The Impact of ESG Integration Strategies on Portfolio Performance: Evidence from the European Equity Market

Giovanni Battista Del Basso

Dissertation written under the supervision of Christophe Thibierge

Dissertation submitted in partial fulfilment of requirements for the MSc in International Management, at Universidade Católica Portuguesa and for the MSc in Management at ESCP Business School, 15/05/2026.

Abstract

The growing interest in Environmental, Social and Governance (ESG) considerations has moved the debate from whether to integrate this information in portfolio construction to how it should be integrated.

Using data from the STOXX Europe 600 index, this thesis investigates how portfolio return and volatility vary when the ESG score is used as a primary screening criterion versus as a weighted factor within a multifactor framework. In addition, the study analyses how sensitive the performance outcomes are with different ESG weighting schemes.

The empirical results show that ESG integration does not have a uniform effect on performance, but rather depends on how it is implemented. The multifactor portfolio delivers the highest cumulative return in the in-sample period, with consistent risk-adjusted performance in the out-of-sample. By contrast, the static ESG screening strategy underperforms the benchmark. The third portfolio introduces the concept of ESG momentum and this strategy produces intermediate results, outperforming the benchmark but showing weaker statistical significance.

A key finding is that portfolio performance is mainly driven by risk metrics, while ESG contributes at the margin to increase the returns. The sensitivity analysis reveals that investors can adjust ESG exposure within a certain range without compromising performance, demonstrating the flexibility of ESG integration. Furthermore, the market beta is almost always negative and statistically significant, indicating the defensive characteristics of the strategies.

Overall, the study contributes to the literature by shifting the focus from whether ESG investing pays to how ESG should be integrated into portfolio construction.

Portoguese version

O crescente interesse por fatores Ambientais, Sociais e de Governança (ESG) deslocou o debate de se essa informação deve ser integrada na construção de portfólios para como deve ser incorporada.

Utilizando dados do índice STOXX Europe 600, esta tese investiga como o retorno e a volatilidade dos portfólios variam quando o score ESG é usado como critério principal de seleção, em comparação com seu uso como fator ponderado em um modelo multifatorial. Além disso, o estudo analisa a sensibilidade do desempenho a diferentes esquemas de ponderação ESG.

Os resultados empíricos mostram que a integração ESG não produz um efeito uniforme, dependendo da forma de implementação. O portfólio multifatorial apresenta o maior retorno acumulado no período in-sample, com desempenho ajustado ao risco consistente no período out-of-sample. Em contraste, a estratégia estática de triagem ESG apresenta desempenho inferior ao benchmark. Um terceiro portfólio introduz o conceito de momentum ESG, gerando resultados intermediários, superando o benchmark, mas com menor significância estatística. Um resultado central é que o desempenho dos portfólios é principalmente impulsionado por métricas de risco, enquanto o ESG contribui marginalmente para os retornos. A análise de sensibilidade mostra que é possível ajustar a exposição ao ESG dentro de certos limites sem comprometer o desempenho. Além disso, o beta de mercado é geralmente negativo e significativo, indicando características defensivas.

De forma geral, o estudo contribui ao focar não em se o investimento ESG compensa, mas em como deve ser integrado na construção de portfólios.

List of Keywords

ESG investing, Portfolio construction, ESG integration, Multifactor model, ESG screening, ESG momentum, Risk-adjusted return, Portfolio performance, Volatility, Maximum drawdown, Expected shortfall, Sharpe ratio, Sortino ratio, Sensitivity analysis, Alpha generation

Title of the Thesis

The Impact of ESG Integration Strategies on Portfolio Performance: Evidence from the European Equity Market

Author

Giovanni Battista Del Basso

Table of Contents

Introduction

Today, the question is no longer if investors should include or not ESG metrics in their investment decisions, the real question is how they should include these indicators. This introduction sets the stage for this Master's thesis, which examines the relationship between portfolio return and volatility under different integrations of ESG metrics.

Investors rely on ESG scores to assess non-financial risks and long-term value creation potential. The growing demand for ESG data is driven by stricter regulations and shifting investor preferences toward responsible investment strategies. As sustainability factors become more financially material, ESG performance is no longer viewed as purely ethical, but as an essential component of comprehensive risk assessment and strategic capital allocation.

