significance and practical relevance.

9.1 Aim and Position of Discussion

Rather than presenting additional technical details, the discussion focuses on the implications of the results and their significance for the application of artificial intelligence in financial analysis.

The core objective of the project was to examine whether modern AI technologies, particularly Large Language Models, are capable of meaningfully supporting or partially automating financial analysis and investment-related decision processes. This chapter therefore evaluates the extent to which this objective was achieved and derives conclusions from the observed strengths and limitations of the proposed approach.

The discussion explicitly differentiates between the outcomes of the financial search and analysis engine and the advisory-oriented implementation path.

9.2 Interpretation of the Results

The evaluation clearly demonstrated that AI-based systems can provide valuable support in the analysis of financial information. In particular, the system proved effective at aggregating large volumes of data, extracting relevant information from unstructured sources and presenting insights in a condensed and accessible form. These capabilities address a key challenge faced by private investors: information overload.

At the same time, the results show that the strengths of AI lie primarily in interpretation and summarization rather than in decision-making. While the system was able to identify general trends and sentiments, it lacked the ability to reliably satisfy simultaneous requirements for semantic correctness and structural determinism, which are essential for autonomous financial recommendations. This gap is not a consequence of poor implementation but rather reflects fundamental characteristics of current AI models.

The findings therefore suggest that AI should be understood as a supporting analytical tool rather than a substitute for human judgment. This distinction is critical, particularly in domains such as finance where decisions can have significant economic consequences. (Raisch & Krakowski, 2021).

9.3 Implications for AI-Based Financial-Systems

From a technical perspective, this work highlights both the potential and the structural limitations of integrating Large Language Models into financial systems. The evaluation shows that AI components are highly effective in processing unstructured textual data, identifying patterns and summarizing large volumes of information that would be difficult to analyze manually. These capabilities make LLMs well suited for supporting exploratory financial analysis and reducing information overload.

At the same time, the findings demonstrate that current language models operate probabilistically and do not provide guarantees of deterministic behaviour. In applications where outputs must be reproducible, verifiable and complete, this characteristic imposes significant constraints. The observed inability to reliably satisfy simultaneous requirements for semantic correctness and strict output structure represents a critical limitation for automated financial pipelines.

Attempts to compensate for this behaviour through stricter prompt design, validation logic or recursive correction mechanisms increase system complexity but do not fundamentally resolve the underlying issue. As a result, AI-generated outputs should remain clearly separated from deterministic system components and be framed as advisory rather than authoritative. Future financial systems should therefore emphasize transparency, traceability and user control, combining AI-based interpretation with traditional rule-based validation and domain knowledge.

9.4 Reflection on the Abandoned Advisory Component

One of the most significant outcomes of this thesis emerged from the advisory-oriented implementation path, which was intentionally limited after systematic evaluation revealed fundamental reliability and accountability issues. Although initially conceived as a core component, its development demonstrated limitations that could not be responsibly resolved within the scope of the project.

The central challenge was not merely inconsistent output, but the inability of the language model to reliably fulfill simultaneous constraints on content correctness and output structure. In practice, outputs that adhered to strict structural requirements frequently exhibited semantic errors, while semantically plausible results often violated structural constraints. Efforts to enforce convergence through validation and correction loops did not yield stable results, but instead led to oscillating behaviour between these two failure modes.

In the context of financial advisory systems, where correctness, completeness and reproducibility are non-negotiable requirements, this limitation represents an unacceptable risk. The decision to abandon the implementation of a fully automated advisory component therefore does not reflect incomplete development, but rather constitutes a key research finding. It underscores a fundamental mismatch between probabilistic language generation and the deterministic guarantees required for responsible financial decision-making.

9.5 Scientific and Practical Contribution

Despite these limitations, the thesis makes several relevant contributions to both research and practice. It provides a detailed, technically grounded examination of how AI can be integrated into financial analysis workflows. It also offers empirical evidence that challenges overly optimistic assumptions about fully automated financial advisors.

