MASTERARBEIT | MASTER'S THESIS

Titel | Title

Integration of Environmental, Social and Governance Criteria into Portfolio Optimization for Medium-Sized Enterprises

verfasst von | submitted by

Alisa Proshkina BSc

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2025

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 915
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Betriebswirtschaft
Betreut von | Supervisor:	Mag. Dr. Raimund Kovacevic

Abstract

Diese Masterarbeit untersucht die Entwicklung eines praxisnahen Portfolio-Optimierungsmodells für kleine und mittlere Unternehmen (KMU), das sowohl finanzielle Zielsetzungen als auch Environmental, Social and Governance (ESG) Kriterien berücksichtigt.

Angesichts der wachsenden Bedeutung nachhaltiger Investitionen analysiert die Studie, wie ESG-Faktoren das Risiko, die Rendite und die Gesamtperformance eines Portfolios beeinflussen. Hierzu werden verschiedene modellbasierte Ansätze mit und ohne ESG-Integration entwickelt und miteinander verglichen. ESG-Daten von drei anerkannten Anbietern werden zu einem einzigen, gewichteten Komposit-Score zusammengeführt. Die Modelle werden anhand eines rollierenden Backtests mit realen Marktdaten der XETRA-Börse evaluiert.

Die Ergebnisse zeigen, dass ESG-Vorgaben zwar die Portfoliozusammensetzung beeinflussen, jedoch nicht zwangsläufig zu negativen Performanceeffekten führen. Diese Arbeit leistet somit einen Beitrag zur Integration von ESG-Aspekten in die Finanzstrategien von KMU – auch unter eingeschränkten Ressourcenbedingungen.

Abstract

This master's thesis investigates the development of a practical portfolio optimization model tailored for small and medium-sized enterprises (SMEs), incorporating both traditional financial objectives and Environmental, Social, and Governance (ESG) criteria.

Considering the increasing focus on sustainable investing, the study examines how ESG factors influence portfolio risk, return, and overall performance. To achieve this, multiple model-based approaches are implemented and compared, with and without ESG integration. ESG data from three established providers is aggregated into a single, weighted composite score. Model performance is assessed through a rolling backtest using real-world data from the XETRA exchange.

The findings indicate that while ESG constraints influence portfolio structure, they do not inherently compromise financial performance. This research offers practical insights into embedding ESG considerations within SME investment strategies, particularly under conditions of limited resources.

Table of Contents

List of Abbreviations

CAPM: Capital Asset Pricing Model

DAX: Deutscher Aktienindex (German Stock Index)

DEA: Data Envelopment Analysis

ESG: Environmental, Social, and Governance

EWMA: Exponentially Weighted Moving Average

GBM: Geometric Brownian Motion

MAE: Mean Absolute Error

MPT: Modern Portfolio Theory

OOS: Out-of-Sample

SME: Small and Medium-sized Enterprise

LSEG: London Stock Exchange Group

VDA: Verbal Decision Anaysis

List of Figures

List of Tables

1. Introduction

ESG criteria are now commonly used to assess a company's financial potential as well as its wider impact on society and the environment. Instead of focusing only on profits, ESG factors help investors look at how responsibly a company operates, how it treats people, and how it is managed.

The growing interest in ESG is supported by many studies. Whelan, Atz, Van Holt and Clark (2021) demonstrate that ESG investing has shifted from a niche to a mainstream strategy, with evidence from over 1,000 studies showing links to increased financial resilience, risk reduction, and innovation. Pedersen, Fitzgibbons and Pomorski (2021) furthermore note that ESG considerations increasingly take the form of how investors build and manage portfolios, reflecting a wider redefinition of value in financial decisions.

This transformation has expanded to include quantitative models. While the foundation of modern portfolio theory remains grounded in the mean-variance approach introduced by Markowitz (1952), new models increasingly incorporate ESG objectives. Adding sustainability as a third dimension, along with risk and return, reflects the changing priorities of many investors. Pedersen et al. (2021) argue that such integration can improve risk-adjusted performance by aligning portfolios with evolving regulatory standards and stakeholders' expectations. Credible ESG practices for businesses currently significantly influence investor perception and capital access (Whelan et al., 2021).

Yet the evolution of ESG investing has been uneven across different market segments. While institutional investors and large corporations have been the primary focus of academic research and professional development, smaller businesses face unique challenges. As Shalhoob and Hussainey (2023) point out, SMEs often struggle with limited resources and infrastructure, making it difficult for them to implement comprehensive ESG strategies. The absence of standardized frameworks tailored to smaller firms exacerbates this capacity gap. The OECD (2020) suggests that unless ESG strategies are adapted to account for these constraints, many SMEs will remain excluded from the shift toward more sustainable finance.

This thesis addresses this gap by developing and evaluating ESG-integrated portfolio optimization models specifically tailored to the needs and limitations of SMEs. While most academic models assume access to advanced data and computing resources, this work focuses on transparent, accessible, and low-complexity frameworks that SMEs can realistically implement. ESG integration is operationalized via constraint-based methods rather than embedded utility functions, enabling practical alignment with sustainability goals using static ESG ratings and publicly available data.

This thesis explores the integration of ESG criteria into portfolio optimization models, specifically tailored for SMEs. Unlike previous studies that primarily focused on institutional investors, this research addresses the often-neglected SME perspective. In doing so, the thesis contributes to the current debate in two ways. First, it offers a systematic comparison of multiple portfolio optimization approaches, each tested with and without ESG constraints. Second, it introduces realistic performance benchmarking for SMEs by focusing on model interpretability, forecast accuracy, turnover, and exposure to market risk. The thesis demonstrates that including ESG factors can be helpful even in basic models and with limited resources, which is different from most existing studies that look at large institutions and complex frameworks.

The research is guided by the following research questions:

1. Which portfolio optimization approach offers the most transparent and implementable framework for SMEs with limited analytical and data resources?
2. How do different return estimation methods, such as historical averages and market-based models, affect the reliability of portfolio performance in out-of-sample testing?
3. In what ways do static versus dynamic risk estimators influence the interpretability and robustness of ESG-integrated portfolios?
4. How does a One-Fund strategy compare to minimum-risk portfolios under ESG constraints, and does it offer meaningful advantages for investors with varying levels of risk tolerance?

By addressing these questions, the thesis aims to support smaller firms in navigating the shift toward sustainable investing using tools that match their operational realities. The empirical results highlight that ESG integration does not necessarily reduce portfolio performance and that simpler models, when properly configured, can offer both financial and sustainability benefits.

The following chapters are organized as follows: Chapter 2 presents a literature review on ESG investing and portfolio optimization, with particular attention to the challenges faced by SMEs. Chapter 3 outlines the methodology, including the selected portfolio models, performance metrics, and the implementation of ESG constraints. Chapter 4 presents the results of the backtesting analysis across different strategies and models, while Chapter 5 discusses the findings in relation to the research questions, practical applications, and existing literature. The concluding chapter summarizes the key insights, contributions, and potential areas for future research in this field.

2. Literature Review

This literature review explores the incorporation of ESG factors into portfolio optimization models, with a particular emphasis on this applicability to SMEs. By analysing theoretical foundations, prevalent assumptions, and current methodological limitations, the review seeks to develop a more inclusive and tailored investment framework that better addresses the unique needs of SMEs. This research aims to connect ESG-driven investment strategies with the real challenges smaller businesses face, helping to create a better understanding of sustainable finance in different economic situations.

2.1. ESG Integration as a Structural Shift in Investment Strategy

ESG investing has shifted from being mainly about ethics to becoming an important part of financial planning. Today, many investors and companies see that including ESG factors in their decisions helps manage long-term risks and supports financial stability. This shift in perspective has been influenced by evolving investor priorities, regulatory developments, and an expanding stream of studies indicating that firms with higher ESG standards often demonstrate improved financial resilience and competitiveness (Whelan et al., 2021). Once regarded as a secondary concern, ESG is now widely viewed as a fundamental aspect of corporate resilience and strategic planning (Eccles, Ioannou, & Serafeim, 2014).

The evolution of corporate responsibility can be traced back to earlier models of Corporate Social Responsibility (CSR), which initially focused on charitable giving and voluntary ethical commitments, according to the research of Zhang and Hao (2024). As time passed, CSR transformed into a more strategic approach, driven by increasing stakeholder demands for transparency and accountability in various areas, including climate risk, labour rights, and corporate governance. According to Zhang and Hao (2024), ESG factors have a much greater impact on long-term profitability and organizational adaptability when they are embedded into a company's core business model, rather than being treated as isolated or symbolic initiatives.

Empirical research supports this direction. A meta-analysis by Whelan et al. (2021), covering more than 1,000 studies, identifies a consistent association between ESG integration and enhanced financial outcomes, especially among companies addressing material sustainability topics. Khan, Serafim, Yoon and Suh (2016) research challenges the common belief that sustainability and profitability are at odds. Their study indicates that companies prioritizing strong ESG practices tend to excel in areas such as innovation, risk management, and access to capital. This trend has caught the attention of policymakers, who are beginning to see ESG integration as an essential component of a healthy and stable financial system, rather than just a voluntary initiative.

The European Union has become a global leader in formalizing ESG standards, introducing regulations that aim to align capital markets with sustainability objectives. The EU Taxonomy (Regulation (EU) 2020/852) and the Corporate Sustainability Reporting Directive (CSRD) set straightforward rules for how companies should report on sustainability and how investments should be categorized (European Parliament & Council, 2020; European Commission, 2023).

Regulation (EU) 2020/852, also known as the EU Taxonomy, establishes a framework to classify environmentally sustainable economic activities and seeks to channel capital toward investments that contribute meaningfully to environmental goals (European Parliament & Council, 2020, Art. 3). The taxonomy establishes clear guidelines for assessing sustainability, which creates a uniform standard for incorporating ESG factors into financial decision-making across various markets (European Parliament & Council, 2020, Arts. 9-11).

In addition, the Corporate Sustainability Reporting Directive (CSRD) builds upon the Non-Financial Reporting Directive by introducing more comprehensive and specific disclosure obligations. Under the CSRD, companies must disclose how ESG factors are integrated into their operations, including risks and impacts. These reports are subject to standardized formats to improve transparency, reliability, and comparability across firms and sectors (European Commission, 2023; European Parliament & Council, 2020, Recitals 12-14).

Recent regulatory changes aimed at improving consistency and transparency in ESG reporting have sparked debate, especially among SMEs. The implementation of ESG standards often comes with significant financial and administrative burdens. Hancock (2024) points out that SMEs, unlike larger firms, typically do not have the same level of in-house expertise or financial resources, making it challenging for them to fulfil these new requirements without experiencing considerable pressure. In addition to that, the continued lack of uniformity between various national and international ESG frameworks adds to the overall uncertainty in the business environment (European Parliament & Council, 2020, Recital 14).

ESG factors are increasingly influencing financial decision-making due to evolving regulations, particularly in the European Union. As these rules take shape, companies of all sizes are expected to show concrete ESG results. This trend highlights the importance of developing flexible, efficient methods for incorporating ESG principles into business practices, especially for smaller firms with limited resources. The growing focus on ESG has sparked interest among both academics and investors in understanding how ESG performance relates to financial success. As the regulatory landscape continues to change, businesses will need to adapt their strategies to meet new sustainability requirements while maintaining profitability.

2.2. ESG and Financial Performance

The evidence linking ESG practices to financial performance remains inconclusive, with research showing a wide range of results. While some studies indicate that implementing sustainable business practices can lead to long-term stability, others argue that the expenses associated with ESG compliance might exceed short-term financial benefits. However, a significant portion of recent research suggests that companies focusing on material ESG issues often experience enhanced resilience and lower risk exposure (Whelan et al., 2021).

Recent research has delivered significant insights on the relationship between ESG practices and financial performance. A comprehensive meta-analysis by Whelan et al. (2021) examined over 1,000 empirical papers revealed that ESG adoption is associated with improved financial performance in most cases. Khan et al. (2016) further support this finding, demonstrating that companies focusing on material ESG issues often see enhanced stock valuation and operational efficiency. However, the financial benefits of ESG implementation may not be immediate. Jung and Yoo (2023) describe a U-shaped effect, where companies initially face higher costs during ESG implementation before realizing performance benefits. Ahmad, Yaqub and Lee (2024) expand on this concept, noting that ESG disclosures contribute to business sustainability and innovation, which can strengthen financial outcomes over time.

ESG practices are becoming increasingly relevant for managing business risks. Research by Eccles, Ioannou and Serafim (2014) showed that companies with well-developed ESG strategies often benefit from lower costs of capital and better financial resilience during market turbulence. Similarly, Whelan et al. (2021) further supported this, noting that ESG integration has helped companies as a buffer during financial crises. In sector-specific research, Wanday and Ajour El Zein (2022) found that ESG-oriented investment approaches in the energy industry not only produced higher returns but also helped investors navigate complex regulatory and reputational challenges. Drei et al. (2019) further observed that ESG factors are influencing asset valuation in the Eurozone, where they are increasingly viewed as central to financial decision-making processes.

The financial consequences of incorporating ESG principles vary widely depending on the situation. Recent research by Patak, Tewari and Drempetic (2024) suggests that companies facing high compliance costs may experience reduced financial flexibility when ESG policies become mandatory. Son and Suh's (2024) work emphasizes the challenges businesses face in trying to meet ESG requirements while containing their competitive positioning. The effectiveness of ESG disclosures, according to Ahmad, Yaqub and Lee (2024), is largely influenced by factors specific to each company, including their governance structures, what investors expect from them, and the norms within their particular industry sector.

The growing adoption of ESG principles in investment strategies stems from their perceived ability to improve financial performance, making them increasingly attractive to both institutional and individual investors. However, the effectiveness of ESG implementation depends on the specific approach and contextual factors. This raises a critical question for investors: how can ESG factors be systematically incorporated into portfolio construction? Addressing this issue requires a closer examination of portfolio optimization fundamentals and an exploration of how these models can be adapted to incorporate ESG objectives.

Institutional investors have been the primary target audience for many existing investment decision-making tools. To increase clarity, use active voice: This observation underscores the necessity of evaluating potential adjustments for these models in resource-constrained environments like SMEs. To facilitate this, this research reviews the core concepts of portfolio optimization theory, with an emphasis on models that seek to balance return and risk. In view of the typical constraints faced by SMEs, the goal is not to replicate the complexity of institutional frameworks but to identify theoretical elements that can support simplified, ESG-conscious investment decisions. A more detailed examination and the practical implementation of such models form the core of this thesis, while the theoretical background is outlined in the methodology chapter.

2.3. Trade-offs and Constraints in ESG Portfolio Optimization

Incorporating ESG criteria into investment strategies presents challenges when balancing sustainability with financial performance. The most noticeable impact of ESG screening is that it reduces the number of available investment options, which in turn narrows the overall investment universe. This process often leads to uneven sector representation, with tech and healthcare companies typically overrepresented, while energy and industrial firms are underrepresented. These effects are particularly evident in developing markets, where fewer companies meet high ESG standards (Abate, Basile, & Ferrari, 2024; Chen, Zhang, Huang, Xiao, & Zhou, 2021; Spigeleer, Höcht, Jakubowski, Reyners, & Schoutens, 2023).

The inconsistent ESG ratings across different agencies make it especially difficult for investors. Each rating provider employs unique methodologies and scoring systems, resulting in different ESG scores for the same company. This discrepancy creates conflicting investment signals, making it difficult for investors to construct portfolios effectively. Relying on mismatched data may lead to poor investment decisions (Berg, Kölbel, & Rigobon, 2022; Spigeleeret al., 2023). To address these issues, researchers have proposed alternative approaches, such as Data Envelopment Analysis and semivariance-based models, which aim to balance ESG considerations with portfolio diversification (Chen et al., 2021). Also, since ESG constraints change the efficient frontier, investors need to use multi-objective optimization techniques to keep their performance strong while including non-financial goals (Abate et al., 2024).

From a risk perspective, ESG integration has both advantages and drawbacks. On one hand, portfolios that focus on ESG principles usually show better stability during market ups and downs, as companies with high ESG ratings often have stronger finances and are less likely to face reputation or regulatory issues (Abate et al., 2024; Chen et al., 2021). However, the regular adjustments needed for ESG-focused portfolios can result in higher trading costs, which might reduce some of the benefits of lowering risk (Abate et al., 2024). While ESG portfolios offer competitive risk-return profiles, their performance is heavily influenced by the specific ESG rating methodologies employed and the prevailing economic climate. Portfolios designed to minimize carbon footprint do not necessarily suffer from increased volatility and often show similar drawdown patterns to conventional investment approaches (Spigeleer et al., 2023). These nuanced risk dynamics highlight the complexity of ESG integration in portfolio management.

Researchers have proposed several changes to traditional investment strategies to better incorporate ESG factors. Robust optimization frameworks, for instance, are designed to incorporate ESG preferences while accounting for uncertainty in input parameters (Escobar-Anel & Jiao, 2024). Similarly, adjustments to the traditional Markowitz framework demonstrate that ESG factors can be integrated without necessarily sacrificing financial performance (Zhou, 2024). However, the overall effectiveness of ESG integration still depends on external factors such as rating consistency, market trends, and the investor's specific objectives and constraints (Spigeleer et al., 2023).

Despite these advancements, many of these innovative approaches cater to large institutional investors and corporations. This leaves an important question unanswered: How can smaller businesses and organizations with limited resources participate in ESG investing in a way that is both practical and impactful?

2.4. Small Firms and ESG: Gaps, Challenges, and Emerging Approaches

The previous chapters highlighted a noticeable gap in understanding the application of ESG concepts to smaller businesses. The current ESG frameworks cater to publicly traded companies, often relying on standardized reporting, shareholder engagement, and significant capital movements. SMEs, on another hand, face unique challenges due to their distinct financial structures, regulatory environments, and organizational setups (MacNeil & Esser, 2022; Wagemans, van Koppen, & Mol, 2018).

Institutional investors and SMEs approach ESG considerations differently. Large institutions often focus on ESG to mitigate risks and comply with regulations or maintain their public image. In contrast, smaller businesses typically integrate ESG practices because they align with the company's fundamental values or long-term strategic goals (Luke, 2022).

Institutional investors typically engage with ESG issues through shareholder activism, but this approach does not apply to SMEs that are not publicly traded (Wagemans et al., 2018). Consequently, smaller companies seldom take part in formal ESG evaluation processes. Instead, they often implement sustainability measures directly, such as investing in renewable energy sources, participating in carbon offset initiatives, or developing ethical supply chain practices (Luke, 2022). These approaches reflect ESG principles but do not always fit within conventional ESG assessment models (Garrido-Ruso, Otero-Gonzalez, & Lopez-Penabad, 2024).