Asset managers use ESG criteria to determine long term risk-adjusted returns and to identify growth opportunities such as the energy transition, the circular economy or the financial inclusion. The best example to show this trend is the 2020 letter to CEOs by Larry Fink (CEO of BlackRock), in which he emphasised that climate risk is investment risk and that sustainability should be at the centre of investment strategy, signalling a major reallocation of capital toward more sustainable assets (Fink, 2020).

In terms of figures, the Europe ESG investing market[1] generated revenue of USD 7,204.0 million in 2024, reflecting the strong commitment to sustainable finance and responsible investment practices. The market is projected to grow at a compound annual growth rate (CAGR) of 19.9% from 2025 to 2030. By segment, ESG integration emerged as the largest revenue-generating type in 2024, highlighting the widespread adoption of ESG (Horizon, 2024).

From a broader perspective, Europe accounted for 25.4% of the global ESG investing market revenue in 2024, demonstrating its significant role in shaping the global sustainable investment landscape. However, North America is projected to lead the regional market in terms of revenue by 2030. At the same time, Asia Pacific is expected to be the fastest-growing regional market, with projections indicating it will reach USD 19,528.2 million by 2030, reflecting the rapid expansion of ESG adoption across emerging and developed economies in the region (Horizon, 2024).

The information presented above shows the current situation of the ESG world and it forms the foundation of this thesis, which seeks to examine how ESG metrics can be integrated into portfolio construction to optimise risk-adjusted returns.

Three distinct investment strategies are developed, each applying different weighting schemes,

[1] The European ESG investing market refers to the ecosystem of investors, financial products, and regulatory frameworks within Europe that incorporate environmental, social, and governance (ESG) criteria into investment processes. The aim is to align capital allocation with sustainability objectives while managing financial risk and return.

in order to assess where and to what extent ESG integration contributes to performance enhancement.

The portfolio returns are then evaluated using the Capital Asset Pricing Model (CAPM) and the Fama–French five-factor model. This analysis allows for a decomposition of returns and a search for the presence of alpha.

Overall, the thesis is motivated by the evolving global investment landscape, in which ESG considerations are gaining increasing prominence, investor preferences are shifting, and sustainability is emerging as a central pillar of capital allocation decisions.

1 Research Question and Scope

As the first chapter of the thesis, this section introduces the research problem addressed in the project and it outlines its motivation and relevance. The chapter also defines the scope of the analysis to establish the boundaries of the study.

1.1 Research Question, Motivation and Relevance

The research question is driven by the considerations made in the introduction, and the precise research question is the following:

> *"In the EU equity market, how do portfolio return and volatility differ when ESG scores are used as a primary screening criterion versus as a weighted component within a multifactor framework, and how sensitive are these outcomes to alternative ESG weighting schemes?"*

The primary objective of this thesis is to assess whether investors benefit from incorporating ESG considerations into portfolio construction, or whether doing so entails a measurable sacrifice in terms of return or risk-adjusted performance. To address this question, three portfolios are constructed with varying degrees of ESG integration, allowing for a systematic comparison of performance outcomes. By analyzing differences in returns and volatility across these strategies, the study aims to quantify the potential trade-off or enhancement associated with ESG constraints or contributions in portfolio allocation.

This project is particularly relevant given the growing prominence of ESG integration in modern financial markets. While much of the existing debate focuses on corporate sustainability practices, this thesis adopts the perspective of the investor, examining how ESG metrics influence investment decisions and financial outcomes.

In order to contextualise this analysis, the first part of the study provides a broad overview of ESG principles and the evolving regulatory landscape, including recent European sustainability regulations. Understanding these regulatory developments is essential, as they shape corporate disclosure practices and increase the availability, comparability, and reliability of ESG information.

Building on this foundation, the focus shifts from corporate obligations to investor strategy, exploring the different approaches available for ESG-related investment. Given that ESG metrics represent a relatively recent addition to the set of tools available to investors, it is important to assess whether they should serve as a primary screening criterion or rather as a complementary factor within a broader portfolio framework.

1.2 Scope of the Analysis

The thesis focuses exclusively on European equities, specifically the companies included in the STOXX Europe 600 Index. The decision to limit the analysis to European stocks is motivated by considerations of data availability, consistency, and quality.