From a scientific perspective, the work contributes to the growing body of literature that emphasizes the importance of human-AI collaboration rather than substitution. From a practical perspective, it demonstrates how AI can be used responsibly to support, but not replace, financial decision-making.

9.6 Summary of Discussion

In summary, the discussion demonstrates that the central research question of this thesis can be answered in a differentiated manner. Artificial intelligence and Large Language Models in particular, are well suited for supporting financial analysis through data aggregation, interpretation and contextualization of complex information.

However, the results also show that current AI models cannot reliably meet the combined requirements of semantic correctness, structural determinism and accountability necessary for autonomous financial decision-making. This limitation is not the result of implementation shortcomings, but reflects fundamental characteristics of probabilistic language models.

The findings therefore emphasize the importance of cautious integration, clearly defined system boundaries and sustained human oversight. Rather than replacing human judgment, AI should be applied as a transparent and controlled decision-support tool. (Raisch & Krakowski, 2021).
These insights form a robust foundation for the concluding chapter, which summarizes the overall contributions of this work and outlines directions for future research and development.

10 Conclusion and Outlook

The findings of this work ultimately allow a comprehensive assessment of the role artificial intelligence can realistically assume in financial decision-support systems. Building upon the conceptual design, technical implementation and critical evaluation presented in the previous chapters, the following section consolidates the main insights, reflects on their broader implications and outlines potential directions for future development.

10.1 Summary of the Thesis

The objective of this thesis was to investigate the feasibility, limitations and responsible application of artificial intelligence, particularly Large Language Models, in the context of financial analysis and investment-related decision support. Motivated by increasing complexity in financial markets and the growing availability of AI technologies, the work aimed to move beyond theoretical discussion by designing, implementing and critically evaluating a prototype system under realistic technical constraints.

The project followed a dual-path approach. The first path focused on the development of a financial search and analysis engine that integrates structured market data with AI-supported news interpretation to support exploratory analysis. The second path explored the feasibility of an AI-based advisory component intended to generate structured, personalized financial insights.

Rather than pursuing full automation, the thesis deliberately adopted a decision-support perspective. This allowed for a differentiated evaluation of what can be reliably achieved with current AI technologies and where fundamental limitations emerge. As a result, the primary contribution of this work lies not in delivering a market-ready system, but in empirically identifying the boundaries between feasible assistance and unreliable automation in financial contexts.

10.2 Answering the Research Question

The central research question of this thesis was:

To what extent are modern AI models suitable for supporting or automating financial analysis and investment-related decision-making and which limitations arise in practice?

Based on the implementation and evaluation results, this question can be answered in a differentiated manner.

Modern AI models, particularly Large Language Models, are well suited for supporting financial analysis. They are highly effective at aggregating large volumes of information, interpreting unstructured textual data and providing contextual summaries that reduce cognitive load for users. Within the scope of exploratory analysis, these capabilities offer substantial practical value.

However, the results also demonstrate that current AI models are not suitable for autonomous or fully reliable financial decision-making. Their probabilistic nature prevents them from consistently satisfying simultaneous requirements for semantic correctness, structural determinism and reproducibility. In financial applications, where accountability and precision are essential, this limitation represents a decisive barrier to full automation. (Bommasani et al., 2022).

Accordingly, AI should currently be understood as an assistive technology that augments human judgment rather than replaces it. (Shneiderman, 2020).

10.3 Key Findings and Contributions

First, the thesis demonstrates that AI-supported financial analysis is technically feasible and practically valuable when applied in a controlled and transparent manner. The implemented search and analysis engine illustrates how structured market data and unstructured financial news can be integrated into a coherent decision-support framework.