Another persistent challenge for SMEs is the difficulty in linking ESG investments to measurable financial returns. As Garrido-Ruso et al. (2024) observe, that while adopting sustainable practices may strengthen a company's long-term resilience, smaller businesses often find it challenging to see quick financial benefits. These obstacles can make ESG initiatives less appealing compared to traditional investment. MacNeil and Esser (2022) and Luke (2022) point out that SMEs generally do not have the same technical expertise or financial resources as larger corporations to effectively measure and report their ESG performance.

One of the key issues SMEs encounter is the availability of suitable financial instruments that align with their ESG goals. Many existing ESG investment products are designed with larger companies in mind, potentially leaving smaller firms at a disadvantage. The following section will describe the current landscape of ESG investment products and evaluate how well they meet the specific needs and constraints of SMEs.

2.4.1. Limited Accessibility of ESG Instruments for SME Investors

SMEs are becoming more interested in ESG investing, but they still face serious challenges in accessing the right financial tools. Many ESG-related investment products, such as green bonds, sustainability-linked funds, and ESG indices, are mainly created for large institutional investors and big public companies. These products often require detailed reports, high levels of disclosure, or large financial commitments, which are difficult for SMEs to manage as either investors or issuers (Chien et al., 2021; Orzechowski & Bombol, 2022). Because of this, SMEs looking to invest responsibly find themselves limited to options that don’t fit their size or capabilities.

The way ESG investment markets are built often leaves SMEs out. For example, green bonds are usually issued by large companies with strong credit ratings and the ability to complete complex verification steps (Orzechowski & Bombol, 2022). This makes it hard for smaller investors to use them for low-scale sustainable investing. Likewise, ESG funds tend to focus on companies that already have strong ESG scores and advanced reporting systems, leaving out many firms that SMEs might prefer because of local or industry relevance (Chien et al., 2021; Zeidan, 2022).

Problems with data and transparency also make things more complicated. ESG benchmarks and indices mostly rely on public data from big, listed companies, which means SMEs are often left out of the ESG data landscape (MacNeil & Esser, 2022).

Without access to specific, relevant information, SME investors find it hard to evaluate the ESG performance of smaller or local firms. This forces them to work with incomplete or inconsistent data, limiting their ability to build well-informed, ESG-focused portfolios.

Another issue is that many SMEs lack the knowledge and support needed to properly assess ESG risks and benefits from an investment perspective. As noted by Chien et al. (2021), the lack of clear policies and the existence of fragmented ESG standards make things confusing and discourage participation. Many SMEs are not fully aware of the advantages of using ESG principles in investment and often cannot find ESG investment options that are simple and affordable enough to use.

Overall, the ESG investment system still mainly serves large financial institutions. It focuses on scale, easy trading (liquidity), and standardized reporting, which often leaves SMEs unable to take part, whether as issuers or investors (MacNeil & Esser, 2022; Garrido-Ruso et al., 2024). Although some SMEs are starting to explore thematic ESG funds or easier screening tools, their ability to join in on sustainable investing largely depends on future policy changes and the creation of ESG systems designed for smaller businesses (Wagemans et al., 2018).

2.4.2. Limitations of Conventional ESG Models for SME Investors

Many mainstream ESG portfolio models do not reflect the real conditions and limitations of SMEs. Chien et al. (2021) explain that SMEs face disadvantages because they lack the economies of scale that large corporations benefit from. Institutional investors often receive volume discounts, have better access to financial markets, and operate within structures that support more efficient and lower-cost ESG investing. As a result, most ESG investment products and portfolio models are built around the needs of large, publicly listed companies. This leaves smaller investors with fewer practical tools for sustainable investing.

This issue goes beyond the availability of ESG data, it also relates to the way ESG portfolio models are designed. Standard ESG strategies usually assume that investors can meet certain performance benchmarks and diversify across many sectors and asset types. However, SMEs often have limited financial resources and smaller sets of investment options. These conditions make it difficult to apply conventional ESG models without facing high concentration risk (Miranda-García & Segovia-Vargas, 2024).

Advanced ESG investment methods have been proposed in academic research. For example, some studies suggest using Monte Carlo simulations or scenario-based optimizations (Hibiki, 2006). These methods require significant data access, computational resources, and long-term capital planning. Such tools are usually suitable for large institutions rather than smaller firms. Other approaches, like Verbal Decision Analysis, depend on expert guidance and detailed ESG ratings (Silva, Pinheiro, & Poggi, 2017). These features are often inaccessible or unaffordable for SMEs. From the perspective of a smaller investor, the complexity of these tools and the lack of customized support make it difficult to set clear ESG goals or build a portfolio that reflects specific sustainability values.

The gap between ESG model assumptions and SME realities is not only a technical matter. It also reflects structural challenges. According to Martinez-Cillero et al. (2020), SMEs tend to rely heavily on internal financing, face limited access to alternative capital sources, and are more sensitive to investment risks. These conditions reduce the likelihood that SMEs will participate in investment models that involve high initial costs or delayed financial returns. Although the study focuses on general investment behaviour, the same factors affect how SMEs engage with ESG investing.

In conclusion, many ESG portfolio frameworks are based on assumptions about stable funding, access to high-quality ESG data, and broad diversification. These assumptions do not match the financial and operational conditions of most SMEs. To make ESG investing more accessible and relevant to this group, portfolio models should be simplified. They should require less data, allow for smaller investments, and include flexible criteria that reflect the real-world behaviour of SME investors. These changes would not only increase fairness but also encourage wider participation in sustainable investment.

2.4.3. Simplified Optimization Strategies for ESG Integration

This section focuses specifically on how ESG optimization methods used by institutional investors can be adapted for SMEs acting as investors, not ESG implementers.

Examining how institutions have incorporated ESG factors into their models provides valuable insights for adapting these frameworks to other contexts. While the focus has been on large-scale institutional applications, the principles and approaches used can serve as a foundation for broader implementation across different investment strategies and scales.

For instance, Abate et al. (2024) expand on the classic mean-variance approach by adding ESG as a third factor alongside risk and return. Their research indicates that better ESG scores can lower downside risk, and ESG-focused portfolios tend to have reduced transaction costs over time, potentially enhancing cost efficiency. Nevertheless, SMEs often struggle with limited data and computational resources, making it impractical to fully implement these complex models. As an alternative, authors emphasise that SMEs might consider using simpler ESG screening techniques, such as industry-based filters or preset sustainability weightings. This approach allows SMEs to integrate ESG criteria into their portfolio selection process while remaining within their financial and technical limits.

A different angle on simplification comes from Liu and Jing (2024), who apply cardinality constraints to ESG portfolios. Their approach restricts the number of assets selected, which helps avoid over-diversification and keeps transaction costs in check. This method could be particularly beneficial for SMEs that often face limitations in capital allocation and implementation capabilities. By focusing on a select few ESG-aligned assets, these companies may find a more practical and achievable way to incorporate sustainability into their investment strategies. This targeted approach enables SMEs to build ESG-conscious portfolios in a manageable way, without compromising financial performance or resource limits.

Recent research, including a study by Chen et al. (2021), has investigated advanced techniques like Data Envelopment Analysis (DEA) for enhancing ESG portfolio performance. DEA provides a comprehensive assessment that considers both financial returns and sustainability metrics. However, this approach demands extensive data and significant computing resources, which may be out of reach for many smaller companies. As an alternative, simpler strategies such as eliminating the bottom 20% of ESG-scored investments can be an effective compromise. These straightforward methods require minimal technological infrastructure while still allowing firms to align their investment choices with their sustainability objectives. This approach offers a practical solution for businesses looking to incorporate ESG factors into their investment decisions without the need for complex analytical tools.

Regional variations in ESG portfolio performance are evident, as highlighted by Cesarone, Martino and Carleo (2022). Their research found that ESG constraints enhanced profitability in the United States, while European markets showed no financial gains from this approach. However, portfolios incorporating ESG factors demonstrated increased resilience during market downturns in both regions. These findings indicate that SMEs might find value in adopting a flexible approach to ESG thresholds by adjusting their exposure based on market conditions, investor demands, and local regulations, rather than sticking to strict inclusion criteria.

Alternative optimization techniques, like those studied by Liagkouras, Metaxiotis, and Tsihrintzis (2022), look at different ways to optimize that use penalty functions to find a balance between financial performance and ESG objectives. This approach differs from strict exclusionary practices by gradually incorporating ESG factors into investment decisions. Instead of mandating absolute compliance, it discourages significant deviations from sustainability targets. SMEs, which typically prioritize financial stability, may benefit from using such models. By treating ESG as a flexible constraint within portfolio optimization, they can align investments with sustainability preferences while preserving competitive performance. This balanced approach allows for a more flexible and realistic integration of ESG principles into business operations.

As illustrated in the studies, institutional investors typically use sophisticated models with multiple objectives for optimization. However, SMEs require a simpler and more affordable method to integrate ESG into their investment goals. Research has shown that, while large institutions can afford and benefit from intricate financial modelling, SMEs need practical solutions that are easier to implement and maintain without sacrificing effectiveness.

2.5. Towards an SME-Compatible ESG Optimization Model

The integration of ESG factors into financial markets continues to evolve, but many existing portfolio optimization models are still primarily designed for large institutional investors. While research has shown potential benefits like improved risk-adjusted returns and long-term stability, the methods used often rely on assumptions that do not apply to smaller businesses. Many SMEs lack access to comprehensive ESG data, diverse asset pools, and advanced computational resources, making it challenging to implement these complex frameworks effectively.

To address this issue, researchers have proposed alternative approaches that maintain the core principles of portfolio optimization while adapting them for simpler environments.

Some strategies developed to reduce model complexity without sacrificing sustainability objectives include rule-based screening, limiting the number of assets, and incorporating softer ESG penalties. These modifications suggest that smaller investors can achieve meaningful ESG integration, as long as they tailor the tools to their specific needs and constraints.

A balanced approach is necessary to address the challenges faced by SMEs in implementing effective ESG strategies. This approach should combine analytical rigor with practical applicability, considering the unique constraints and requirements of SMEs. These businesses often have access to a smaller pool of assets, struggle with incomplete or inconsistent data, and are more risk-averse when it comes to financial decisions. Despite these limitations, it is extremely important to establish a system that allows SMEs to pursue sustainability objectives in a methodical and straightforward way.

The next section presents a conceptual framework that modifies traditional portfolio theory to better suit the needs of SMEs in their ESG endeavours. This adapted framework aims to provide a more accessible and relevant tool for smaller businesses looking to integrate sustainability into their investment strategies.

3. Methodology

This chapter lays the groundwork for the portfolio optimization models used in this thesis. It delves into both traditional and contemporary methods of predicting returns and assessing risk, with a particular focus on their applicability to SMEs looking to integrate ESG factors into their investment approaches. The discussion examines each technique's theoretical merits, data needs, and practical challenges in implementation. By thoroughly analysing these aspects, the chapter aims to provide a comprehensive understanding of the various tools available for portfolio optimization, enabling SMEs to make informed decisions about which methods best align with their specific goals and resources in the context of ESG-conscious investing.

This chapter's structure aligns with the thesis's main goal: creating a clear and practical optimization framework for investors with limited resources, using publicly accessible data. To achieve this, the chapter addresses several key research questions that guide the investigation and analysis. These sub-questions help break down the complex topic into manageable parts, allowing for a thorough exploration of the subject matter. By focusing on these specific areas of inquiry, the chapter aims to provide a comprehensive understanding of the optimization framework and its potential applications in real-world investment scenarios.:

1. Which portfolio optimization approach offers the most transparent and implementable framework for SMEs with limited analytical and data resources?
2. How do different return estimation methods, such as historical averages and market-based models, affect the reliability of portfolio performance in out-of-sample testing?
3. In what ways do static versus dynamic risk estimators influence the interpretability and robustness of ESG-integrated portfolios?
4. How does a One-Fund strategy compare to minimum-risk portfolios under ESG constraints, and does it offer meaningful advantages for investors with varying levels of risk tolerance?

3.1. Conceptual Foundations of ESG-Oriented Portfolio Optimization

This section introduces the main concepts underlying the portfolio models developed in this thesis. It explores the fundamental principles behind forecasting returns, evaluating risk, and integrating ESG factors into investment decisions. The discussion outlines how these elements come together within an optimization framework, creating a cohesive approach to portfolio construction. The goal is to design a practical model that can be readily implemented, with a specific emphasis on making it transparent and accessible for small and medium-sized enterprises. By breaking down complex concepts into manageable components, this chapter aims to provide a clear concept for businesses looking to incorporate sustainable investing practices into their financial strategies.

3.1.1. Portfolio Optimization Foundations: The Markowitz Model

The Markowitz Portfolio Theory, developed by Harry Markowitz in 1952, revolutionized investment strategies by offering a mathematical approach to balancing risk and potential returns. This concept emphasizes the importance of diversification in creating an optimal investment portfolio. By selecting a mix of assets with varying risk levels and expected returns, investors can potentially maximize their profits while minimizing their exposure to market volatility. The theory posits that investors can construct a portfolio optimized to deliver the highest possible returns for a given level of risk tolerance (Markowitz, 1952).

Mean-variance optimization is a widely used investment strategy that assumes investors prefer to minimize risk while maximizing returns. This method evaluates potential investments based on their expected returns and the variability of those returns, usually measured by variance or standard deviation. The model analyses a group of assets to determine the optimal combinations that form what is known as the efficient frontier. These portfolios along the frontier represent the best possible balance between risk and return at various levels, allowing investors to choose the mix that aligns with their personal risk tolerance and financial goals (Cornuejols, Pena, & Tütüncü, 2018, S. 91-93).

To build such portfolios, investors calculate three core metrics: expected return, portfolio variance, and the correlations between assets. The expected return of a portfolio is computed as:

$$E(R_p) = \sum w_i E(R_i)$$

where $E(R_p)$ is the expected return of the portfolio, w_i represents the weight of asset I in the portfolio, and $E(R_i)$ is the expected return of the asset i (Cornuejols et al., 2018, S. 94).

The total portfolio risk is measured by portfolio variance:

$$\sigma_p^2 = \sum w_i^2 \sigma_i^2 + \sum \sum w_i w_j \sigma_{ij}$$

Here, σ_p^2 is the portfolio variance, σ_i^2 is the variance of the asset i, and σ_{ij} is the covariance between assets i and j. Lower correlations between assets lead to greater diversification benefits and reduced portfolio volatility (Cornuejols et al., 2018, S. 94).

In applying the Markowitz model, certain constraints must be defined. The most basic is the budget constraint, ensuring full investment of available capital:

$$\sum w_i = 1$$

This guarantees that portfolio weights sum to 100%. Additional constraints may include non-negativity (no short selling), expressed as $w \geq 0$. and sector or asset class limits to prevent overexposure and maintain diversification (Cornuejols et al., 2018, S. 94).

Modern Portfolio Theory (MPT) has significantly shaped investment strategies, but it relies on assumptions that do not always reflect real-world market conditions. As Mangram (2013) noted, MPT presumes that all investors have equal access to perfect information, which is rarely the case in practice. The theory also relies on the notion of efficient markets, which assume that asset prices encompass all available information. However, the occurrence of financial crises and market bubbles casts doubt on this assumption, highlighting the limitations of MPT in accurately predicting market behaviour and guiding investment decisions.

Zhang, Li and Guo (2018) emphasize that market behaviour frequently departs from the ideal of efficiency. Real-world disruptions, such as unexpected policy changes or macroeconomic shocks, can lead to sudden shifts in asset prices, undermining the stability assumed by the model. The COVID-19 pandemic, for example, introduced economic volatility not captured by historical data or past asset correlations.

In addition to its market efficiency assumptions, the Markowitz model overlooks practical factors such as transaction costs and taxes, even though these elements significantly impact investment decisions (Mangram, 2013). In more dynamic market environments, where trading occurs frequently, it becomes less feasible to assume that rational investors will rebalance their portfolios continuously when each transaction incurs a cost. By neglecting these practical considerations, the model presents an oversimplified view of investor behaviour, reducing its relevance in situations where such costs are unavoidable. Furthermore, the model's reliance on the assumption that asset returns follow a multivariate normal distribution has been widely contested by empirical research. In reality, return distributions often exhibit more complex characteristics, such as fat tails, high peaks, and asymmetry, which deviate substantially from the smooth, bell-shaped curves predicted by the normal distribution (Zhang, Li, & Guo, 2018).

Traditional portfolio models face challenges when applied to more complex investment scenarios, especially those involving ESG considerations. While the Markowitz framework provides a basic approach to balancing risk and return, it assumes reliable input data, stable asset correlations, and ample diversification opportunities. SMEs often fail to meet these conditions, and the inclusion of ESG criteria further restricts their investment options. As a result, there is a growing need for simplified or adapted optimization methods that better address the constraints faced by smaller companies. These modified approaches should account for the limited resources and data available to SMEs while still allowing for effective portfolio management within ESG parameters. Developing such tailored solutions could significantly improve investment outcomes for smaller firms navigating the complexities of modern financial markets and sustainability requirements.

3.1.2. Forecast Quality Evaluation with Out-Of-Sample Testing

This thesis uses out-of-sample (OOS) testing with a rolling-window method to assess how reliable return forecasting techniques are and how practical portfolio strategies can be. The methodology replicates real-world investment scenarios, where future returns are uncertain and models require continuous updates based on the latest available information. Unlike in-sample backtesting, which uses identical data for model fitting and assessment, OOS testing distinguishes between estimation and validation stages. This separation provides a more accurate picture of a model's capacity to extrapolate to new, unseen data, offering valuable insights into its predictive power and robustness in dynamic market conditions.

The rolling-window method is a common technique used to assess model performance over time. It involves regularly updating models using a fixed amount of historical data and then testing them on subsequent periods. For example, Ban, Karoui and Lim (2018) conducted extensive rolling simulations to evaluate how well complex machine learning models could predict returns and risks across various time frames. Although their approach is computationally demanding and designed for large institutions, this thesis adapts a simpler version of the model. This modified approach is more suitable for SMEs, which typically require solutions that are easy to understand, implement, and maintain without incurring significant costs.

Guidolin, Hansen and Lozano-Bandaz (2018) highlight the effectiveness of OOS Sharpe ratios as a reliable metric for comparing predictive models. Their research, which spans nearly five decades of monthly U.S. equity portfolio data from 1968 to 2016, evaluates one-month-ahead return forecasts for various factor models. They assess performance by examining realized Sharpe ratios over rolling OOS periods. In contrast, this thesis employs a different approach, utilizing rolling quarterly evaluation windows to measure realized return, volatility, Sharpe ratio, and turnover. This method ensures consistency across all models while still capturing the ever-changing dynamics of market conditions.