The European market provides a relatively mature and well-regulated ESG reporting environment, supported by standardised disclosure frameworks and stringent regulatory requirements. The data provider is the London Stock Exchange Group (LSEG Data & Analytics, formerly Refinitiv). The financial information is retrieved through an API call made with Python. For the ESG score, the data are obtained from Datastream and the score is used in its totality without using the individual pillar scores[1]. This automated data extraction process enhances transparency and ensures full reproducibility of the analysis.

The timespan is between January 2015 to March 2026, and the analysis covers three different data windows. The first three years are used as a warm-up period, during which the necessary statistics are analysed and financial ratios are constructed to establish the initial portfolios. The period between January 2018 and December 2024 is the main evaluation period over which the portfolios are compared and for one of the three portfolio, an additional out-of-sample period runs from January 2025 to March 2026.

The size of the sample is 600 stocks that are inside the index and cover the most important stocks in Europe.

As previously mentioned, the thesis analyses three distinct portfolio construction strategies. The first strategy adopts a multi-criteria approach, combining financial performance metrics (Sharpe ratio and Sortino ratio), risk measures (Expected Shortfall and Maximum Drawdown), and ESG score. The second strategy selects only firms that rank in the top decile (90th percentile and above) of ESG scores within each year. The third portfolio applies a momentum strategy based on the ESG score and it also selects only stocks that are in the top decile.

The first portfolio is rebalanced monthly, while the others two are rebalanced yearly and for each change in position a transaction fee is applied to show comparable net results with the benchmark. The fee applied is 5 basis points (bps) per trade and it reflects the normal fee for large-cap stocks.

[1]Refinitiv provides single score for the environmental, social and governance dimensions but in this thesis only the aggregation of the three score is used.

2 The ESG World

The chapter presents the definition of ESG and explores the development of the concept. Moreover, it outlines the regulatory framework in order to understand which types of information are available and the current stage of development of the sector. Finally, the chapter distinguishes between different types of investment strategies, providing all the possibilities for an investor.

2.1 Definition of ESG

Before having an extensive review of the regulations and the different types of investment, it is important to be aligned on the definition of ESG, so what is ESG? ESG stands for environmental, social, and governance. It refers to the way in which an organisation manages its impacts on, risks related to, and opportunities arising from, environmental, social and governance issues (Institute, 2021). Analyzing the three pillars:

- Environmental: covers both the organisation's impact on the natural environment and the environment's impact on the organisation, including climate change impacts, direct impacts (the organisation's own activities) and indirect impacts (throughout its value chain, including suppliers and customers or investments).
- Social: covers the organisation's interactions with, responsibilities towards, and impacts on a range of stakeholders, including society at large.
- Governance: is the process by which decisions are made and implemented in an organisation. It covers the compliant and ethical conduct of an organisation's activities and the controls in place for monitoring them.

2.2 Evolution of ESG in Financial Markets

The current situation of ESG in financial markets is not the result of a sudden shift, but the outcome of a long and progressive evolution in the relationship between finance, society and environment. What began as ethically motivated investment choices has evolved into a structured framework in which non-financial factors are considered as financial material. Understanding this historical trajectory is necessary to explain why ESG considerations are today inside the investment practices.

The modern concepts of Corporate Social Responsibility (CSR) and Socially Responsible Investing (SRI) have their roots in the religious movements of eighteenth and nineteenth centuries in both Europe and North America. In this period, Protestant communities pioneered the practice of excluding the "sin industries" from their investment portfolios, this meant refusing to hold investments in industries linked to the slave trade, liquor and war. This exclusionary logic set the groundwork for what would later be defined as negative screening in portfolio management.

During the 1960s and 1970s, the scope of responsible investing went beyond religious ethics. Public protests against the war (Vietnam war) prompted the creation of funds that explicitly excluded military and nuclear efforts. Moreover, anti-apartheid movement accelerated the trend to treat all employees equally, leading to selective divestments by institutional investors in companies that did not follow these principles. These episodes showed the fact that investors could use their capital to pressure corporations to change their behaviours.