Second, the implementation and evaluation of Large Language Models revealed a fundamental trade-off between analytical flexibility and operational reliability. The results show that LLMs are capable of producing contextually meaningful interpretations and summaries of financial information. However, their probabilistic output behaviour limits reproducibility and deterministic integration within automated evaluation pipelines. This limitation persisted despite extensive prompt engineering, validation logic and output correction mechanisms, highlighting current technical constraints in production-oriented financial applications.

Third, the thesis emphasizes the importance of architectural and conceptual system design. The transition from an initially AI-centered qualitative evaluation approach to a hybrid deterministic scoring architecture represents a key contribution of this work. This design preserves the analytical benefits of qualitative information processing while ensuring reproducibility, interpretability and system stability. The decision to maintain AI as a supportive analytical component rather than an autonomous advisory mechanism proved to be both technically and ethically justified.

Finally, the thesis provides practical insight into challenges associated with real-world financial data processing, including issues of data quality, availability, standardization and dependency on external providers. These findings underline the importance of robust data handling strategies in the development of financial decision-support systems.

10.4 Limitations of the Work

Despite its contributions, this thesis is subject to several limitations.

The system relies on publicly available data sources, which restricts both accuracy and completeness. Access to professional-grade financial data could improve analytical depth but was beyond the scope of this project.

In addition, the evaluation was conducted without a formal user study. While technical behaviour and system reliability were examined in detail, user experience, trust and decision-making effectiveness could not be empirically assessed.

Furthermore, the language model used was a general-purpose Large Language Model rather than a finance-specific variant. While this choice allowed for a realistic assessment of widely available AI technology, it also limited domain-specific precision.

These limitations underline the exploratory nature of the work and suggest that the findings should be interpreted as indicative rather than definitive.

10.5 Outlook and Future Research Directions

While the identified limitations currently constrain the use of AI in fully automated financial advisory systems, they should not be interpreted as permanent barriers. The rapid pace of development in artificial intelligence, driven by advances in model architectures, training methods and alignment techniques, suggests that some of the observed challenges may be addressed in the foreseeable future.

Ongoing research into more controllable and deterministic language models, as well as hybrid approaches that combine probabilistic AI with symbolic reasoning and rule-based validation, represents a particularly promising direction. In addition, domain-specific models trained explicitly on financial data may improve semantic accuracy and reduce ambiguity in analytical outputs.

Beyond model development, future work could also explore alternative system architectures for data acquisition. During the conceptual phase of this project, a continuous crawling approach was considered in which asset-related information would be collected and updated persistently in the background. Such an architecture could improve scalability and reduce dependency on external data providers, but was not implemented due to resource constraints related to computation, storage and maintenance.

As AI technologies mature, it is reasonable to expect that AI-based advisory systems will become more prevalent. However, their adoption will depend not only on technical feasibility, but also on regulatory frameworks, transparency requirements and societal trust. Future research should therefore focus on combining technological progress with explainability, accountability and human oversight to ensure responsible integration of AI into financial decision-making.

10.6 Final Remarks

This thesis demonstrates that while artificial intelligence has reached a remarkable level of technical capability, its application in financial decision-making must be approached with realism and responsibility. The results show that AI excels at supporting analysis and interpretation but currently falls short when strict guarantees of correctness and determinism are required.

Rather than viewing this as a limitation of the approach, the findings emphasize the importance of aligning system design with the actual capabilities of the underlying technology. In this sense, the developed prototype serves both as a functional system and as an empirical case study of responsible AI integration.

As AI technologies continue to evolve at a rapid pace, their role in finance will undoubtedly expand. The central challenge for future systems will not be automation alone, but the careful balance between technological potential, transparency and human control.