Kresta and Wang's (2020) research describes the common discrepancy between in-sample and OOS model performance. They emphasize the importance of OOS testing to avoid overconfidence in model predictions, especially in volatile market conditions. This approach is particularly relevant for evaluating models intended for use in unstable or rapidly changing financial environments.

SMEs often face resource constraints that limit their ability to frequently update models or employ sophisticated techniques. In this context, a rolling OOS framework provides a practical and realistic method for assessing model performance. This approach allows for forward-looking analysis while considering the operational limitations typically faced by smaller investors, thus offering a more accurate reflection of their real-world constraints and capabilities.

The model development period from 2020 to 2024 encompasses diverse economic conditions, including the COVID-19 pandemic, subsequent recovery, geopolitical tensions, and inflation. This timeframe provides a varied backdrop for assessing model resilience. Additionally, it coincides with significant ESG regulatory changes in the EU, ensuring that sustainability data used aligns with current standards. The analysis employs a 63-day (quarterly) rolling window for both training and testing phases. This frequency matches typical corporate reporting cycles and strikes a balance between statistical validity and adaptability to changing circumstances. For smaller businesses, quarterly rebalancing helps reduce transaction costs and operational difficulties.

Overall, the OOS testing framework offers a comprehensive but practical approach to evaluating portfolio strategies under real-world constraints. Instead of pursuing theoretical perfection, the methodology aims to identify models that are transparent, consistent, and usable in environments with limited financial and technical resources.

3.1.3. Return Forecasting Models

This chapter outlines the return estimation methods used in the portfolio models developed in this thesis. Three approaches are considered: historical average returns, exponentially weighted moving averages (EWMA), and the Capital Asset Pricing Model (CAPM). Each method reflects a unique way of dealing with available data and expected risk. By applying them within the same optimization setup, the analysis examines how return forecasts affect portfolio outcomes, with a focus on models that remain practical for medium-sized investors.

3.1.3.1. Historical Method

The traditional mean-variance model introduced by Markowitz (1952) remains a widely adopted baseline in portfolio optimization, particularly due to its intuitive structure and transparency. In the historical variant of this model, the expected return of a portfolio is calculated as a weighted sum of the mean asset returns estimated from historical data:

$$E[R_p] = \sum_{i=1}^{N} x_i \cdot E[R_i] = \mathbf{x}^\top \boldsymbol{\mu}$$

where $\mathbf{x} = (x_1, \dots, x_n)^\top$ denotes the vector of portfolio weights and $\boldsymbol{\mu} = (\mathbb{E}[R_1], \dots, \mathbb{E}[R_N])^\top$ represents the vector of historical mean returns (Markowitz, 1952, S. 81). The associated risk is measured by the variance of portfolio returns, given by:

$$\sigma_p^2 = \mathbf{x}^\top \Sigma \mathbf{x}$$

with Σ being the covariance matrix of asset returns, also estimated from historical observations (Markowitz, 1952, S. 81). Despite ongoing debates about its sensitivity to estimation errors, the historical mean-variance approach remains a widely used benchmark in both theoretical studies and practical applications (Fletcher & Hillier, 2001).

For instance, Kresta and Wang's (2020) portfolios constructed using historical data to minimize variance effectively reduced risk during turbulent market periods, such as the 2008 financial crisis. These portfolios also showed robust performance over a ten-year timeframe when tested OOS. Similarly, portfolios introduced by Guidolin et al. (2018), which are based on historical data, continued to generate significant Sharpe ratios when assessed against multifactor models. Their findings suggest that this approach serves as a practical bridge between theoretical asset pricing models and real-world investment strategies.

For these reasons, the historical mean-variance model is used as the foundational comparator in this thesis. It offers a clear standard to compare more complicated portfolio strategies that consider ESG factors or adjust for risk, particularly for SMEs that do not have advanced estimation tools.

3.1.3.2. EWMA

To better account for changing market risks than fixed historical estimates, this thesis uses the Exponentially Weighted Moving Average (EWMA) model to estimate how asset returns vary together over time. The EWMA framework, as formalized in the RiskMetrics™ Technical Document by (J.P. Morgan, 1996), updates volatility recursively using the following structure:

$$\Sigma_t = \lambda\Sigma_{t-1} + (1-\lambda)\,\mathbf{r}_{t-1}\mathbf{r}_{t-1}^{\top}$$

Here, Σ_t denotes the conditional covariance matrix at time t, $\mathbf{r}_{t-1}$ is the return vector from the previous time step, and λ is the smoothing parameter (J.P. Morgan, 1996, S. 83). Consistent with RiskMetrics' original recommendation for daily data, this research sets $\lambda = 0.94$, which places higher weight on recent observations while retaining information from past market movements through exponential decay (J.P. Morgan, 1996, S. 101).

Research on Singapore equity indices by Kuen and Hoong (1992) demonstrated that EWMA forecasts outperformed both simple historical variance and GARCH(1,1) models in terms of OOS mean-squared error. This pattern has been observed in numerous markets since then. While the standard EWMA model does not account for asymmetric volatility or fat-tailed return distributions, it remains a reliable benchmark in empirical finance. Recent studies, such as by Li and Zhu (2020) acknowledge its fundamental role in risk forecasting models, despite its known limitations.

For the purposes of this thesis, EWMA is particularly well-suited for generating dynamic covariance estimates across rolling windows. Its simplicity, stability, and ability to quickly incorporate latest information make it a viable choice for constructing adaptively optimized portfolios.

3.1.3.3. CAPM

The Capital Asset Pricing Model (CAPM) is a financial theory that was developed by Sharpe (1964) and Lintner (1965) (Fama & French, 2004, S. 25). It builds on the work of Harry Markowitz (1952), who introduced the concept of portfolio diversification. CAPM proposes a straightforward relationship between an asset's expected return and its market risk, measured by beta. The model is based on several key assumptions about investor behaviour and market conditions. These include the ideas that investors are rational and try to avoid risk, they focus on short-term gains, and they can borrow or lend money at a risk-free rate. CAPM also assumes that all investors have the same expectations about future asset returns and that financial markets operate without any barriers or inefficiencies (Elbannan, 2015).

In the CAPM framework, the expected excess return of an asset is defined by its sensitivity to the excess return of the market portfolio, formally captured by beta (Fama & French, 2004, S. 28):

$$E\left[r_i - r_f\right] = \beta_i \cdot E\left[r_m - r_f\right]$$

In practice, β_i is estimated using historical data via the following standard formula (Fama & French, 2004, S. 28):

$$\beta_i = \frac{\mathrm{Cov}\left(R_i - R_f,\ R_m - R_f\right)}{\mathrm{Var}\left(R_m - R_f\right)}$$

This formulation aligns with the approach used by Fletcher and Hillier (2001), who apply CAPM-based expected returns to construct and evaluate portfolio strategies. Their technique involves calculating beta coefficients using rolling windows of historical returns and combining these with the average market excess return observed over time. This approach enables them to predict asset-specific returns while accounting for evolving risk characteristic.

CAPM posits that an asset's expected return above the risk-free rate is solely determined by its systematic risk, measured by beta. In practice, the theoretical world market portfolio is not directly observable, so financial professionals often use broad market indices as substitutes. However, this approach can introduce bias if beta and the market risk premium are calculated using different proxies, as noted by Bartholdy and Peare (2003). This study employs the traditional CAPM framework, utilizing rolling regressions to estimate beta over a historical period and applying the average market excess return from the same index. While acknowledging the potential bias discussed by Bartholdy and Peare (2003), this model is retained as a benchmark due to its simplicity, ease of interpretation, and its established position in portfolio theory literature.

Fletcher and Hillier (2001) further support the practical benefits of CAPM-based strategies. Their analysis shows that using CAPM forecasts leads to better Sharpe ratios for portfolios created through resampling, especially when longer time periods are used for estimation, compared to traditional mean-variance optimizations.

This thesis incorporates the CAPM as an empirically validated method for estimating expected returns. Unlike historical average-based models that rely solely on past performance, CAPM introduces a forward-looking element by connecting return expectations to systematic market risk. This characteristic is particularly valuable in situations where historical return data may be unstable or limited but estimates of market-related risk remain relatively consistent.

3.1.4. Efficiency-Oriented Optimization

In designing portfolio strategies suited for SMEs, the focus must extend beyond minimizing risk to include efficient capital deployment. These businesses need to make the most of their limited capital. SMEs typically have tighter budgets and less financial cushion compared to large institutional investors. As a result, focusing only on minimizing risk could lead to overly conservative strategies that do not align with the company's growth goals. Instead, a more balanced approach using the Sharpe ratio can help optimize returns while still managing risk. This method considers both potential gains and losses, providing a more comprehensive view of investment efficiency for SMEs with constrained resources.

The Sharpe ratio is a widely recognized metric in finance used to assess performance by comparing a portfolio’s excess return to its level of risk, measured as volatility. It is formally defined as (Sharpe, 1994):

$$\mathrm{Sharpe}(w) = \frac{\mathbf{w}^\top \boldsymbol{\mu} - r_f}{\sqrt{\mathbf{w}^\top \boldsymbol{\Sigma} \boldsymbol{w}}}$$

where $\mathbf{w}$ represents the vector of asset weights, $\boldsymbol{\mu}$ the vector of expected returns, $\boldsymbol{\Sigma}$ the covariance matrix of returns, and r_f the risk-free rate, set at 2% annually in this analysis (Cornuejols et al., 2018, S. 102). This ratio effectively communicates how much return is earned per unit of risk, which represent an important consideration for SMEs with limited risk tolerance.

As noted by Cornuejols et al. (2018, S. 102), incorporating a risk-free asset into the portfolio model allows for a more practical approach to optimization. By treating the expected returns vector $\boldsymbol{\mu}$ as excess returns, one can evaluate portfolios against a meaningful benchmark. This adjustment is particularly valuable in situations where decision-making is constrained.

To extend the analysis further, the optimization framework incorporates the One-Fund Theorem. This principle, discussed in the book by Cornuejols et al. (2018, S. 97-98), posits that when a risk-free asset is available, all efficient portfolios can be constructed by combining it with a single tangency portfolio, the one that achieves the maximum Sharpe ratio. These portfolios take the form:

$$w_{\alpha} = \alpha \cdot \mathbf{w}^{*} + (1 - \alpha) \cdot r_f$$

where $\mathbf{w}^{*}$ denotes the tangency portfolio and $\alpha \in [0,1]$ reflects the investor's risk preference. A fully risk-averse investor selects $\alpha = 0$, allocating entirely to the risk-free asset, while $\alpha = 1$ indicates full investment in the tangency portfolio.

In this thesis, the One-Fund approach is applied to all quarters of the year 2024, using the projected risks and returns derived from Sharpe-optimized models based on the historical method of return estimation. By examining these results across different risk tolerance levels, the analysis offers valuable insights into how businesses can adapt their financial planning to achieve their goals while managing risk effectively in a dynamic economic environment.

3.1.5. ESG Integration into Portfolio Optimization

Portfolio managers typically employ two main strategies when incorporating sustainability metrics into their models: using ESG criteria as a strict requirement or integrating it into the overall investment goals. The first approach is straightforward, involving either setting a minimum ESG score for the entire portfolio or eliminating assets that do not meet specific sustainability standards. This method fits well with traditional optimization techniques and is particularly suitable for firms with limited analytical resources. Research shows that using ESG constraints like this usually keeps the basic risk-return balance of portfolios, but it might slightly raise tracking error (Abate et al., 2024; Chen et al., 2021).

However, this approach can lead to sector imbalances by narrowing the pool of potential investments and may exclude companies that are improving their ESG performance but have not yet reached the required thresholds (Spigeleer et al., 2023).

Integrating ESG factors directly into investment objectives shifts the focus from simply following rules to optimizing preferences. This approach involves weighing ESG scores alongside traditional metrics like expected returns and risk, creating a more comprehensive framework for sustainable investing. By doing so, investors can pursue a more nuanced integration of sustainability goals into their portfolios (Escobar-Anel & Jiao, 2024; Zhou, 2024). However, this method comes with its own set of challenges. One significant hurdle is the need for investors to explicitly assign weights to their ESG preferences, which often involves a subjective and case-specific calibration process. As Varmaz, Fieberg and Poddig (2024) point out that embedding ESG considerations within the utility function requires ongoing specification and re-estimation of the trade-offs between risk, return, and ESG factors.

Furthermore, the reliability of ESG data poses a significant challenge to portfolio optimization strategies. Berg et al. (2022) demonstrate substantial discrepancies among ESG scores from major providers, with more than half of the variation stemming from measurement errors. This phenomenon introduces considerable noise into optimization processes that rely heavily on ESG scores. As a result, portfolio construction methods that directly incorporate these scores may be compromised, potentially leading to unreliable or suboptimal investment decisions.

Finally, the embedded-objective method requires increased clarity in explaining how distinct factors are balanced against each other. The research done by Pedersen et al. (2021) on determining the most efficient ESG portfolios highlights that while optimization can produce a range of viable investment options, the ultimate choice must align with an institution's specific goals and values. Their findings emphasize the importance of openly discussing the potential impact on performance metrics, such as the Sharpe ratio or overall returns, when pursuing various levels of ESG-focused investing.

Given the obstacles faced, constraint-based approaches continue to be the more practical and understandable choice, especially for SMEs working with restricted data or computing resources. While ESG metrics are becoming more detailed and comparable, integrating them into objective functions may provide more customized solutions in the future. However, these advancements demand infrastructure and specialized knowledge that are not widely accessible to all organizations at present.

3.1.6. Other Model Considerations

Multi-factor models like the Fama-French three-factor approach have shown better results than CAPM for large companies, but they may not be as useful for small and medium-sized businesses. Sarisoy, de Goeij and Werker (2024) note that the implementation of such models is not simple and straightforward, as both the factor loadings (β) and factor premia (λ) must be estimated jointly. This leads to issues with accuracy, especially when different time periods, weighting methods, and less significant factors are coming into consideration.

Fama and French (2004) themselves acknowledge that the size and value factors at its core are highly volatile and difficult to estimate accurately. This leads to significant uncertainty in forward-looking applications, with expected premiums varying widely between different samples. Moreover, the model's applicability across diverse market contexts remains debatable.

Given the challenges related to parameter calibration and the inconsistent replication of multi-factor models across different regions and firm sizes, this thesis adopts more straightforward benchmarking techniques. Historical mean return estimation, EWMA covariance structures, and the single-factor CAPM are used as reference models. These methods offer greater interpretability and stability, which is particularly important when optimizing portfolios for SMEs, where transparency and simplicity are essential.

3.2. Practical Implementation

The following sections provide an in-depth look at the models employed in this thesis, including the data sources, parameter settings, and backtesting methods used. The focus lies in two main investment strategies: minimum-risk portfolios and maximum Sharpe ratio portfolios. The former aims to protect capital and maintain stability, while the latter seeks to maximize returns relative to risk. These approaches represent different but widely used optimization goals, offering different perspectives on how ESG constraints impact portfolio performance.

3.2.1. Technical Setup and Code Environment

This part of the thesis describes how the portfolio optimization models were put into practice. All the simulations were created using Python version 3.10 and run on Google Colab. This cloud-based platform was chosen because it offers a reproducible environment and allows easy access to the files stored on Google Drive. The use of these tools ensured that the work could be easily replicated and verified by other researchers in the field.

The code is organized into different notebooks, each showing a specific model type: historical, EWMA, and CAPM, with strategies for both minimum risk and maximum Sharpe ratio. ESG-constrained and unconstrained versions are implemented in parallel. The code relies on the following core packages:

Table 1: Packages Used in Optimization Process

Package	Purpose
pandas, numpy	Data handling, matrix calculations
cvxpy	Convex optimization (min-variance models)
scipy.optimize.minimize	Non-convex optimization (Sharpe-max strategies)
matplotlib	Visualizations of performance

Rolling-window backtesting was implemented manually with fixed-length training and testing periods (63 trading days each). Each model produces quarterly portfolios, expected and realized performance metrics, turnover, and benchmark comparisons. The logic, structure, and constraints of the optimization models are documented in Appendix along with code references that mirror the methodology description.

3.2.2. Data Sources and Preprocessing

This research is based on a sample of twenty publicly traded companies listed on the XETRA exchange (n.d.). The selection includes firms from four key sectors: consumer goods and retail (Adidas AG, Puma SE, Beiersdorf AG, Fielmann Group AG, Henkel AG & CO. KGaA), chemicals and pharmaceuticals (Bayer AG, Evonik Industries AG, Merck KGaA, Qiagen N.V., Sartorius AG), industrial manufacturing and engineering (Siemens AG, KION Group AG, Knorr-Bremse AG, Krones AG, Rheinmetall AG, Jungheinrich AG), and technology and IT services (SAP SE, Infineon Technologies AG, Aixtron SE, Bechtle

AG). The data sample was intentionally diversified across sectors to mitigate risk and capture the range of ESG reporting quality and sustainability performance in different industries. This approach reflects the real-world challenges investors face when evaluating companies. The sample includes firms known for their strong ESG practices, as well as those with average or below-average transparency, providing a comprehensive view of the market.

By focusing on XETRA-listed companies, the research ensures consistent euro-based pricing, eliminating the need for currency conversions and simplifying the comparison of returns. This choice also grounds the analysis in the European economic and regulatory landscape, aligning with the thesis's institutional context and considering ESG standards that are particularly relevant to investors in the region.

The data covers the period from January 1, 2020, to December 31, 2024. This time frame includes both highly volatile and more stable periods, such as the COVID-19 pandemic, the war in Ukraine, and post-crisis recovery phases. Including these events allows the research to assess how ESG-integrated optimization strategies perform under varied market conditions. The period also coincides with a rise in ESG-related regulation and investor interest, making it suitable for a sustainability-focused analysis.

Daily price data was collected from Yahoo Finance (n.d.). Logarithmic returns were calculated using the standard formula, due to their additive properties and suitability for continuous compounding (J.P. Morgan, 1996, S. 47):

$$r_t = \ln\left(\frac{p_t}{p_{t-1}}\right)$$

Log returns are preferred in financial modelling for their stability and aggregation properties, which make them useful in portfolio-level risk analysis. This approach is also consistent with the assumption that asset prices follow a log-normal distribution, implying that log returns are normally distributed (J.P. Morgan, 1996, S. 46, 49-50). For the implementation see *Appendix 1: Data Loading and Preprocessing*.