The 1980s and 1990s marked a decisive turn toward environmental concerns. The early consensus on global warming and the depletion of the ozone layer brought the ecological risk in the political discourse. In 1987, the United Nations Brundtland Commission published its report "Our Common Future", which introduced the concept of sustainable development, defined as: "development that meets the needs of the present without compromising the ability of future generations to meet their own needs (Keeble, 1988)." This report provided the basis for the start of an international movement that includes also the Rio de Janeiro Earth Summit of 1992, which defined the principle of "do not significant harm" to the environment as a basic license to operate and put the management of negative externalities at the core of corporate responsibility.

The third pillar of ESG gained importance with a series of huge corporate scandals in the early 2000s. The collapse of Enron and other major companies exposed systematic failures in financial and audit integrity, prompting a change in the legislation. The new act, the Sarbanes-Oxley Act marked a turning point in corporate governance by transforming accountability from a normative expectation into a legally enforceable obligation.

The consolidation of environmental, social and governance considerations under the ESG umbrella is dated to 2004, when a United Nations report entitled "Who Cares Wins" first called for the integration of ESG issues into capital markets. This was followed in 2005 by the UNEP Finance Initiative's "Freshfields" report, which defined the legal framework for treating ESG considerations as part of the duty of institutional investors, the first time that ESG was viewed not as a charitable activity but as a component of prudent financial management.

In 2006, principles for responsible investing (PRI) were launched, committing signatories to incorporate ESG issues into investment analysis and decision-making process, to seek appropriate disclosure from investment companies and to report on their own activities and progress. By 2015, the sustainable development goals provided a new normative framework, extending ESG logic into a global agenda of seventeen interconnected objectives covering poverty, inequality, climate and institutional integrity.

Yet the path toward ESG integration has not been linear or without contestation and the most famous argument against ESG can be found in Milton Friedman's thought. For him, corporate executives who used shareholder resources toward social ends were effectively imposing a tax on investors, employees and customers. The challenge of Friedman's profit-maximisation thesis has not been fully resolved but it has evolved in how non-financial data should be incorporated into financial analysis.

Retracing this history is important to reveal that today's frameworks are not arbitrary constructions but are the cumulative product of a long journey. Each historical episode added a layer of complexity and urgency to the concept of responsible investment and it shaped its meaning.

2.3 ESG Regulatory Framework

The increasing integration of ESG considerations in financial markets has been accompanied by a growth in the regulatory framework governing the sustainability disclosures. This chapter demonstrates that ESG is no longer a voluntary choice but is increasingly becoming a regulatory requirement for companies. The regulations that are going to be covered are the EU Taxonomy, the CSRD, the SFDR (Articles 6, 7, 8, and 9) and the MiFID II. Before examining in details the regulations, the relationship between them is explained by the following graph:

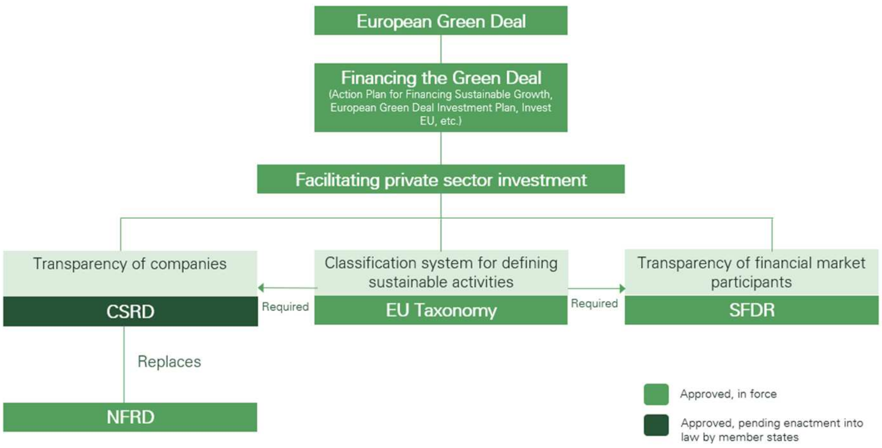

Figure 2.1: Relationship between regulations

2.3.1 EU Taxonomy

What makes possible to disclose ESG information is a common taxonomy, this regulation helps investors to understand the level of sustainability of an economic activity and it also supports the transition to a low-carbon economy by redirecting capital flows. The EU Taxonomy's objectives are linked to those of the EC Action Plan on Financing Sustainable Growth (Commission, 2018), published in March 2018, the three objectives are:

- Reorient capital flows towards sustainable investment, in order to achieve sustainable and inclusive growth.