List of Diagrams

Diagram 1: Distribution of Employment Status (Multiselection possible) 8
Diagram 2: Experience with Investing 8
Diagram 3: Interest in Personal Investing 8
Diagram 4: Risk Awareness (1-10) 9
Diagram 5: Personal Assessment 9

List of Tables

Table 1: Survey Questions Overview 32

List of Figures

Figure 1: Landing Page of Search Engine 27
Figure 2: Search Function on Landing Page 27
Figure 3: Loading State - Search 27
Figure 4: Loading State - News 27
Figure 5: Loading State - Forwarding 27
Figure 6: Asset Overview Page - General Information 30
Figure 7: Asset Overview Page - Assessment 28
Figure 8: Asset Overview Page - News 31
Figure 9: Survey Questions (Multiple Choice) 32
Figure 10: Calculation Waiting Screen 33
Figure 11: Dashboard - Overview Page 33
Figure 12: Dashboard - Risk Assessment Page 34

References

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., Center for Research on Foundation Models, Stanford Institute for Human-Centered Artificial Intelligence, Stanford University & Liang, P. (2022). On the Opportunities and Risks of Foundation Models. In Center For Research On Foundation Models (S. 1 - 2). https://arxiv.org/pdf/2108.07258.pdf

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

Fischer, T. & Krauss, C. (2017). Deep learning with long short-term memory networks for financial market predictions. European Journal Of Operational Research, 654–669. https://doi.org/10.1016/j.ejor.2017.11.054

Heaton, J. B., Polson, N. G. & Witte, J. H. (2016). Deep learning for finance: deep portfolios. Applied Stochastic Models in Business And Industry, 3–12. https://doi.org/10.1002/asmb.2209

Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. Business Horizons, 577–-586. https://doi.org/10.1016/j.bushor.2018.03.007

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y. & Hong Kong University of Science and Technology. (2023). Survey of Hallucination in Natural Language Generation. In ACM Computing Surveys (Bd. 55, Nummer 12, S. 38). https://doi.org/10.1145/3571730

Jung, D., Dorner, V., Glaser, F. & Morana, S. (2018). CATCHWORD Robo-Advisory Digitalization and Automation of Financial Advisory. In Bus Inf Syst Eng (S. 81–86). https://doi.org/10.1007/s12599-018-0521-9

Kearney, C. & Liu, S. (2013). Textual Sentiment Analysis in Finance: A Survey of Methods and Models. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.2213801

Merkle, C. & Weber, M. (2014). Do investors put their money where their mouth is? Stock market expectations and investing behaviour. Journal Of Banking & Finance, 372–386. https://doi.org/10.1016/j.jbankfin.2014.03.042

Raisch, S. & Krakowski, S. (2021). Artificial Intelligence and Management: The Automation-Augmentation Paradox. Academy Of Management Review, 192 - 210. https://doi.org/10.5465/amr.2018.0072

Shneiderman, B. (2020). Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy. International Journal Of Human-Computer Interaction, 495 - 504. https://doi.org/10.1080/10447318.2020.1741118

Tetlock, P. C. (2007). Giving Content to Investor Sentiment: The Role of Media in the Stock Market. The Journal Of Finance, 1139–1168. https://doi.org/10.1111/j.1540-6261.2007.01232.x

Wirtschaftsuniversität Wien
Vienna University of Economics and Business

List of aids for seminar paper/thesis

Title of thesis: **FINWIKI**
Author(s): **KEISLER Raphael, H12309392**

Aids/tools used	Type(s) of use	Affected areas/chapters
ChatGPT	**Rephrasing, Restructuring and linguistic Refinement**	**Whole Thesis**
Gemini	**Rephrasing, Restructuring and linguistic Refinement**	**Whole Thesis**
ChatGPT	**Code Generation; Explanation of technical Concepts; Optimization of existing Code; Bug Handling**	**Practical Part**
Gemini	**Code Generation; Explanation of technical Concepts; Optimization of existing Code; Bug Handling**	**Practical Part**
Ollama	**Used in Software**	**Practical Part**

I hereby declare that I have listed all the aids I have used in the list above. If no aids have been used, it is also indicated in the list (to be listed under “Aids/tools used: none”).



09.02.2026

Date

Signature(s)