ESG scores were manually collected for all 20 firms using publicly available data from LSEG, S&P Global, and Sustainalytics (London Stock Exchange Group, n.d.; Sustainalytics, n.d.; S&P Global, n.d.). Long-term analysis of ESG performance was constrained by limited data accessibility and the prohibitive costs of specialized databases. To address this, the research relied on the most recently available ESG scores, applying them uniformly across the full evaluation period. Although this static approach does not reflect changes in firm-level ESG performance over time, it realistically represents the constraints faced by smaller investors who typically base decisions on single, snapshot ratings.

Different ESG rating providers use varying approaches to assess companies' ESG performance. LSEG and S&P employ a system where higher scores indicate better ESG alignment. In contrast, Sustainalytics uses a risk-based approach, with lower scores signifying less unmanaged ESG risk exposure. To create consistency across these rating systems for aggregation purposes, the Sustainalytics scores were inverted:

$$\text{Adjusted Sustainalytics Score} = 100 - \text{Original Sustainalytics Score}$$

This adjustment reorients the scale so that higher values correspond to stronger ESG performance, aligning Sustainalytics' risk-based measure with the performance-based frameworks of the other two providers. The updated Sustainalytics score was then included in the calculation of the overall ESG score, together with the original scores from LSEG and S&P Global.

To create a unified ESG measure, data from multiple providers was combined using a weighted average approach. LSEG scores were given the highest importance at 50%, followed by S&P Global at 30%, and Sustainalytics at 20%. These weights were chosen based on factors like data availability, methodology clarity, and topic coverage. The most weight was attributed to the LSEG due to its comprehensive approach and balanced consideration of environmental, social, and governance aspects. S&P Global's contribution was valued for its industry comparisons and detailed thematic analysis through its Corporate Sustainability Assessment. Sustainalytics provided a risk-focused perspective, but we inverted its scores to match the performance-based nature of the other sources. The final composite ESG score remained constant throughout all portfolio optimization scenarios examined in this research.

The ESG score aggregation was implemented in Python (see *Appendix 2: ESG Score Loading and Weighted Composite Calculation*), using data from all three providers mentioned above.

3.2.3. Model Setup

This part of the thesis describes the mathematical framework for optimizing investment portfolios. Two main objectives are explored: minimizing risk through variance reduction and maximizing portfolio efficiency via the Sharpe ratio. Both approaches are subject to the same set of constraints. These strategies are applied to three different forecasting models, each with its own method for calculating expected returns and risk. These formulations provide the basis for the backtesting procedures that will be discussed in detail in the next sections of the thesis.

3.2.3.1. Optimization Objectives

Portfolio managers commonly employ two key strategies when building investment models. The first approach focuses on minimizing overall risk by constructing a portfolio with the lowest possible variance in returns. This minimum-variance strategy aims to create a more stable and predictable investment outcome. The second method, which maximizes the Sharpe ratio, attempts to optimize the efficiency of the portfolio by maximizing the portfolio's excess returns relative to the level of risk taken.

The minimum-variance optimization problem is formulated as:

$$\min_{\mathbf{w}} \mathbf{w}^\top \mathbf{\Sigma} \mathbf{w}$$

where $\mathbf{w}$ is the vector of asset weights and $\mathbf{\Sigma}$ is the covariance matrix of asset returns.

The Sharpe-optimal portfolio maximizes the expected excess return per unit of risk:

$$\max_{\mathbf{w}} \frac{\mathbf{w}^\top \boldsymbol{\mu} - r_f}{\sqrt{\mathbf{w}^\top \mathbf{\Sigma} \mathbf{w}}}$$

where $\boldsymbol{\mu}$ is the vector of expected returns and $r_f = 2\%$ is the fixed annual risk-free rate (Sharpe, 1994).

By comparing these two portfolio optimization approaches, this analysis aims to identify which strategy offers the most practical and suitable solutions for SMEs when making investment decisions, particularly in the context of sustainable portfolios.

3.2.3.2. Constraints

All optimization problems are subject to a consistent set of constraints that reflect realistic investment limitations and sustainability requirements.

- **Budget constraint:**
  The portfolio must be fully invested:

$$\sum_{i=1}^{N} w_i = 1$$

- **Non-negativity constraint:**
  Short selling is not allowed; all weights must be non-negative:

$$w_i \geq 0, \forall_i \in \{1, \dots, N\}$$

- **Maximum allocation per asset (Sharpe strategy only):**
  To ensure diversification, the weight of each individual asset is capped at 20%.

- **ESG constraint (ESG-constrained models only):**
  The portfolio's weighted average ESG score must meet or exceed a specified threshold. This is modelled as a linear inequality:

$$\mathbf{w}^{\top}\mathbf{s} \geq \mathrm{ESG}_{\mathrm{threshold}}$$

The primary adjustment involves introducing a minimum ESG score requirement at the portfolio level. This ensures that the selected assets collectively meet a predetermined sustainability threshold. By integrating this constraint, investors can align their portfolios with both financial objectives and sustainability goals, reflecting the evolving landscape of responsible investing and regulatory expectations.

A comparison is then made between the constrained and unconstrained versions of the historical, EWMA, and CAPM-based models. This approach helps to clarify how ESG integration influences portfolio structure and performance. By holding all other modelling conditions constant, the impact of sustainability requirements can be assessed more directly.

Each asset in the investment universe is assigned a composite ESG score, constructed according to the formula described in the previous section. These scores are held constant throughout the sample period due to data limitations, reflecting the static nature of ESG information often available to smaller investors. The ESG scores are used not to forecast returns or adjust risk assumptions, but rather to filter portfolio allocations through a minimum sustainability requirement. This logic is implemented uniformly across all optimization frameworks.

The incorporation of ESG criteria adds a new dimension to portfolio optimization. Instead of being a mere preference, the ESG threshold becomes a strict requirement, limiting investment options to portfolios that meet or exceed a specified ESG score. This approach naturally favours assets with higher ESG ratings, unless lower-rated ones are essential for risk reduction within the given constraints.

3.2.4. Model Variants

The chapter outlines various methods for predicting returns, which are utilized in the portfolio optimization models developed throughout the thesis. The approaches range from purely historical analysis to more sophisticated techniques that incorporate forward-looking, market-based assumptions. This systematic progression enables a structured comparison of how each model performs under different market conditions and with varying data inputs. The subsequent sections provide a detailed explanation of the design and rationale behind each forecasting approach, offering insights into their respective strengths and limitations.

The chapter directs to consult the Appendix for in-depth technical details regarding the implementation within the coding environment. However, it is worth noting that only versions without an ESG constraint were included in the Appendix. This decision was made because the sole difference between the models lies in the inclusion of the constraint, while the fundamental optimization setup remains consistent across both versions.

3.2.4.1. Historical Model

The mean-variance model applied in this thesis builds on the original concept introduced by Markowitz (1952), where portfolio optimization is framed as a quadratic programming problem. Expected returns and covariances are estimated from a fixed historical window, and the model calculates asset weights by minimizing portfolio variance while satisfying standard investment constraints. This setup is used as the benchmark model against which alternative strategies are compared.

Following section describes the setup for the historical approach.

Let:

- N be the number of assets in the investment universe,

- T be the number of training days in the rolling window,
- $\mathbf{r}_t \in \mathbb{R}^N$ denote the vector of daily log returns across N assets on day *t,*
- $\mathbf{R} \in \mathbb{R}^{T \times N}$ denote the matrix of log returns over the current training window.

Logarithmic returns are computed as:

$$r_{t,i} = \ln\left(\frac{p_{t,i}}{p_{t-1,i}}\right), \text{ for } i = 1, \dots, N, \ t = 2, \dots, T$$

Where $p_{t,i}$ is the adjusted closing price of asset i at time t. Log returns are used due to their time additive property, approximate normality over short horizons, and compatibility with continuous compounding (J.P. Morgan, 1996, S. 47).

Now the model inputs need to be estimated, starting with the expected returns.

Let $\boldsymbol{\mu} \in \mathbb{R}^N$ be the vector of expected returns. The historical model estimates μ_i as the arithmetic mean of past daily returns (Cornuejols et al., 2018, S. 106):

$$\mu_i = \frac{1}{T}\sum_{t=1}^{T} r_{t,i}$$

To convert these daily returns into annualized form (to match annualized risk), the expected returns are scaled by the approximate number of trading days per year $D = 252$:

$$\mu_i^{\text{annual}} = D \cdot \mu_i$$

In this model, expected returns are not used as optimization targets but are retained for OOS performance comparison.

As a next step, the covariance matrix should be defined.

Let $\Sigma_{ij} \in \mathbb{R}^{N \times N}$ be the sample covariance matrix of asset returns, estimated from our training window (Cornuejols et al., 2018, S. 106):

$$\Sigma_{ij} = \frac{1}{T-1}\sum_{t=1}^{T}(r_{t,i} - \mu_i)(r_{t,j} - \mu_j)$$

or in matrix notation:

$$\boldsymbol{\Sigma} = \frac{1}{T-1}(\mathbf{R} - \overline{\mathbf{R}})^\top(\mathbf{R} - \overline{\mathbf{R}})$$

where $\bar{\mathbf{R}} \in \mathbb{R}^{T \times N}$ is a matrix with each column i filled with the mean return μ_i. As with returns, the covariance matrix is annualized:

$$\mathbf{\Sigma}^{annual} = D \cdot \mathbf{\Sigma}$$

This matrix is used as the risk input for optimization.

Appendix 3: Historical Min Risk (No ESG) - Optimization and Backtesting outlines the full implementation of the historical minimum-variance model, including portfolio optimization via cvxpy, a rolling backtest with quarterly rebalancing, and evaluation of expected and realized performance metrics.

Appendix 4: Historical Max Sharpe (No ESG) - Optimization and Backtesting extends the historical framework to a Sharpe-ratio optimization setup, using scipy.optimize for numerical minimization. The rolling backtest structure mirrors that of the risk-minimizing strategy to ensure comparability.

3.2.4.2. EWMA-Based Model

The historical approach provides a clear starting point but assumes volatility remains steady over time. To better reflect real-world market dynamics, an alternative model incorporates an exponentially weighted covariance structure using EWMA. This approach accounts for the fact that market volatility evolves over time, giving more weight to recent observations. Although expected returns are still computed using a simple average over the estimation window, the EWMA method enables a more adaptive evaluation of risk. This adaptability is particularly valuable in markets characterized by fluctuating volatility levels.

The EWMA model retains the same structure and constraints as the historical approach, allowing for a controlled comparison where only the risk estimation method differs.

As in the historical model, let:

- N be the number of assets in the investment universe,
- T be the number of training days in the training window,
- $\mathbf{r}_t \in \mathbb{R}^N$ denote the vector of daily log returns across N assets on day *t,*
- $\mathbf{R} \in \mathbb{R}^{T \times N}$ denote the return matrix in the training window.

The expected return calculation was not changed and is identical to the historical approach, where the expected returns for each asset are calculated as an arithmetic mean of returns during the training period. To allow a direct comparison, the annual return was also included in this model, taking into consideration 252 trading days.

In contrast to the equal-weighted sample covariance, the EWMA matrix assigns higher weights to more recent returns using a decay parameter $\lambda \in (0,1)$. The EWMA covariance matrix (J.P. Morgan, 1996, S. 83) is defined as:

$$\Sigma_t = \lambda \boldsymbol{\Sigma}_{t-1} + (1-\lambda)(\mathbf{r}_t - \mu)(\mathbf{r}_t - \boldsymbol{\mu})^\top$$

In practice, the matrix is initialized using the first observation and updated iteratively over the training window. The final matrix is then annualized using the same factor:

$$\boldsymbol{\Sigma}^{annual} = D \cdot \boldsymbol{\Sigma}_{\text{EWMA}}$$

A commonly used value for the decay factor is $\lambda = 0.94$, which reflects a memory of roughly 16-20 days, though other values can be evaluated depending on market context.

The code implementation for the EWMA-based minimum-variance strategy is provided in *Appendix 5: EWMA Min Risk (No ESG) - Optimization and Backtest*. It blends short-term EWMA returns with long-term historical averages.

Appendix 6: EWMA Max Sharpe (No ESG) - Optimization Setup and Backtesting presents the corresponding Sharpe-optimized version of the EWMA model. This version applies the same rebalancing schedule as the minimum-risk variant but prioritizes return efficiency.

3.2.4.3. CAPM-Based Model

Unlike historical averages or exponentially weighted moving averages, CAPM calculates expected returns based on an asset's relationship to market performance. This approach provides a more forward-looking perspective on potential returns. However, the model still relies on historical data for estimating the covariance matrix used in portfolio optimization.

Despite the changes in return estimation, the fundamental structure of the optimization problem remains consistent with previous models. The same constraints and overall framework are applied, allowing for continuity in the portfolio construction process while incorporating this new methodology for forecasting returns.

As before, let:

- N be the number of assets in the investment universe,

- T the number of trading days in the trianing window,
- $\mathbf{r}_t \in R^N$ denote the vector of daily log returns across N assets on time t,
- $\mathbf{R} \in R^{T \times N}$ the return matrix over the training window.

In this setup, asset returns are expressed in excess terms by subtracting a constant risk-free rate R_f (2% annualized, divided by 252 trading days to match daily frequency). The expected return of asset i is then computed according to the CAPM formula (Fama & French, 2004, S. 28):

$$\mu_i^{\text{CAPM}} = R_f + \beta_i \cdot \left(\overline{R_m} - R_f\right)$$

where:

- $\overline{R_m}$ is the average daily excess return of the market (DAX index) over the training window,
- β_i is the estimated exposure of asset i to market movements, calculated as (Fama & French, 2004, S. 28):

$$\beta_i = \frac{\text{Cov}\left(R_i - R_f,\ R_m - R_f\right)}{\text{Var}\left(R_m - R_f\right)}$$

To compute these estimates, a rolling window of T training days is used to dynamically re-estimate each asset's beta in each quarter. Daily excess returns for both the asset and the market index are first aligned and filtered to ensure equal length and data availability. This process produces a vector $\boldsymbol{\mu}^{\text{CAPM}} \in \mathbb{R}^N$ of expected returns for all assets of forward-looking return estimates for all available assets in the current training window.

Appendix 7: CAPM Min Risk (No ESG) - Optimization Setup and Backtesting contains the implementation details of the CAPM-based minimum-variance strategy, including dynamic beta estimation from rolling windows and portfolio optimization based on forecasted returns and historical covariance.

Appendix 8: CAPM Max Sharpe (No ESG) - Optimization and Backtesting outlines the Sharpe-maximizing version of the CAPM model. It uses the same expected return estimation via CAPM but alters the optimization objective, relying on numerical maximization of excess Sharpe under defined constraints.

3.2.5. OOS Testing and Evaluation

The practical relevance of the suggested portfolio models is evaluated through an OOS backtesting methodology. By distinctly separating the model estimation phase from performance assessment, this method helps mitigate overfitting issues and provides a more accurate representation of how the models would perform in real-world scenarios. The evaluation process includes not only the tracking of portfolio returns and risk but also incorporates a diverse set of performance metrics to facilitate comparisons across different models, strategies, and ESG constraints. In the subsequent sections, a comprehensive explanation of the testing procedure and the chosen performance indicators will be described, offering readers a thorough understanding of the evaluation framework employed in this research.

3.2.5.1. Rolling Window Design

The OOS testing process uses a rolling window approach with a fixed time frame. Each optimization model is trained on 63 trading days (about one calendar quarter) of data, and then the resulting portfolio allocations are applied to the next 63-day test period. The window then moves forward by one test period, and the process repeats until the end of the available data.

For each window, expected returns and risk parameters are estimated using only the training data. Portfolio allocations are optimized based on these estimates and remain unchanged throughout the test period. There is no re-optimization conducted during the test phase, ensuring that all performance metrics truly reflect OOS results. This procedure is applied uniformly across all model types and strategy variations to maintain consistency in our evaluation.

3.2.5.2. Backtesting Procedure

The backtesting process follows a consistent pattern for each rolling window. Initially, model-specific parameters like expected returns and covariance matrices are calculated using only the training data. These calculations vary based on the chosen approach, such as historical averages, exponentially weighted covariance, or CAPM-based return expectations.

Using these inputs, the portfolio is then optimized according to either the minimum-variance or maximum Sharpe ratio strategy. During this optimization, all predefined constraints are applied, including full investment, non-negativity, asset allocation limits, and ESG thresholds where relevant.

After obtaining the optimized portfolio weights, they are applied to the returns in the subsequent test period. These weights remain constant throughout this phase, simulating a passive holding strategy based on the previous quarter's information.

Lastly, performance metrics are computed using the actual returns from the test window. These results are recorded and combined across all rolling periods, allowing for a thorough evaluation of each model's performance over time. This backtesting methodology is consistently applied across all portfolio models and strategy types to ensure fair and comparable assessments.

3.2.5.3. Evaluation Metrics

To evaluate the performance of various portfolio optimization models, a uniform set of performance indicators is computed for each strategy across all backtesting periods. These indicators are based on OOS results, ensuring that the reported figures represent realized performance grounded in prior estimates, free from any forward-looking bias. The indicators include measures of return, risk, forecasting accuracy, market exposure, and implementation effort.

Return is reported as the annualized geometric mean of realized portfolio performance over each test window. It reflects the cumulative gains a strategy would achieve over time, assuming returns compound daily. Volatility is computed as the annualized standard deviation of daily returns and serves as a proxy for total portfolio risk.

The Sharpe ratio is calculated as the excess return over the risk-free rate of 2%, divided by the realized volatility (Sharpe, 1994):

$$\text{Sharpe} = \frac{R - R_f}{\sigma}$$

The Sortino ratio is used as an alternative measure of risk-adjusted performance, penalizing only downside volatility (Sortino & Van Der Meer, 1991):

$$\text{Sortino} = \frac{R - R_f}{\sqrt{252} \cdot \text{ std}(\min, (r_t, 0))}$$

Beyond performance, forecasting quality is evaluated using Mean Absolute Error (MAE) between expected and realized values for both return and Sharpe ratio:

$$\mathrm{MAE}_{\mathrm{return}} = \frac{1}{n}\sum_{i=1}^{n}\left|R_i^{\mathrm{expected}} - R_i^{\mathrm{realized}}\right|$$

$$\mathrm{MAE}_{\mathrm{Sharpe}} = \frac{1}{n}\sum_{i=1}^{n}\left|Sharpe_i^{\mathrm{expected}} - Sharpe_i^{\mathrm{realized}}\right|$$

These values are derived from the training (expected) and test (realized) windows and provide insight into each model's predictive stability (Hyndman & Koehler, 2006).

To capture market exposure, each portfolio's beta is calculated relative to the DAX index by regressing daily portfolio returns against market returns in the training period. This allows an interpretation of systematic risk and helps determine whether ESG constraints or optimization logic shift portfolios away from broad market dynamics.