- Manage financial risks stemming from climate change, environmental degradation and social issues.

- Foster transparency and long-termism in financial and economic activity.

The EU Taxonomy is built around four main principles:

- It follows the principle of technology neutrality, meaning it does not favour specific technologies but evaluates economic activities based on their environmental performance.

- It is designed as a dynamic and evolving tool, allowing updates over time as technologies advance and climate objectives develop.

- It aims to be easy to understand and use, providing a clear and practical classification system for sustainable activities.

- Instead of applying a strict binary distinction between "green" and "not green," it enables transition activities.

Under this approach, the taxonomy distinguishes three types of economic activities. Substantial contribution activities are those that significantly contribute to environmental objectives based on their own performance, such as renewable energy activities that have a low environmental impact and can replace more polluting alternatives. Transitional activities refer to sectors where low-carbon alternatives are not yet fully available but where certain activities perform better than the sector average, for example best-in-class cement production with lower greenhouse gas emissions. Enabling activities are those that support other sectors in making a substantial environmental contribution through their products or services, such as manufacturing renewable energy technologies or installing energy-efficient features in buildings.

Overall, the EU Taxonomy provides a structured, performance-based framework that supports the transition to a more sustainable and low-carbon economy while remaining adaptable over time.

In practice, the first step is determining whether an activity is eligible under the taxonomy. To be officially recognised as "taxonomy-aligned," an activity must satisfy three key conditions.

First, it must make a substantial contribution to at least one of the six environmental objectives defined in the regulation. These objectives include climate change mitigation, climate change adaptation, sustainable use and protection of water and marine resources, transition to a circular economy, pollution prevention and control, and the protection and restoration of biodiversity and ecosystems. In doing so, the activity must comply with the technical screening criteria established for each objective.

Second, the activity must do no significant harm to any of the other environmental objectives.

Third, it must comply with minimum social safeguards, meaning it aligns with international standards such as the OECD Guidelines for Multinational Enterprises, the UN Guiding Principles on Business and Human Rights, and International Labour Organisation conventions.

Activities that do not meet these requirements are considered non-aligned and are therefore not recognised as sustainable under the EU Taxonomy. In the Appendix, there are the sectors covered by the taxonomy and how the taxonomy is shown by the companies.

2.3.2 CSRD

In the European Union, the regulation of the ESG disclosure is the Corporate Sustainability Reporting Directive (CSRD), which represents the most comprehensive corporate disclosure obligations until now. The rule entered into force in 2023 and it replaced the Non-Financial Reporting Directive (NFRD), expanded the number of companies subject to reporting obligations covering nearly 50,000 companies. The CSRD includes corporations with more than 1,000 employees and €50 million or more in revenues; from 2026, the obligation of disclosure extends to large corporations not previously covered by the NFRD. Regarding non-EU companies, they are targeted if they generate more than €150 million in turnover within the EU starting from 2028. Small companies remain out of scope, but they can do a voluntary disclosure to answer to the pressure from the stakeholders.

The CSRD introduced detailed and standardised reporting requirements through the European Sustainability Reporting Standards (ESRS), which cover environmental (ESRS E), social (ESRS S) and governance (ESRS G) topics. These standards generate 1,200 observations that are structured in four main disclosure areas:

- Strategy to align with the 1.5°C Paris Agreement.
- Management of impacts, risks and opportunities.
- Governance structure and compensation incentives linked to ESG goals.
- Metrics and targets to achieve.

A major innovation of the CSRD is the double materiality, in the sense that companies must report both on how sustainability issues affect their financial performance and on how their activities impact society and environment. Before, the companies were not required to disclose this double materiality but they had to disclose only how they were affected financially.

Another important change is the requirement of external auditors, treating sustainability information at the same level of financial data.

In essence, the CSRD improves quality, consistency and availability of sustainability data from companies, providing the foundation for investors to include these dimensions into investment decisions.

2.3.3 SFDR and MiFID II

The CSRD covers the companies perspective, the regulations of this section change the perspective and cover specifically the financial sector. The Sustainable Finance Disclosure Regulation

(SFDR) and the Markets in Financial Instruments Directive II (MiFID II) form a regulatory layer that includes sustainability directly into investment processes, portfolio creation and client advisory.