In addition to standard performance indicators, portfolio turnover is calculated for each rebalancing point to evaluate trading intensity and potential implementation costs. Turnover is defined as the sum of absolute differences in asset weights between two consecutive test periods:

$$\mathrm{Turnover}_t = \sum_{i=1}^{N}\left|w_{i,t} - w_{i,t-1}\right|$$

All these metrics are consistently applied across models and strategies, forming the basis for the comparative analysis presented in Chapter 4.

3.2.6. One-Fund Implementation

Based on the Sharpe-optimized portfolios, a One-Fund model applies to model investor-specific portfolio choices. Following this theorem, any efficient portfolio in the presence of a risk-free asset can be represented as a linear combination of the risk-free return and a single tangency portfolio, represented in this case by the Sharpe-maximizing portfolio identified for each model.

To make this concept usable in practice, investor preferences are simulated using a parameter α that ranges from 0 to 1. This parameter reflects investor's risk preferences, where $\alpha = 0$ is representing risk-averse behavior (100% investment into the risk-free portfolio) and $\alpha = 1$ being the representation of risk-seeking investor (100% investment into the risky Sharpe-optimized portfolio).

For each value of α, the portfolio return is computed using:

$$R_\alpha = \alpha \cdot (\mathbf{w}^{*\top}\boldsymbol{\mu}) + (\mathbf{1} - \alpha) \cdot r_f$$

where $\mathbf{w}^*$ is the Sharpe-optimal portfolio, $\boldsymbol{\mu}$ is the return estimate and $r_f = 2\%$ is the fixed risk-free rate.

This model is applied across all quarters of the year 2024 to capture a range of possible outcomes, rather than focusing on a single point in time. This broader application allows for a more realistic assessment of how investors might respond to changing market conditions and evolving ESG dynamics throughout the year. For each quarter, the Sharpe-optimal portfolio is combined with the risk-free asset to create One-Fund portfolio, depending on the diverse levels of risk preference. This approach is then compared directly with the minimum risk portfolio. For simplicity reasons, this method was performed only on historical model, due to its transparency and ease in implementation.

The evaluation is applied across all backtesting periods of the year 2024 and for both ESG-constrained and unconstrained model variations. Finally, the comparison between these two strategies was conducted based only on the historical approach for simplicity reasons.

3.2.7. ESG Threshold Sensitivity Analysis

To gain insights into the impact of increasingly stringent ESG criteria on portfolio outcomes, a comprehensive analysis was conducted by systematically testing various ESG thresholds while keeping model parameters consistent.

This section builds upon the previously established ESG-constrained optimization framework, examining how varying levels of sustainability requirements affect portfolio composition and performance over time. The historical mean-variance model was used as the analytical foundation, providing a stable basis for assessing the influence of ESG thresholds on portfolio structure and results.

The sensitivity analysis was conducted by incrementally adjusting the ESG constraint from a threshold of 60 to 83, increasing by one point at each step. For every level, the optimization problem was solved under the same minimum-variance objective, with an added condition that the portfolio's weighted average ESG score must meet or exceed the specified threshold. All other model inputs and assumptions were held constant, allowing for a controlled assessment of how tighter sustainability constraints shape portfolio decisions. OOS performance was evaluated for each threshold and rebalancing period using a consistent test window.

The primary goal of this analysis is to provide insights into the compromises that must be made between minimizing risk and adhering to ESG principles. Since the same optimization method is used as in the unconstrained model, any differences in performance or asset selection can be directly attributed to the ESG constraint. This consistency ensures that the effects of sustainability requirements are isolated, without being obscured by other modelling differences. The analysis is based entirely on historical model, because of its transparency and simplicity in use.

4. Results

This chapter details the empirical outcomes which were obtained from the ESG-integrated portfolio optimization framework established earlier in this thesis. Expanding upon the methodology introduced in Chapter 3, the primary objective here is to examine how the inclusion of ESG constraints influences portfolio results under various optimization strategies. The analysis consistently compares the performance of models with and without ESG factors, applying a unified OOS evaluation framework across all model variations.

The research is centred around three main portfolio models, historical mean-variance, EWMA, and CAPM, each tested under two different objectives: minimizing variance and maximizing the Sharpe ratio. For every model pairing, ESG-constrained and unconstrained portfolios were constructed, delivering a total of twelve core model variants. Additionally, a simplified One-Fund approach was applied to the Sharpe-optimized configurations to highlight how ESG considerations might influence investor decisions across varying levels of risk preferences.

The models were assessed through a rolling, OOS backtesting process covering the period from 2020 to 2024, with quarterly portfolio rebalancing. Key performance metrics include expected versus realized returns, volatility, Sharpe and Sortino ratios, turnover, beta (relative to the DAX), and the average ESG score. Forecasting accuracy was obtained using the MAE between predicted and actual outcomes.

To ensure consistency and avoid distortions from extreme values, fixed evaluation caps were applied across all Sharpe-maximizing strategies. A minimum volatility of 5% was enforced to stabilize the Sharpe ratio in low-risk environments, while annualized returns were capped at 10% to reflect realistic performance boundaries. These adjustments were implemented during the evaluation phase and not embedded within the optimization itself.

The findings are presented in three main segments. Initially, each model's individual performance is outlined, comparing ESG and non-ESG versions side by side. This is followed by a cross-comparison that highlights the central trade-offs observed, particularly regarding returns, risk, and asset allocation under ESG constraints. Lastly, a sensitivity analysis examines how tightening ESG thresholds affects the size of the investment universe and alters the balance between risk and return.

The chapter concludes with a synthesis of the key findings, which serve as a foundation for the following discussion. Emphasis is placed on the practical implications for SMEs, which, despite their constrained modelling capabilities, increasingly face expectations aligned with ESG-driven financial standards.

4.1. Historical Model

The analysis in this section examines how the historical return-based model performs when constructing portfolios using two different methods: minimizing variance and maximizing the Sharpe ratio. Both methods were assessed with and without ESG constraints. This evaluation is especially relevant for SMEs, as they typically have restricted resources for complex modelling. These businesses need concrete, data-driven evidence to support any decision to incorporate ESG criteria into their investment strategies.

4.1.1. Minimum Risk Strategy

The minimum-variance portfolio showed stable and consistent results during the 2020 - 2024 backtesting period. Since the strategy aims to reduce volatility, it led to steady asset allocations and moderate trading levels.

The ESG-constrained version clearly performed better than the unconstrained one. It delivered an average annual return of 12.9%, compared to 8.9% for the non-ESG portfolio. Sharpe ratios also improved with ESG integration, reaching 0.80 versus 0.43. This gain in performance came without added risk, as both versions maintained similar annual volatility around 13%.

In terms of forecast accuracy, the ESG portfolio had a slightly higher mean absolute error in return prediction (2.5% vs. 0.2%). Trading activity stayed moderate for both, with the ESG version showing slightly lower quarterly turnover (0.87 vs. 0.91). Portfolio analysis showed that applying ESG filters shifted allocations toward firms with better ESG scores, without reducing diversification too much. Overall, the minimum-variance strategy remained effective and stable when combined with a fixed ESG constraint.

4.1.2. Maximum Sharpe Strategy

The Sharpe-maximizing version of the historical model presented greater challenges. Without adjustments, it often led to unstable portfolios caused by high asset concentration and overly optimistic return estimates. To improve stability in OOS tests, two constraints were applied: a 20% cap on individual asset weights, return capping up to 10% and a minimum volatility limit of 5%.

With these adjustments, the model produced more reasonable results. The unconstrained version achieved a realized return of 9.5%, while the ESG-constrained portfolio returned 9.4%. Despite strong in-sample projections, both versions only reached moderate OOS Sharpe ratios of about 0.50-0.51. This gap points to potential overfitting and reflects the difficulty of using Sharpe optimization when return estimates are noisy.

Volatility was notably higher than in the minimum-variance case, ranging from 18.7% to 18.9%, and turnover also increased significantly to between 1.32 and 1.35 per quarter. Such high trading levels may be impractical for SMEs. While the ESG constraint slightly improved forecast consistency, it did not lead to clearly better performance. The mean absolute error for return estimation went up to 0.58 for both strategies. Compared to the minimum-risk strategy, this approach added complexity and sensitivity without delivering reliable gains.

4.1.3. Cross-Strategy and Benchmark Comparison

The table below presents key performance indicators for the different versions of the historical model. An equal-weighted portfolio is also shown for comparison. While its average return of 2.71% and volatility of 19.14% was affected by high volatility during early 2020, its Sharpe ratio remained relatively steady (0.23) and serves as a useful reference point. The table below summarizes the key performance indicators for all historical model scenarios:

Table 2: Cross-Strategy Comparison of Historical Model

Strategy	*ESG*	*Return*	*Volatility*	*Sharpe*	*Sortino*	*MAE (Return)*	*MAE (Sharpe)*	*Beta*	*Turnover*
Min Risk	No	8.9%	13.2%	0.44	0.73	0.22	1.79	0.50	0.91
Min Risk	Yes	12.8%	13.8%	0.80	1.47	0.25	1.92	0.61	0.87
Max Sharpe	No	9.5%	18.7%	0.51	1.62	0.58	3.59	0.74	1.32
Max Sharpe	Yes	9.5%	18.9%	0.50	1.59	0.58	3.50	0.78	1.35

Table 2: Cross-Strategy Comparison of Historical Model compares all four historical model setups across key performance indicators: return, volatility, Sharpe ratio, Sortino ratio, forecast accuracy (MAE), beta exposure, and portfolio turnover. Each row represents a specific combination of strategy type (Min Risk or Max Sharpe) and ESG constraint (Yes or No), enabling a direct assessment of how ESG integration influences performance and stability across optimization goals.

In both strategies, the ESG-constrained models matched or outperformed their unconstrained versions, with the strongest results seen in the minimum-risk setup. This suggests that adding ESG filters can improve risk-adjusted returns without increasing portfolio volatility or trading frequency. Sortino ratios, which focus on downside risk, also supported the better performance of ESG-constrained models. In the minimum-risk strategy, the ESG version scored higher (1.47 vs. 0.73), showing stronger protection against losses. In the Sharpe-maximizing strategy, the ESG version also did slightly better, though the difference was smaller.

Beta values reveal further differences in market exposure. The minimum-risk models had relatively low beta values (around 0.50 to 0.61), meaning their returns were less tied to broader market movements, as measured by the DAX index. This lower exposure supports the strategy's goal of reducing volatility and may help protect capital during market downturns. On the other hand, the maximum Sharpe strategies showed higher beta values (roughly 0.74 to 0.78), reflecting stronger links to market trends but also greater exposure to systematic risk. ESG constraints had a small effect on beta, generally increasing it slightly in both strategies. This was likely due to excluding lower-rated firms that tend to be less correlated with the market.

For SMEs, this distinction is particularly important, as their investment choices are influenced by limited resources, a strong focus on capital preservation, and a lower tolerance for volatility, especially in uncertain economic conditions. In such cases, ESG-integrated, low-beta strategies may offer a more practical and risk-aware approach to sustainable investing. In contrast, higher-beta portfolios may suit investors aiming for growth and willing to take on more risk. ESG integration did not reduce market exposure sharply but maintained or slightly increased it without weakening other aspects of portfolio performance.

4.2. EWMA Model

This section reviews the performance of the EWMA-based portfolio model, which estimates returns and covariance using an exponentially weighted moving average to better reflect recent market behaviour. Two strategies were evaluated, minimum risk and maximum Sharpe ratio, each with and without ESG constraints. Compared to the historical model, the EWMA method reacts more quickly to market changes but is also more sensitive to short-term fluctuations.

4.2.1. Minimum Risk Strategy

The EWMA minimum-risk strategy delivered consistent results during the test period. The ESG-constrained version performed better, with an average annual return of 10.5%, compared to 8.0% for the unconstrained version. Risk levels stayed similar, with annual volatility at 14.1% (ESG) and 13.1% (non-ESG). Sharpe ratios also improved under ESG constraints, rising from 0.45 to 0.63. The Sortino ratio followed the same trend, increasing from 0.72 to 1.07, indicating better downside risk protection. Forecast accuracy remained steady, with return MAE at 0.23 for both.

Both versions had relatively low turnover (0.54 ESG, 0.68 non-ESG), showing the strategy did not require frequent trades. ESG constraints further reduced turnover slightly, suggesting more stable allocations. Beta values went up from 0.57 to 0.70 after the ESG-constraint introduction, indicating increased yet moderate market exposure for ESG-portfolios.

4.2.2. Maximum Sharpe Strategy

Like the historical model, the Sharpe-maximizing EWMA strategy needed extra constraints - a 5% volatility floor, return fluctuations up to 10% and a 20% cap on individual assets to avoid unstable or unrealistic results.

Even with these adjustments, performance was more volatile. The non-ESG version achieved an average annual return of 10.0%, but the Sharpe ratio was at 0.53 due to higher volatility (18.5%) and larger prediction errors. The ESG version return was at 10.5% and had slightly better risk-adjusted results with a Sharpe of 0.55 and a Sortino of 1.70, compared to 1.62 for the non-ESG model.

Forecasting was less accurate for both model setups. MAE(Return) resulted into a value of 0.58 for the ESG-model and 0.56 for non-ESG version, highlighting a gap between projected and actual performance especially under Sharpe targeting. Turnover was higher than in the minimum-risk strategy (1.41 ESG, 1.34 non-ESG), showing a more active trading pattern.

4.2.3. Cross-Strategy Comparison

The table below summarizes key performance metrics for all EWMA-based portfolio versions, covering both expected and realized result:

Table 3: Cross-Strategy Comparison of EWMA Model

Strategy	*ESG*	*Return*	*Volatility*	*Sharpe*	*Sortino*	*MAE (Return)*	*MAE (Sharpe)*	*Beta*	*Turnover*
Min Risk	No	8.0%	13.1%	0.46	0.72	0.23	1.75	0.58	0.68
Min Risk	Yes	10.5%	14.1%	0.64	1.07	0.23	1.58	0.70	0.54
Max Sharpe	No	10.0%	18.5%	0.54	1.62	0.56	3.69	0.75	1.35
Max Sharpe	Yes	10.5%	18.9%	0.55	1.70	0.58	3.68	0.79	1.41

Table 3: Cross-Strategy Comparison of EWMA Model summarizes performance metrics across all EWMA model configurations, covering both minimum-risk and Sharpe-maximizing strategies with and without ESG constraints.

In the minimum-risk strategy, ESG integration led to a higher annual return (10.5% vs. 8.0%) and a stronger Sharpe ratio (0.64 vs. 0.46). The slightly higher volatility (14.1% vs. 13.1%) was offset by improved downside protection, as shown by the Sortino ratio (1.07 vs. 0.72), and lower turnover (0.54 vs. 0.68). This indicates that ESG constraints improved performance and portfolio stability without increasing trading activity.

In the Sharpe-maximizing strategy, both versions reached similar returns (around 10.0–10.5%) and Sharpe ratios (0.54–0.55), but with significantly higher volatility (18.5%–18.9%) and turnover (1.35 ESG, 1.34 non-ESG). Forecast accuracy also worsened, with MAE(Return) increasing to 0.56–0.58. These results suggest that while the strategy achieved higher returns, it was more sensitive to input instability and required more frequent rebalancing.

Overall, the ESG-constrained minimum-risk variant offered the best trade-off between return, risk, and implementation effort, making it a practical option for SMEs focused on stability and sustainability.

4.3. CAPM Strategy

This section evaluates the performance of the CAPM-based portfolio optimization approach under two portfolio construction strategies: minimum variance and maximum Sharpe ratio. Each strategy is evaluated in both ESG-constrained and unconstrained formats. The aim is to examine whether CAPM's forward-looking return estimation improves portfolio outcomes and how well it aligns with SME investment needs, especially under sustainability constraints.

4.3.1. Minimum Risk Strategy

The CAPM minimum-variance strategy delivered competitive performance during the 2020-2024 backtesting period. With a realized return of 12.9% in the ESG-constrained version, it closely matched the performance of the historical model under the same conditions. Sharpe ratios were identical at 0.81 for both ESG and non-ESG variants, suggesting that the forward-looking return estimation did not significantly alter risk-adjusted outcomes when volatility was minimized.

Volatility remained moderate at 13.8% for the ESG model and 13.2% for the non-ESG model. Turnover was slightly higher in the unconstrained case (0.91 vs. 0.87), while beta values pended around 0.61 and 0.50, respectively. This indicates a relatively low level of market dependence, especially in the non-ESG variant. Interestingly, the ESG-constrained version yielded a higher Sortino ratio (1.47 vs. 0.73), highlighting stronger downside risk protection. These results suggest that CAPM-based risk minimization can be both effective and ESG-compatible for SMEs.

4.3.2. Maximum Sharpe Strategy

The Sharpe-optimized version of the CAPM model showed obvious signs of instability. The ESG-constrained portfolio yielded a realized return of just 4.9%, compared to an even lower 2.9% for the non-ESG version. Despite high expected returns based on in-sample forecasts, the OOS Sharpe ratios fell sharply - 0.27 for ESG and just 0.25 for non-ESG indicating a disconnect between forecasted performance and actual outcomes.

Volatility reached 20.1% in the ESG model and 19.1% in the non-ESG model, making both portfolios riskier and less efficient than their minimum-risk counterparts. Turnover also increased notably (1.08-1.11), and beta values rose to 0.99 and 0.93, respectively, suggesting a higher sensitivity to market movements. These characteristics make the CAPM Sharpe-maximizing approach a poor fit for SMEs with limited tolerance for volatility or model risk. While ESG constraints marginally improved stability, they did not resolve the underlying issues tied to unstable return forecasts.

4.3.3. Cross-Strategy Comparison

The table below summarizes the key performance metrics across all CAPM-based portfolio variants:

Table 4: Cross-Strategy Comparison of CAPM Model

Strategy	*ESG*	*Return*	*Volatility*	*Sharpe*	*Sortino*	*MAE (Return)*	*MAE (Sharpe)*	*Beta*	*Turnover*
Min Risk	No	8.9%	13.2%	0.44	0.73	0.20	1.63	0.50	0.91
Min Risk	Yes	12.9%	13.8%	0.80	1.47	0.21	1.56	0.61	0.87
Max Sharpe	No	2.8%	19.0%	0.24	1.02	0.35	1.56	0.93	1.11
Max Sharpe	Yes	4.8%	20.1%	0.27	0.89	0.0.34	1.41	0.99	1.07

In contrast to the historical and EWMA models, the CAPM strategy revealed a more pronounced gap between the two optimization approaches. The minimum-risk variant delivered solid and consistent performance, particularly in the ESG-constrained version, which achieved the highest return (12.9%) and the strongest Sortino ratio (1.47) among all CAPM setups. In contrast, the Sharpe-maximizing version failed to convert its expected return forecasts into realized gains. Both ESG and non-ESG variants produced weak Sharpe ratios (0.27 and 0.24) and low realized returns (4.8% and 2.8%), reflecting instability in beta estimation and market premium assumptions under short training windows.