The SFDR imposes transparency obligations for the participants of financial market (asset managers, pension funds, financial advisers...) requiring them to disclose how they integrate sustainability risks into their investment decisions and, where it is possible, to report the principal adverse impacts (PAIs) on environmental and social factors reporting indicators as greenhouse gas emissions intensity, biodiversity impact etc.

At the product level, the SFDR introduces a classification framework that categorise products according to sustainability ambition[1]. Each category has its own rules of disclosures, and a central rule for the last 2 categories is the "Do No Significant Harm" principle, which stipulates that an investment contributing to one environmental or social objective must not materially harm any other objective.

The MiFID II sustainability amendments extend the SFDR into the retail and institutional advisory relationship, ensuring that the sustainability characteristics of financial products are not only disclosed but actively matched with the preferences of end investors. Under this framework, investment firms are required to integrate a new dimension in the sustainability assessment process: the sustainability preferences. These preferences cover three dimensions:

- The desired minimum portion of EU Taxonomy aligned investments
- The minimum proportion of SFDR sustainable investments
- The desired consideration of PAIs

The effective functioning of both the SFDR and the MiFID II is dependent on the availability of reliable, standardised and comparable data. The progressive implementation of CSRD is expected to be the start of the transformation of sustainable finance.

2.4 ESG Investment Strategies

The last part of this chapter is dedicated to the investor point of view with the scope to analyse all the possible alternatives that an investor has. The recognition of climate risk as a financial factor has changed the investment practice, and investors have progressively integrated ESG considerations into portfolio construction. These new strategies created by the integration of ESG differ in their objectives, methodology and degree of intentionality.

The principal approaches are the following:

- ESG integration

[1] The first type are products that integrate sustainability risk but make no explicit sustainability claims; the second type are products that promote Environmental or social characteristics as part of their investment strategy; Finally, there are products that pursue sustainable investment as their core objective

- Screening investing
- thematic investing
- Impact investing
- Philanthropy

The different strategies are shown from the least to the most ESG-engaged, the full picture is the following:

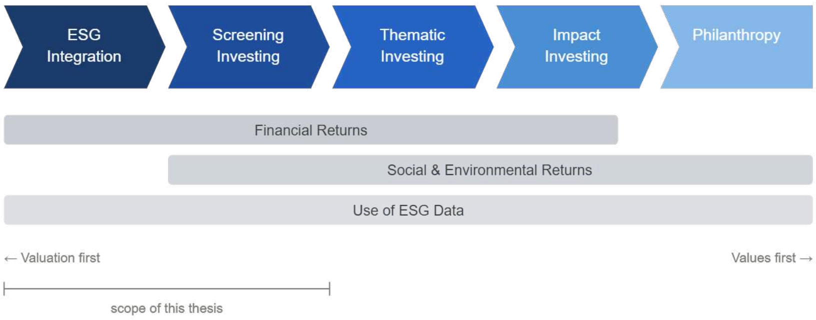

Figure 2.2: Responsible investing spectrum

2.4.1 ESG Integration

ESG integration represents the incorporation of ESG data into investment analysis and portfolio construction. Rather than a standalone screening, it is embedded in the investment life-cycle. Climate related factors are particularly important in this framework: transition and physical risk metrics such as carbon intensity, stranded asset exposure, disrupt operations, exposure to extreme events are increasingly treated as financially material.

The idea behind ESG integration is the hypothesis of incomplete market pricing. If the markets do not reflect the financial consequences of physical and transition risks, then investors with superior ESG capabilities may identify mispriced securities. This framing positions ESG integration as consistent with standard investment strategy rather than a different approach.

2.4.2 Screening Strategy

Screening is one of the best-known tools of responsible investing; it operates by applying predefined criteria to delimit the eligible investment universe.

Exclusionary screening removes exposure to sectors or activities incompatible with climate objectives (coal or other high-carbon energy sources for example).

Positive or best in class screening, by contrast, selects only companies demonstrating superior climate performance relative to sector competitors, rewarding the ones that are managing transition risk in the best way.

Norms based screening applies international standards such as the one from the Paris Agreement as a baseline filter.