While CAPM remains a useful framework for constructing low-risk, ESG-compliant portfolios, its application becomes unreliable in return-seeking strategies. The model's sensitivity to input fluctuations and reliance on forward-looking assumptions make it less practical for SMEs with limited modelling capacity and lower tolerance for volatility.

4.4. One-Fund Model

This section evaluates the performance of the One-Fund model, which combines a risk-free asset with a single tangency portfolio derived from Sharpe ratio optimization. The goal is to examine whether investor-specific preferences, represented by different risk levels ($\alpha \in [0, 1]$), can generate balanced portfolios that align with sustainability goals while preserving performance. In contrast to the minimum-variance and pure Sharpe-optimized portfolios, the One-Fund strategy allows for risk scaling without re-optimizing at each point.

The analysis is based on the historical return forecasting approach, selected for its transparency and superior risk-adjusted performance in earlier tests. The One-Fund strategy is constructed for the four quarters of 2024, enabling comparison across varying market conditions. For each quarter, the One-Fund portfolios are compared against the minimum-risk strategy under both ESG-constrained and unconstrained setups.

4.4.1. Performance Across Risk Preferences

As shown in the figure below, One-Fund portfolios exhibit a wide dispersion in realized returns, strongly dependent on the chosen risk preference α. Portfolios with low α values (e.g., 0.0 0.25) delivered stable but modest returns in all quarters, closely aligning with the fixed risk-free rate of 2%. As α increases toward 1.0, return variability also increases significantly.

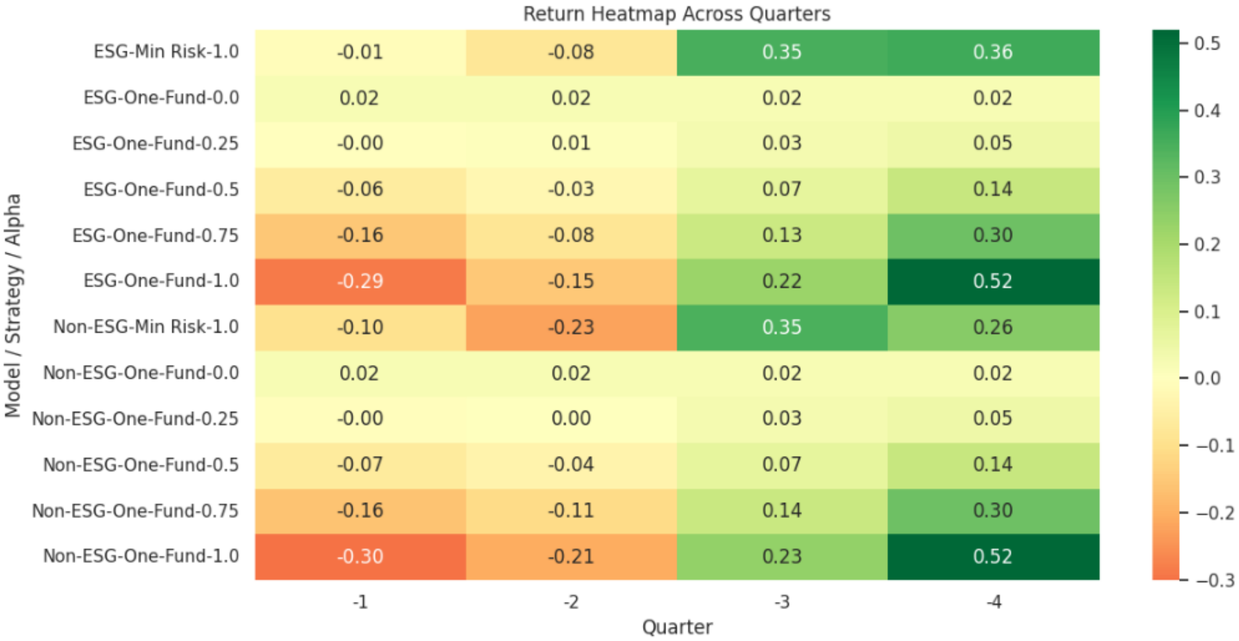

Figure 1: One-Fund Model Compared with Minimum Risk Strategy

Notably, the ESG-constrained One-Fund model with α = 1.0 yielded the highest return (52%) in Q4 but also experienced the sharpest losses (-29%) in Q1. The same trend holds for the non-ESG setup, with strong gains in late 2024 but severe drawdowns earlier in the year. These outcomes demonstrate that risk-seeking configurations are overly sensitive to market cycles and may not be suitable for SMEs with lower volatility tolerance.

In contrast, moderate risk preferences (α = 0.5-0.75) delivered more balanced outcomes. ESG One-Fund portfolios at α = 0.5 achieved 14% in Q4 with relatively muted losses earlier in the year. These levels closely approximate the average performance of minimum-risk portfolios, suggesting that the One-Fund model can approximate a similar outcome while offering greater flexibility.

4.4.2. ESG vs. Non-ESG Comparison

Across all quarters and α levels, ESG-constrained One-Fund portfolios generally mirrored the performance patterns of their non-ESG counterparts. However, ESG portfolios tended to exhibit slightly lower drawdowns in early quarters (e.g., Q1), particularly at moderate risk levels (α = 0.5 or 0.75), which may reflect improved downside resilience driven by sustainability-focused asset selection. Conversely, in aggressive portfolios (α = 1.0), ESG filtering did not prevent steep losses, indicating that the benefits of ESG integration may diminish when full exposure to risky assets is permitted.

When compared to the minimum-risk strategy, the One-Fund model offers no clear performance advantage but does allow for tailoring portfolio outcomes based on individual investor preferences. This may be especially relevant in SME contexts where decision-makers have varying appetites for volatility but limited resources for repeated optimization.

4.4.3. Summary of Findings

The One-Fund model, when applied to historical Sharpe-optimized portfolios, provides a flexible framework for adjusting risk exposure without modifying the underlying optimization structure. While low-α strategies deliver stability, high-α strategies introduce significant return swings, both positive and negative, particularly in volatile quarters.

Moderate values of α (e.g., 0.5-0.75) appear to have a reasonable balance, offering improved returns over the risk-free baseline without reaching the instability of full-risk positions. ESG integration offers slight improvements in downside protection but does not radically shift the overall risk-return landscape.

These findings suggest that for SMEs with moderate growth objectives and some tolerance for risk, the One-Fund model could serve as a pragmatic extension to a minimum-risk approach, especially when paired with transparent historical return estimates.

While the One-Fund approach offers useful flexibility across different investor preferences, its effectiveness relies entirely on the underlying Sharpe-optimized portfolio. In cases where return forecasts are unstable, as seen in the CAPM model, this strategy inherits the same weaknesses. For SMEs, it remains a valuable extension only when based on stable and interpretable input models.

4.5. ESG Threshold Sensitivity Analysis

This section examines how different ESG threshold levels affect portfolio performance when using the historical minimum-variance model. The ESG constraint is increased step-by-step from 60 to 83, while all other model parameters remain unchanged. This allows for a controlled comparison of how tighter sustainability requirements influence returns, risk, and risk-adjusted performance.

Each threshold level is applied during rolling OOS backtesting from 2020 to 2024. For every case, the realized return, annualized volatility, Sharpe ratio, and achieved average

ESG score are recorded. This setup ensures that any performance change can be linked to the ESG filter.

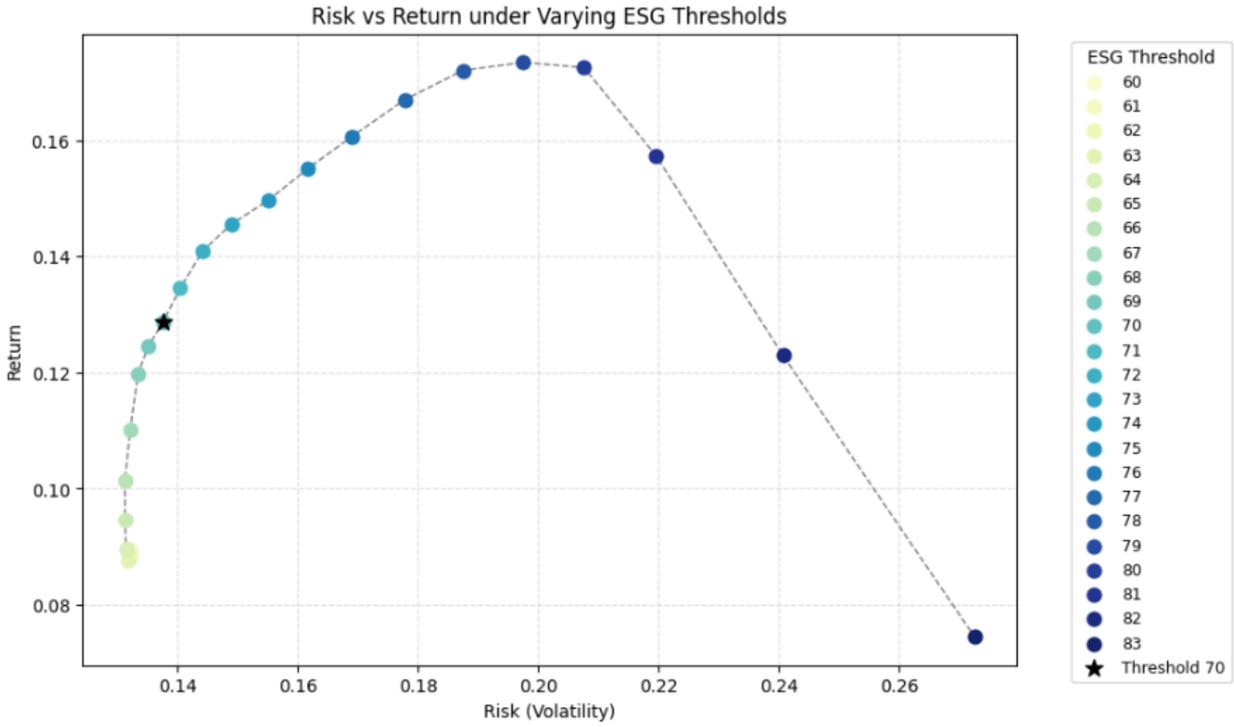

Figure 2: ESG Threshold Sensitivity Analysis

The results show that increasing the ESG threshold from 60 to around 70 improves return without a notable rise in volatility. However, thresholds above 72 lead to sharply higher volatility without additional return benefits. The most extreme case, with a threshold of 83, results in the highest volatility (27.3%) and the lowest return (7.4%).

In contrast, thresholds between 68 and 72 show the most balanced outcomes. Portfolios in this range achieve relatively high returns (up to 17.2%) while keeping volatility moderate. Sharpe ratios remain close to or above 0.9, which indicates efficient risk-adjusted performance. The threshold of 70, marked in the plot, lies near this efficient zone.

These results suggest that while a basic ESG constraint can improve portfolio quality, pushing the threshold too far reduces diversification and leads to weaker outcomes. For investors with limited flexibility like SMEs a threshold in the 68-72 range appears to offer a reasonable balance between sustainability goals and portfolio stability for this particular set of assets.

4.6. Summary of Results

This section brings together the key findings across all 12 portfolio optimization models and the One-Fund extension. The goal is to highlight which strategies performed best under different constraints, assess the impact of ESG integration, and identify practical trade-offs that matter for implementation, especially in SME settings.

To support the comparison, Table 4 summarizes performance across the main models using consistent metrics: realized return, Sharpe ratio, volatility, turnover, beta, and forecast accuracy (MAE between expected and realized returns and Sharpe ratios).

Table 5: Summary of Results

Strategy	*ESG*	*Return*	*Volatility*	*Sharpe*	*Sortino*	*MAE (Return)*	*MAE (Sharpe)*	*Beta*	*Turnover*
Historical Min Risk	No	8.9%	13.2%	0.44	0.73	0.22	1.79	0.50	0.91
Historical Min Risk	Yes	12.8%	13.8%	0.80	1.47	0.25	1.92	0.61	0.87
Historical Max Sharpe	No	9.5%	18.7%	0.51	1.62	0.58	3.59	0.74	1.32
Historical Max Sharpe	Yes	9.5%	18.9%	0.50	1.59	0.58	3.50	0.78	1.35
EWMA Min Risk	No	8.0%	13.1%	0.46	0.72	0.23	1.75	0.58	0.68
EWMA Min Risk	Yes	10.5%	14.1%	0.64	1.07	0.23	1.58	0.70	0.54
EWMA Max Sharpe	No	10.0%	18.5%	0.54	1.62	0.56	3.69	0.75	1.35
Max Sharpe	Yes	10.5%	18.9%	0.55	1.70	0.58	3.68	0.79	1.41
CAPM Min Risk	No	8.9%	13.2%	0.44	0.73	0.20	1.63	0.50	0.91
CAPM Min Risk	Yes	12.9%	13.8%	0.80	1.47	0.21	1.56	0.61	0.87

CAPM Max Sharpe	No	2.8%	19.0%	0.24	1.02	0.35	1.56	0.93	1.11
CAPM Max Sharpe	Yes	4.8%	20.1%	0.27	0.89	0.0.34	1.41	0.99	1.07

4.6.1. Best Overall Performance

The Historical Min-Risk ESG model consistently achieved the best balance between return and risk. It reached a realized return of 12.9%, a Sharpe ratio of 0.81, and kept volatility at a moderate 13.8%. Turnover remained relatively low (0.87), which makes this strategy not only effective but also practical in environments with limited capacity for frequent rebalancing. Forecasting error was slightly higher than in the unconstrained version, but well within acceptable bounds.

4.6.2. Instability in Sharpe-Maximizing Strategies

Sharpe-maximizing models showed significantly more volatility and instability in performance. While they were designed to maximize risk-adjusted return, they were extremely sensitive to return forecasts. For example, the CAPM Sharpe ESG variant produced high expected returns but delivered a weak realized Sharpe ratio of just 0.27, with volatility over 20%. These findings support the view that Sharpe-max optimization often leads to unstable weights and overfitting, especially when expected returns are difficult to estimate reliably.

4.6.3. Impact of ESG Integration

Across most models, ESG-constrained portfolios outperformed or matched their unconstrained counterparts. This was especially true in the minimum-risk strategies, where ESG integration not only preserved returns but slightly improved risk-adjusted outcomes. In contrast, Sharpe-maximizing ESG portfolios showed mixed results, with only modest gains in some setups and no significant improvement in others. Importantly, ESG constraints did not increase volatility or turnover significantly, making them viable even under limited rebalancing capacity.

4.6.4. Forecast Accuracy and Turnover

Forecasting accuracy, measured through mean absolute error between expected and realized metrics, varied across models. CAPM and EWMA Sharpe models showed the largest gaps, while Historical Min-Risk variants were more stable. This confirms that models with more complex or unstable return forecasts may not offer consistent benefits for SMEs.

Turnover was also an important differentiator. Sharpe-maximizing strategies generally had higher turnover (1.3-1.4) than min-risk models (0.5-0.9), which may make them less suitable in real-world scenarios where transaction costs or resource limitations are a concern.

4.6.5. One-Fund Strategy and Investor Preferences

The One-Fund model, based on the historical Sharpe-optimized portfolio, allowed for adjusting risk exposure without re-optimizing weights. Results showed that α = 0.5 or 0.75 provided a good compromise, avoiding the sharp losses seen in fully risky portfolios (α = 1.0), especially in volatile quarters. The heatmap confirms that moderate risk levels delivered more stable and positive returns, particularly in the second half of 2024.

This approach can be especially useful for investors seeking flexibility but lacking the tools or data to build fully customized portfolios. However, the One-Fund model still relies on a stable underlying Sharpe-optimized portfolio, which is not always achievable in unstable forecasting environments (as seen in CAPM).

Taken together, the findings confirm that ESG integration does not inherently diminish performance. Instead, the modelling framework and return assumptions are more decisive. Stable, low-risk strategies with transparent estimation methods consistently produced the best results across models, especially when ESG constraints were moderately applied.

5. Discussion

This section analyses the findings from the previous chapter in accordance with the goals of this thesis, real-world applications, and existing scholarly research. It seeks to question the reasons behind the performance differences among models, the impact of incorporating ESG factors, and the implications for SME investment strategies. The discussion is structured around the original research questions, then looks at real-world effects, compares with previous studies, and critically reviews the research methods used. By contextualizing the results within the broader academic landscape and industry practices, this chapter aims to provide a comprehensive understanding of the contributions and limitations of this thesis.

5.1. Revisiting the Research Questions

This chapter analyses the empirical results of portfolio optimization strategies implemented in this thesis in connection with the four research questions presented at the beginning of the research. The findings from all models are compared to address each question, focusing on technical explanations for the observed variations and their implications for SMEs. The interpretation aims to provide an understanding of the data and its relevance to the target audience, offering insights into the practical applications of the research outcomes for SMEs in the current business landscape.

RQ 1: Which portfolio optimization approach offers the most transparent and implementable framework for SMEs with limited analytical and data resources?

The results show that the historical minimum-variance model with ESG constraints was the most consistently reliable approach across all periods. This strategy delivered the highest Sharpe ratio (0.81), strong realized returns (12.9%), and moderate volatility (13.8%) with low turnover (0.87). The strength of this approach lies in several key factors from a technical perspective. First, the optimization problem's convex structure provides a solid foundation. Additionally, the method relies on historical return estimates, which offer stability and abundant data points. Finally, the minimum-variance objective demonstrates resilience against short-term fluctuations, contributing to the overall robustness of the system. These elements combine to create a reliable and effective framework for analysis and decision-making in this context.

The historical minimum-risk approach offers a simpler alternative to more sophisticated models. It relies on basic inputs like past returns and a covariance matrix, making it accessible to SMEs that may lack advanced forecasting capabilities or extensive resources. This method's straightforward implementation is particularly beneficial for companies with limited access to complex financial tools or specialized expertise. Additionally, the lower trading frequency associated with this approach helps reduce operational costs and complexity, which is especially valuable for businesses operating with constrained human or financial resources. Overall, this strategy provides a practical solution for organizations seeking to manage risk without the need for elaborate systems or substantial investments in infrastructure.

RQ 2: How do different return estimation methods, such as historical averages and market-based models, affect the reliability of portfolio performance in OOS testing?

The comparison across models revealed that forward-looking estimators such as EWMA and CAPM did not improve OOS performance. Instead, they often introduced instability, particularly in Sharpe-maximizing strategies. The CAPM-based models showed the weakest results, with the ESG Sharpe-max variant achieving a Sharpe ratio of only 0.27 and a realized return of 4.9%, despite high expected return forecasts.

This can be explained by the structure of CAPM itself: expected returns are a function of the market premium and estimated β, both of which were derived from rolling windows. The β estimates were unstable due to the short training period (63 days), and the DAX index, used as a market proxy, exhibited its own volatility and sector biases. As a result, CAPM-based return estimates frequently over- or underreacted to recent market movements, creating extreme and unreliable portfolio weights.

The EWMA model, despite its theoretical adaptability, did not perform better than the basic historical model. Using exponentially weighted returns made the model more responsive, but it did not improve its predictive capabilities. Additionally, this approach led to higher turnover and increased complexity without significantly enhancing the Sharpe ratio, a key measure of risk-adjusted performance.

These results highlight the challenges of forecasting expected returns with limited data. For strategies aligned with SMEs, maintaining consistency and ease of understanding often takes precedence over pursuing theoretical precision. The findings underscore the importance of simplicity and reliability in financial modelling, especially when dealing with smaller datasets and more focused investment strategies.

RQ 3: In what ways do static versus dynamic risk estimators influence the interpretability and robustness of ESG-integrated portfolios?

Research shows that using static risk assessment with past data made portfolios easier to understand and more stable than dynamic methods like EWMA. The historical minimum-variance model consistently delivered stable results with moderate volatility and low turnover, even when ESG constraints were applied. This can be explained by the static covariance structure's ability to capture enduring relationships between asset returns, which remain relatively constant over time and are less susceptible to short-term fluctuations.

On the other hand, the EWMA-based model employed dynamic weighting of recent return data, aiming to adapt more quickly to changing market conditions. However, this adaptability did not result in improved OOS performance. While volatility estimates were updated more frequently, the resulting portfolio weights were more reactive and less predictable. This negatively impacted both interpretability, as weight changes became harder to explain, and robustness, as minor market fluctuations often triggered significant rebalancing. The EWMA Sharpe-maximizing strategy exhibited higher turnover and a greater discrepancy between expected and realized performance.

From a technical standpoint, dynamic models like EWMA require additional parameter tuning (e.g., decay factor λ), which introduces more degrees of freedom and potential overfitting. SMEs generally prefer straightforward approaches that are easy to understand and implement. In this context, static estimators provide more predictable outcomes and are less affected by short-term market volatility. The results show that static models are usually the better and more dependable choice for building ESG-integrated portfolios unless there is a strong reason to use dynamic risk adjustment methods. This approach aligns well with the needs of SMEs, who often lack the resources or expertise to manage more complex, dynamic models effectively.

RQ 4: How does a One-Fund strategy compare to minimum-risk portfolios under ESG constraints, and does it offer meaningful advantages for investors with varying levels of risk tolerance?

The One-Fund strategy combines a risk-free asset with a Sharpe-optimized portfolio, allowing investors to adjust risk exposure without completely restructuring their portfolios. While theoretically solid, its real-world effectiveness largely depends on the quality and consistency of the underlying portfolio.

This thesis applied the One-Fund approach to a historical Sharpe-maximizing model, chosen for its transparency and accessible return estimates. Findings revealed that high-risk settings (α = 1.0) resulted in unpredictable and fluctuating returns, especially in early 2024. In contrast, more moderate α values (0.5-0.75) yielded balanced outcomes, with returns and Sharpe ratios comparable to the ESG-constrained minimum-risk model.

These results suggest that while the One-Fund model does not necessarily surpass the minimum-risk strategy, it can achieve similar performance at moderate risk levels. For investors with specific risk preferences or regulatory constraints, this approach offers a valuable method to adjust risk exposure without building new portfolios from the ground up.

However, the model's effectiveness is constrained by the stability of the underlying Sharpe-optimized portfolio. When this foundation is volatile, as seen in CAPM or EWMA models, the One-Fund results can become equally unreliable. Thus, the One-Fund model's advantages are more accurate when based on a stable, well-performing core portfolio. For SME investors, it can serve as a flexible risk adjustment tool, provided the base model is carefully chosen and regularly evaluated.

5.2. Practical Implications for SMEs

SMEs looking to incorporate sustainability into their investment strategies can benefit from several key insights derived from this research. These findings primarily concern model selection, user-friendliness, and practical ESG implementation.

Historical minimum-variance models emerge as the most suitable option for SMEs. These models require only past return data and can be implemented using basic tools like Excel and Python. Their consistent performance, low trading frequency, and straightforward logic make them ideal for businesses with limited resources and time for frequent rebalancing or complex analysis.

While more sophisticated models like CAPM and EWMA may seem attractive theoretically, they often produce unreliable results when applied to small datasets or short timeframes. This can lead to confusion and mistrust, particularly for those without extensive financial expertise. In many instances, these complex approaches may hinder rather than help decision-making. The research indicates that simpler models are not only more manageable but often more effective.

Incorporating ESG constraints into minimum-risk models did not negatively impact performance and, in some cases, even slightly improved returns. This challenges the notion that sustainability considerations always come at the expense of financial performance. A simple rule, such as requiring the portfolio's average ESG score to exceed a certain threshold, proved easy to implement and did not cause instability, even with a limited stock selection. This suggests that SMEs can apply basic ESG filters using freely available ratings without the need for expensive software.

The One-Fund approach offers a flexible method for adjusting portfolio risk based on investor preferences. While it may not outperform the lowest-risk models in absolute terms, it provides adaptability that could be valuable for SMEs with diverse decision-maker preferences. This approach is relatively straightforward to implement without requiring multiple optimization iterations.

Nevertheless, it is necessary to pick an appropriate ESG threshold, as its increase does not necessarily mean additional increase in performance. Sensitivity analysis of ESG thresholds revealed that setting overly strict criteria could compromise diversification and lead to suboptimal portfolios. Thresholds in the 68-72 range appeared to offer the best balance in this analysis. This indicates that ESG should be treated as a flexible parameter rather than a rigid criterion, allowing for fine-tuning based on specific goals and available options. However, this range can differ depending on the companies in the investment universe.

In conclusion, the research demonstrates that SMEs can develop ESG-aware portfolios without sophisticated tools or expert guidance. By employing basic models, carefully adjusting ESG thresholds, and utilizing publicly available data, they can create investment strategies that are both sustainable and practical.

5.3. Literature Comparison

The research findings in this thesis align with and expand upon previous studies in sustainable investing and portfolio optimization. The observed positive effects of incorporating ESG factors on risk-adjusted returns, especially in minimum-variance strategies, support earlier research by Whelan et al. (2021) highlighting the long-term performance benefits of ESG integration. While most existing studies focus on large institutional investors, this research demonstrates that similar advantages can be achieved in smaller portfolios using simplified, accessible methods suitable for SMEs.

The performance of ESG-constrained portfolios in this research relates to the work of Pedersen et al. (2021), who describe how ESG criteria can alter the efficient frontier. This effect was evident in the results, particularly in the historical model, where ESG constraints often led to portfolios with improved Sharpe ratios and enhanced downside protection without increasing volatility or turnover. These results show that ESG-aligned portfolios can do better than traditional ones in certain situations, highlighting the increasing significance of sustainable investing strategies.

By contrast, this thesis challenges concerns raised by Zhou (2024) and Abate et al. (2024), who suggest that ESG integration may lead to increased turnover and complexity. In this analysis, the ESG constraint had no destabilizing effect in the minimum-risk strategies and did not significantly increase rebalancing needs. This can be attributed to precise implementation of ESG constraints using a straightforward linear inequality and a composite ESG score. This approach also addresses concerns raised by Berg et al. (2022) about inconsistencies between different ESG rating providers. Nonetheless, a more extensive analysis on bigger investment universe should be performed to support or to neglect the concern about increased turnover and complexity caused by ESG constrains.

The mediocre performance of the CAPM- and EWMA-based strategies reinforces existing research warning against overreliance on unstable return forecasts. For example, Guidolin et al. (2018) note that forward-looking models can underperform when based on short time windows or volatile market proxies, which was the case here with the DAX as a benchmark. The CAPM-based models showed especially high sensitivity to estimation errors in beta and market returns, resulting in poor realized Sharpe ratios and unstable asset weights.

This thesis presents a practical approach to ESG portfolio optimization tailored for SMEs. It addresses the concerns raised by recent studies, particularly those of Chien (2021) and Whelan (2021), who noted a lack of accessible ESG tools for smaller businesses. Unlike many complex, data-heavy solutions found in academic literature, this research offers a straightforward and transparent modelling method. The findings demonstrate that ESG portfolio optimization can be effectively implemented by firms with limited resources, challenging the notion that such strategies are only useful for large corporations. By carefully selecting and applying appropriate models, SMEs can successfully integrate ESG factors into their investment decisions, bridging the gap between theoretical concepts and real-world applications in sustainable finance.

5.4. Methodological and Practical Limitations

This thesis presents strategies designed for transparency and SME suitability, but it is important to acknowledge its limitations when considering broader applications.

The portfolio models underwent evaluation through historical backtesting, employing rolling OOS windows. While this approach enhances robustness compared to static performance snapshots, it falls short of fully replicating real-time uncertainties, structural market shifts, and behavioural factors that impact decision-making processes. As a result, the findings are inherently linked to the particular market conditions and asset behaviours observed during the 2020-2024 timeframe. These constraints should be considered when assessing the framework's overall effectiveness and potential for wider implementation across various contexts and time periods.

The ESG integration process utilized static scores, calculated by averaging ratings from three providers. This method addressed discrepancies between scores and simplified implementation but failed to capture evolving company sustainability profiles. Remarkable events like ESG controversies, improvements in sustainability practices, or policy changes were not reflected, potentially limiting the accuracy of ESG-based constraints over time. In fact, the constant nature of the ESG scores, as the data is either not publicly available or has not been recorded by the providers, decreases the number of possible strategies which can be tested by SMEs. One possibility would be to integrate the ESG score into the objective function near return or risk, however the plausibility of this strategy is questionable when using static ESG scores, as they do not depict the performance evolution through the time.

When looking at different ways to estimate returns and risks, the results showed that advanced models like CAPM and EWMA did not always perform better than the simpler historical benchmark. The increased sensitivity and parameter dependence of these complex models led to greater fluctuations in portfolio weights without clear performance benefits. These results highlight the well-known challenges of predicting returns with limited data and suggest that increased complexity does not necessarily lead to improved outcomes for investors with constrained modelling resources.

The optimization process did not account for real-world factors like transaction costs, slippage, and taxes. Even though turnover metrics were used to evaluate how practical the strategies were, not considering costs might have resulted in an inflated view of profits for high-frequency strategies, particularly those aimed at maximizing the Sharpe ratio.

The analysis in this thesis focused solely on stocks listed on the XETRA exchange. This choice ensured data consistency and relevance to European SMEs but limited the diversity of sectors and geographical regions. As a result, the findings may not apply to global multi-asset portfolios without further research. Similarly, the ESG threshold analysis was specific to this dataset's distribution of ESG scores and might not be applicable in other regions or for different investment objectives.

The One-Fund approach was ultimately applied solely to the portfolio optimized using historical Sharpe ratios. Expanding this method to incorporate additional models remains an area for future research, especially in more turbulent market conditions or during periods of significant economic shifts.

While these limitations exist, the choices made in developing the model reflect the real-world challenges faced by SMEs with restricted analytical resources. By prioritizing simplicity, clarity, and ease of implementation, the model caters to the practical needs of its intended users. The emphasis on simplicity, transparency, and operational feasibility aligns with the practical context of the target users and supports the overall credibility of the findings within that scope.

6. Conclusion

This thesis explored how to incorporate ESG factors into investment strategies for smaller companies with limited resources. The research compared various portfolio optimization models under different ESG scenarios to find approaches that balance financial returns, sustainability objectives, and practical implementation.

The results indicate that including ESG criteria does not necessarily hurt performance and can sometimes enhance risk-adjusted returns. This was particularly noticeable in minimum-variance strategies, where ESG constraints produced well-diversified portfolios without increasing volatility or frequent trading. In contrast, models based on return predictions, like those maximizing the Sharpe ratio or using the Capital Asset Pricing Model, often led to unstable portfolios and were more susceptible to estimation errors. These findings underscore the value of simple, robust models over complex theoretical approaches, especially for organizations lacking extensive data or sophisticated forecasting capabilities.

This thesis offers more than just performance benchmarks - it provides a practical approach for SMEs to incorporate sustainability into their investment strategies. By utilizing static ESG ratings, rolling-window backtesting methods, and constraint-based modelling techniques, the framework can be implemented using readily available tools and clear, understandable logic. Additionally, the sensitivity analysis of ESG thresholds revealed that sustainability constraints should be carefully calibrated rather than applied as a one-size-fits-all solution. This finding highlights the importance of finding a balance, as too strict ESG criteria may increase portfolio risk and reduce diversification benefits.

The conducted research has some constraints to consider, such as relying on fixed ESG information, not accounting for trading expenses, and concentrating on just one regional stock market. However, the approach taken was designed to address real-world challenges faced by the target audience. Moving forward, researchers could expand on this foundation by integrating evolving ESG indicators, broadening the range of assets examined, or delving deeper into individual investor priorities. This could be especially relevant when developing comprehensive investment solutions or tackling complex optimization problems with multiple objectives.

Overall, this thesis shows that incorporating ESG factors into portfolio optimization is feasible for SMEs, not just large institutions. By choosing appropriate models and implementing them clearly, smaller companies can create investment strategies that balance financial goals with sustainability objectives. The process does not have to be overly complex, making it easier for a wider range of businesses to engage in responsible investing practices. This approach helps expand the reach and impact of sustainable investing beyond just major corporations and institutional investors.

Bibliography

Abate, G., Basile, I., & Ferrari, P. (2024). The integration of environmental, social and governance criteria in portfolio optimization: An empirical analysis. *Corporate Social Responsibility and Environmental Management, 31*(3), 2054-2065.

Ahmad, H., Yaqub, M., & Lee, S. H. (2024). Environmental-, social-, and governance-related factors for business investment and sustainability: A scientometric review of global trends. *Environment, Development and Sustainability, 26*(2), 2965-2987.

Ban, G.-Y., Karoui, N. E., & Lim, A. E. (2018). Machine Learning and Portfolio Optimization. *Management Science, 64*(3), 1136-1154.

Bartholdy, J., & Peare, P. (2003). Unbiased Estimation of Expected Return Using CAPM. *International Review of Financial Analysis, 12*(5´1), 69-81.

Berg, F., Kölbel, J., & Rigobon, R. (2022). Aggregate Confusion: The Divergence of ESG Ratings. *Review of Finance, 6*, pp. 1315-1344.

Cesarone, F., Martino, M. L., & Carleo, A. (2022). Does ESG Impact Really Enhance Portfolio Profitability? *Sustainability, 14*(4), 2050.

Chen, L., Zhang, L., Huang, J., Xiao, H., & Zhou, Z. (2021). Social responsibility portfolio optimization incorporating ESG criteria. *Journal of Management Science and Engineering, 6*(1), 75-85.

Chien, F., Ngo, Q.-T., Hsu, C.-C., Chau, Yin, K., & Iram, R. (2021). Assessing the mechanism of barriers towards green finance. *Environmental Science and Pollution Research, 28*, 60495-60510.

Cornuejols, G., Pena, J., & Tütüncü, R. (2018). *Optimization Methods in Finance.* Cambridge University Press.

Diamond, S., & Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research, 17*(83), 1-5. Retrieved from http://jmlr.org/papers/v17/15-408.html

Drei, A., Le Guenedal, T., Lepetit, F., Mortier, V., Roncalli, T., & Sekine, T. (2019). ESG investing in recent years: New insights from old challenges. *Amundi Asset Management*.

Eccles, R. G., Ioannou, I., & Serafeim, G. H. (2014). The impact of corporate sustainability on organizational processes and performance. *Management Science, 60*(11), 2835-2857.

Elbannan, M. A. (2015). The Capital Asset Pricing Model: An Overview of the Theory. *International Journal of Economics and Finance, 7*(1), 216-228.

Escobar-Anel, M., & Jiao, Y. (2024). Robust Portfolio Optimization with Environmental, Social, and Corporate Governance Preference. *Risks, 12*(33).

European Commission. (2023). Corporate sustainability reporting directive (CSRD): Strengthening the foundations of sustainable investment. Official Journal of the European Union.

European Parliament & Council. (2020). Regulation (EU) 2020/852 of the European Parliament and of the Council of 18 June 2020 on the establishment of a framework to facilitate sustainable investment, and amending Regulation (EU) 2019/2088. Official Journal of the European Union.

Fama, E. F., & French, K. R. (2004). The Capital Asset Pricing Model: Theory and Evidence. *Journal of Economic Perspectives, 18*(3), 25–46.

Fletcher, J., & Hillier, D. (2001). An Examination of Resampled Portfolio Efficiency. *Financial Analysts Journal, 57*(5), 66-74.

Garrido-Ruso, M., Otero-Gonzalez, L., & Lopez-Penabad, M.-C. (2024). Does ESG implementation influence performance and risk. *Corporate Social Responsibility and Environmental Management, 31*(5), 4227-4247.

Guidolin, M., Hansen, E., & Lozano-Bandaz, M. (2018). Portfolio Performance of Linear SDF Models: An Out-of-Sample Assessment. *IGIER Working Paper No. 627*. Milano: Università Bocconi.

Hancock, A. (2024). *Is red tape strangling Europe's growth?* Retrieved March 2025, from Financial Times: https://www.ft.com/content/4e8e6cde-d0ce-4f0a-a7ea-1c913d4dad50

Hibiki, N. (2006). Multi-period stochastic optimization models. *Journal of Banking & Finance, 30*, 365-390.

Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 679-688.

J.P. Morgan. (1996). *RiskMetrics Technical Document.* New York.

Jung, Y. L., & Yoo, H. S. (2023). Environmental, social, and governance activities and firm performance: Global evidence. *Journal of Business Ethics, 30*(6), 2830-2839.

Khan, M., Serafeim, G., Yoon, A. S., & Suh, J. (2016). Corporate sustainability: First evidence on materiality. *The Accounting Review, 91*(6), 1697-1724.

Kresta, A., & Wang, A. (2020). Portfolio Optimization Efficiency Test Considering Data Snooping Bias. *Business Systems Research, 11*(2), 73-85.

Kuen, T. Y., & Hoong, T. S. (1992). Forecasting Volatility in the Singapore Stock Market. *Asia Pacific Journal of Management, 9*(1), 1-13.

Li, D., & Zhu, K. (2020). Inference for Asymmetric Exponentially Weighted Moving Average Models. *Journal of Time Series Analysis, 41*(1), 154-162.

Liagkouras, K., Metaxiotis, K., & Tsihrintzis, G. (2022). Incorporating environmental and social considerations. *Annals of Operations Research, 316*, 1493-1518.

Liu, W., & Jing, K. (2024). ESG portfolio for TDFs with time-varying higher moments and cardinality constraint. *International Transactions in Operational Research, 31*(6), 4270-4295.

London Stock Exchange Group. (n.d.). *ESG scores and sustainable investment data*. Retrieved May 2025, from https://www.lseg.com/en/data-analytics/sustainable-finance/esg-scores

Luke, T. W. (2022). Investment and Rapid Climate Change as Biopolitics:. *Sustainability, 14*(22).

MacNeil, I., & Esser, I.-M. (2022). From a Financial to an Entity Model of ESG. *European Business Organization Law Review, 23*, 9-45.

Mangram, M. E. (2013). A Simplified Perspective of the Markowitz Portfolio Theory. *7*(1), pp. 59-70.

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance, 7*(1), 77-91.

Martinez-Cillero, M., Lawless, M., O'Toole, C., & Slaymaker, R. (2020). Financial frictions and the SME investment gap: new survey evidence for Ireland. *Venture Capital, 22*(3), 239-259.

Miranda-García, M., & Segovia-Vargas, M.-J. (2024). Financial constraints and sustainability in bioeconomy. *Global Policy*, 1-18.

OECD. (2020). ESG Investing: Environmental Pillar Scoring and Reporting. OECD Publishing.

Orzechowski, A., & Bombol, M. (2022). Energy Security, Sustainable Development and the Green. *Energies, 15*(6218).

Patak, T., Tewari, R., & Drempetic, S. (2024). Mandatory CSR regime strips the competitive advantage: A comparative study of pre-post CSR mandate using the Bandwagon-bias effect theory. *Society and Business Review, 19*(4), 695-716.

Pedersen, L. H., Fitzgibbons, S., & Pomorski, L. (2021). Responsible Investing: The ESG-Efficient Frontier. *Journal of Financial Economics, 142*(2), 572-597.

S&P Global. (n.d.). *ESG scores and data solutions*. Retrieved May 2025, from https://www.spglobal.com/esg/solutions/esg-scores-data

Sarisoy, C., de Goeij, P., & Werker, B. J. (2024). Linear Factor Models and the Estimation of Expected Returns. Washington: Finance and Economics Discussion Se- ries.

Shalhoob, H., & Hussainey, K. (2023). Environmental, Social and Governance (ESG) Disclosure and the Small and Medium Enterprises (SMEs) Sustainability Performance. *Sustainability, 15*(1).

Sharpe, W. F. (1994). The Sharpe Ratio. *THe Journal of Portfolio Management*.

Silva, T., Pinheiro, P. R., & Poggi, M. (2017). A more human-like portfolio optimization approach. *European Journal of Operational Research, 256*, 252-260.

Son, E., & Suh, J. (2024). The role of entrepreneurship in alleviating ESG backlash and advancing sustainability. *Journal of the International Council for Small Business.*, 1-19.

Sortino, F. A., & Van Der Meer, H. J. (1991). Downside risk. *Journal of Portfolio Mangement, 17*(4), 27-31.

Spigeleer, J. D., Höcht, S., Jakubowski, D., Reyners, S., & Schoutens, W. (2023). ESG: a new dimension in portfolio allocation. *Journal of Sustainable Finance, 13*(2), 827-867.

Sustainalytics. (n.d.). *ESG risk ratings*. Retrieved May 2025, from https://www.sustainalytics.com/esg-ratings

Varmaz, A., Fieberg, C., & Poddig, T. (2024). Portfolio optimization for sustainable investments. *Annals of Operations Research, 341*, 1151–1176.

Wagemans, F. A., van Koppen, K., & Mol, A. P. (2018). Engagement on ESG issues by Dutch pension funds: is it. *Journal of, 8*(4), 301-322.

Wanday, J., & Ajour El Zein, S. (2022). Higher expected returns for investors in the energy sector in Europe using an ESG strategy. *Frontiers in Environmental Science, 10*.

Whelan, T., Atz, U., Van Holt, T., & Clark, C. (2021). ESG and Financial Performance: Uncovering the Relationship by Aggregating Evidence from 1,000 Plus Studies Published between 2015–2020. NYU Stern Center for Sustainable Business and Rockefeller Asset Management.

Xetra. (n.d.). *Xetra – The market for fast and secure trading. Deutsche Börse Group.* Retrieved May 2025, from https://www.xetra.com/xetra-en/

Yahoo Finance. (n.d.). *Yahoo Finance*. Retrieved May 2025, from https://finance.yahoo.com/

Zeidan, R. (2022). Obstacles to sustainable finance and the covid19 crisis. *Journal of Sustainable Finance & Investment, 12*(2), 525-528.

Zhang, K., & Hao, X. (2024). Corporate Social Responsibility as the Pathway Towards Sustainability: A State-of-the-Art Review. *Discover Sustainability, 5*(348).

Zhang, Y., Li, X., & Guo, S. (2018). Portfolio selection problems with Markowitz's mean–variance framework: a review of literature. *Fuzzy Optimization and Decision Making, 17*, 125-128.

Zhou, X. (2024). From Theory to Practice: Applying Markowitz Model in Stock Portfolio Management under ESG. *International Journal of Global Economics and Management, 2*.

Appendix

Appendix 1: Data Loading and Preprocessing

```
# Load price data
price_data = pd.read_csv(f"{base_path}/xetra_price_data.csv", index_col=0, parse_dates=True)

# Calculate log returns
daily_returns = np.log(price_data / price_data.shift(1)).dropna()

# Filter time range
start_date = "2020-01-01"
end_date = "2024-12-31"
daily_returns = daily_returns.loc[start_date:end_date]
price_data = price_data.loc[start_date:end_date]
```

Appendix 2: ESG Score Loading and Weighted Composite Calculation

```
# Load and parse ESG scores
esg_raw = pd.read_csv(esg_data_path, sep=";", index_col=0)

# Extract and align ESG scores by Ticker
lseg = esg_raw.loc["LSEG", ["Ticker", "Overall Score"]].set_index("Ticker").rename(columns={"Overall Score": "LSEG"})
sp = esg_raw.loc["S&P", ["Ticker", "Overall Score"]].set_index("Ticker").rename(columns={"Overall Score": "S&P"})
sus = esg_raw.loc["Sustainalytics", ["Ticker", "Overall Score"]].set_index("Ticker").rename(columns={"Overall Score": "Sustainalytics"})

# Merge scores into a single DataFrame
esg_scores = lseg.join(sp).join(sus)
esg_scores = esg_scores.apply(lambda col: np.vectorize(lambda x: float(str(x).replace(",", ".")))(col))

# Weighted ESG score: 50% LSEG, 30% S&P, 20% Sustainalytics
esg_weighted = (
    0.5 * esg_scores["LSEG"] +
    0.3 * esg_scores["S&P"] +
    0.2 * esg_scores["Sustainalytics"]
)
```

Appendix 3: Historical Min Risk (No ESG) - Optimization and Backtesting

```
# Optimization
def optimize_portfolio(mu, Sigma):
    mu_np = mu.values
    Sigma_np = Sigma.values
    n = len(mu_np)
    w = cp.Variable(n)
    risk = cp.quad_form(w, Sigma_np)
    constraints = [cp.sum(w) == 1, w >= 0]
    prob = cp.Problem(cp.Minimize(risk), constraints)
    prob.solve()
    return np.array(w.value).flatten() if prob.status == 'optimal' else np.ones(n) / n

# Backtest single quarter
def backtest_single_quarter(train_returns, test_returns, R_f=0.02):
    mu = train_returns.mean() * 252
    Sigma = train_returns.cov() * 252
    weights = optimize_portfolio(mu, Sigma)

    daily_test_returns = test_returns @ weights
    realized_return = np.mean(daily_test_returns) * 252
    realized_std = np.std(daily_test_returns) * np.sqrt(252)
    sharpe_ratio = (realized_return - R_f) / realized_std if realized_std != 0 else 0

    expected_return = mu @ weights
    expected_volatility = np.sqrt(weights.T @ Sigma @ weights)

    return {
        'start': test_returns.index[0],
        'end': test_returns.index[-1],
        'expected_return': expected_return,
        'expected_volatility': expected_volatility,
        'return': realized_return,
        'risk': realized_std,
        'sharpe': sharpe_ratio,
        'weights': weights
    }

# Rolling backtest
```

```
def run_rolling_backtest(returns, train_days=63, test_days=63):
    results = []
    all_dates = returns.index
    max_index = len(all_dates)
    i = 0
    while (i + train_days + test_days) <= max_index:
        train_returns = returns.iloc[i:i+train_days]
        test_returns = returns.iloc[i+train_days:i+train_days+test_days]

        if train_returns.index[0] < pd.Timestamp("2020-01-01"):
            i += test_days
            continue
        if test_returns.index[-1] > pd.Timestamp("2024-12-31"):
            break

        result = backtest_single_quarter(train_returns, test_returns)
        results.append(result)
        i += test_days

    return pd.DataFrame(results)
```

Appendix 4: Historical Max Sharpe (No ESG) - Optimization and Backtesting

```
# Objective function
def neg_excess_sharpe_ratio(w):
    port_return = np.dot(w, mu)
    port_vol = np.sqrt(np.dot(w.T, np.dot(cov, w)))
    excess_return = port_return - (risk_free_rate / annualization_factor)
    return -excess_return / port_vol if port_vol != 0 else np.inf

# Rolling Backtest Loop
for start in range(0, len(daily_returns) - train_window - test_window + 1, test_window):
    train_returns = daily_returns.iloc[start : start + train_window].clip(lower=-0.1, upper=0.1)
    test_returns = daily_returns.iloc[start + train_window : start + train_window + test_window]

    mu = train_returns.mean()
    cov = train_returns.cov()
    tickers_used = mu.index.tolist()

    n = len(mu)
    w0 = np.ones(n) / n
    bounds = [(0, 0.2) for _ in range(n)]
    constraints = [{'type': 'eq', 'fun': lambda w: np.sum(w) - 1}]
    result = minimize(neg_excess_sharpe_ratio, w0, bounds=bounds, constraints=constraints)
    w_opt = result.x if result.success else w0
```

Appendix 5: EWMA Min Risk (No ESG) - Optimization and Backtest

```
# Rolling Backtest Loop
while i + TRAIN_DAYS + TEST_DAYS <= len(dates):
    train_idx = dates[i:i + TRAIN_DAYS]
    test_idx = dates[i + TRAIN_DAYS:i + TRAIN_DAYS + TEST_DAYS]

    train_returns = daily_returns.loc[train_idx]
    test_returns = daily_returns.loc[test_idx]
    tickers = train_returns.columns

    # Covariance Matrix
    shrink = LedoitWolf().fit(train_returns)
    cov = pd.DataFrame(shrink.covariance_, index=tickers, columns=tickers)

    # Expected Returns (EWMA + long-term mean)
    ewma_returns = train_returns.ewm(span=EWMA_SPAN, adjust=True).mean().iloc[-1] * 252
    long_term_mean = train_returns.mean() * 252
    mu_ewma = 0.7 * ewma_returns + 0.3 * long_term_mean

    # Optimization
    port = rp.Portfolio(returns=train_returns)
    port.mu = mu_ewma
    port.cov = cov
    weights = port.optimization(model="Classic", rm="MV", obj="MinRisk", hist=True)
```

Appendix 6: EWMA Max Sharpe (No ESG) - Optimization Setup and Backtesting

```
for start in range(0, len(daily_returns) - train_window - test_window + 1, test_window):
    # Train/test split and clipping
    train_returns = daily_returns.iloc[start : start + train_window].clip(lower=-0.1, upper=0.1)
    test_returns = daily_returns.iloc[start + train_window : start + train_window + test_window]

    # Use historical mean return, but EWMA for covariance
    mu = train_returns.mean()

    # EWMA covariance estimation
    cov_ewma = train_returns.ewm(span=(2/(1 - lambda_) - 1),
     adjust=False).cov(pairwise=True).iloc[-len(train_returns.columns):]
    tickers_used = mu.index.tolist()
    cov_matrix = cov_ewma.values.reshape(len(tickers_used), len(tickers_used))

    # Objective
    def neg_excess_sharpe(w):
        port_return = np.dot(w, mu)
        port_vol = np.sqrt(np.dot(w.T, np.dot(cov_matrix, w)))
        excess_return = port_return - (risk_free_rate / annualization_factor)
        return -excess_return / port_vol if port_vol > 0 else np.inf

    # Optimization
    n = len(mu)
    w0 = np.ones(n) / n
    bounds = [(0, 0.2) for _ in range(n)]
    constraints = [{'type': 'eq', 'fun': lambda w: np.sum(w) - 1}]

    result = minimize(neg_excess_sharpe, w0, bounds=bounds, constraints=constraints)

    if result.success:
        w_opt = result.x
    else:
        print(f"Optimization failed at step {start}, using equal weights.")
        w_opt = w0

    weights_list.append(pd.Series(w_opt, index=tickers_used))
    dates.append(daily_returns.index[start + train_window])
```

Appendix 7: CAPM Min Risk (No ESG) - Optimization Setup and Backtesting

```
def compute_capm_mu(asset_train_returns, market_train_returns, R_f=0.02):
    excess_market = market_train_returns.values.flatten() - (R_f / 252)
    mu = []
    for col in asset_train_returns.columns:
        excess_asset = asset_train_returns[col].values - (R_f / 252)
        beta = np.cov(excess_asset, excess_market)[0, 1] / np.var(excess_market)
        capm_return = R_f + beta * (np.mean(excess_market) * 252)
        mu.append(capm_return)
    return pd.Series(mu, index=asset_train_returns.columns)

# Portfolio optimization
def optimize_portfolio(mu, Sigma):
    mu_np = mu.values
    Sigma_np = Sigma.values
    n = len(mu_np)
    w = cp.Variable(n)
    risk = cp.quad_form(w, Sigma_np)
    constraints = [cp.sum(w) == 1, w >= 0]
    prob = cp.Problem(cp.Minimize(risk), constraints)
    prob.solve()
    return np.array(w.value).flatten() if prob.status == 'optimal' else np.ones(n) / n

# Backtest single quarter
def backtest_single_quarter_capm(asset_train_returns, asset_test_returns, market_train_returns,
     R_f=0.02):
    mu_capm = compute_capm_mu(asset_train_returns, market_train_returns, R_f)
    Sigma = asset_train_returns.cov() * 252
    weights = optimize_portfolio(mu_capm, Sigma)

    test_returns = asset_test_returns @ weights
    realized_return = np.mean(test_returns) * 252
    realized_std = np.std(test_returns) * np.sqrt(252)
    sharpe_ratio = (realized_return - R_f) / realized_std if realized_std != 0 else 0

    expected_return = mu_capm @ weights
    expected_volatility = np.sqrt(weights.T @ Sigma @ weights)

    return {
        'start': asset_test_returns.index[0],
```

```
        'end': asset_test_returns.index[-1],
        'expected_return': expected_return,
        'expected_volatility': expected_volatility,
        'return': realized_return,
        'risk': realized_std,
        'sharpe': sharpe_ratio,
        'weights': weights
    }

# Rolling backtest
def run_rolling_backtest_capm(asset_returns, market_returns, train_days=63, test_days=63):
    results = []
    i = 0
    while (i + train_days + test_days) <= len(asset_returns):
        asset_train = asset_returns.iloc[i:i+train_days]
        asset_test = asset_returns.iloc[i+train_days:i+train_days+test_days]
        market_train = market_returns.iloc[i:i+train_days]

        if asset_train.index[0] < pd.Timestamp("2020-01-01"):
            i += test_days
            continue
        if asset_test.index[-1] > pd.Timestamp("2024-12-31"):
            break

        result = backtest_single_quarter_capm(asset_train, asset_test, market_train)
        results.append(result)
        i += test_days

    return pd.DataFrame(results)
```

Appendix 8: CAPM Max Sharpe (No ESG) - Optimization and Backtesting

```
# Compute CAPM Expected Returns
def compute_capm_mu(asset_train_returns, market_train_returns, R_f=0.02, annualization=252):
    excess_market = market_train_returns.values.flatten() - (R_f / annualization)
    mu_capm = []
    for col in asset_train_returns.columns:
        excess_asset = asset_train_returns[col].values - (R_f / annualization)
        beta = np.cov(excess_asset, excess_market)[0, 1] / np.var(excess_market)
        capm_return = R_f + beta * (np.mean(excess_market) * annualization)
        mu_capm.append(capm_return)
    return pd.Series(mu_capm, index=asset_train_returns.columns)

# Max Sharpe Optimization
from scipy.optimize import minimize


def optimize_max_sharpe_capm(mu, Sigma, R_f=0.02, max_alloc=0.2):
    mu_excess = mu.values - R_f
    Sigma_np = Sigma.values
    n = len(mu)

    def neg_excess_sharpe(w):
        port_ret = np.dot(w, mu_excess)
        port_vol = np.sqrt(np.dot(w.T, np.dot(Sigma_np, w)))
        return -port_ret / port_vol if port_vol != 0 else np.inf

    w0 = np.ones(n) / n
    bounds = [(0, max_alloc) for _ in range(n)]
    constraints = [{'type': 'eq', 'fun': lambda w: np.sum(w) - 1}]

    result = minimize(neg_excess_sharpe, w0, bounds=bounds, constraints=constraints)
    return result.x if result.success else w0

# Parameters
train_window = 63
test_window = 63
annualization = 252
risk_free_rate = 0.02

# Storage
```

```
dates = []
weights_list = []
expected_return_list = []
expected_vol_list = []
expected_sharpe_list = []
realized_return_list = []
realized_vol_list = []
realized_sharpe_list = []

# Rolling Backtest
for start in range(0, len(asset_returns) - train_window - test_window + 1, test_window):
    train_assets = asset_returns.iloc[start : start + train_window]
    train_market = market_returns.iloc[start : start + train_window]
    test_assets = asset_returns.iloc[start + train_window : start + train_window + test_window]

    # Clip outliers
    train_assets = train_assets.clip(lower=-0.1, upper=0.1)

    # Estimate CAPM returns and covariance
    mu_capm = compute_capm_mu(train_assets, train_market, R_f=risk_free_rate,
     annualization=annualization)
    Sigma = train_assets.cov()

    # Optimize portfolio for max Sharpe
    weights = optimize_max_sharpe_capm(mu_capm, Sigma, R_f=risk_free_rate)
```