From an ESG perspective, screening strategies reduce portfolio exposure to transition and physical risk, and they send a capital allocation signal to firms regarding the financial cost of inadequate climate management.

2.4.3 Thematic Investing

Thematic investing seeks to gain exposure to structural trends that are expected to reshape risk and return. The most used themes are category as climate change, energy transition and circular economy models. Unlike ESG integration, which is applied in a diversified portfolio, thematic strategies concentrate the capital in companies positioned to benefit from the low carbon transition.

This strategy is relevant to capture opportunities from the climate change crisis, with companies developing climate adaptation infrastructure, energy efficient technologies or low carbon goods representing the natural targets for thematic investment approaches.

2.4.4 Impact Investing

Almost at the end of the responsible investing framework, impact investing seeks to generate measurable positive social and environmental outcomes in addition to financial returns.

This strategy is different from the others because it has a clear intention to create a change, and this change has to be clearly linked with the investment strategy and it has to be reported with standardised metrics to track results.

The strategy may target assets such as renewable energy infrastructure, green bond issuance or instruments that support the climate change adaptation. The expected financial returns are variable, from below market to competitive, but what is common is the requirement for additionality: the investment must contribute to positive outcome that would not have occurred in its absence.

2.4.5 Philanthropy

Philanthropy represents the complete subordination of financial return to social and environmental goals. Unlike impact investing, which still requires a financial return, this strategy uses capital with no expectation of monetary repayment. While it falls outside the idea of portfolio management, it plays an important role in the sustainable finance world by funding early stage initiatives that later attract impact or thematic capital once the commercial viability is created.

From the perspective of Behavioural portfolio theory, that is going to be shown in the literature review, philanthropy is the top of the aspirational layer, where the utility is derived by non financial returns.

3 Literature Review

In this section, five papers are examined to present previous findings in the field of ESG integration and portfolio creation. Together, these studies provide a clearer understanding of current market conditions, highlight different perspectives within the literature, and identify key limitations in existing research.

3.1 ESG and Financial Performance

Friede et al. (2015) examine the relationship between ESG criteria and corporate financial performance (CFP) by synthesizing more than four decades of empirical research. The objective of this study is to address the claim that the ESG-CFP relationship is mixed or inconclusive. Rather than running another single test, the authors consolidate the accumulated academic evidence and determine whether a clear overall pattern emerges.

The data collection process is based on a search for studies with a quantitative summary of ESG-CFP findings, while narrative reviews are excluded because they do not permit structured aggregation. By focusing on review papers rather than individual articles, the paper maximises the coverage of the literature, and this is the key innovation of Friede et al. (2015), the implementation of a second-level review. Instead of directly aggregating individual primary studies, they collect and synthesize 60 prior review papers, including both vote-count reviews and first-order meta-analyses.

Methodologically, the vote-count analysis categorises findings as positive, negative, neutral, or mixed, and then aggregates the frequency of each category. This method provides a descriptive overview of the ESG-CFP relationship but it does not incorporate information about effect sizes, sample sizes, or statistical precision.

To address these limitations, the authors conduct a second-order meta-analysis. This approach combines the effect sizes of individual studies to estimate the overall strength and direction of a relationship. By using both vote-count and second-order meta-analytic methods, the authors compare whether the general directional patterns observed descriptively are confirmed when effect sizes are formally aggregated.

The main findings strongly support the business case for ESG investing. Approximately 90 percent of the reviewed studies report a nonnegative relationship between ESG and corporate financial performance, and the majority of those findings are positive. This suggests that ESG engagement is not systematically associated with financial underperformance and, in most cases, correlates positively with financial outcomes.

One of the most important distinctions the paper makes, particularly relevant for this project, is between firm-level and portfolio-level studies. The authors find that portfolio-based studies, such as those analyzing ESG mutual funds, ESG indices, or long–short portfolio strategies, tend to report more neutral or mixed results. They argue that portfolio returns may "wash out" firm-level ESG advantages due to noise, overlapping exposure to conventional risk factors, mixed screening approaches combining positive and negative criteria, and implementation frictions such as management fees and transaction costs. This distinction is crucial because it suggests that the strength of the ESG–performance relationship depends partly on how ESG is implemented in practice.

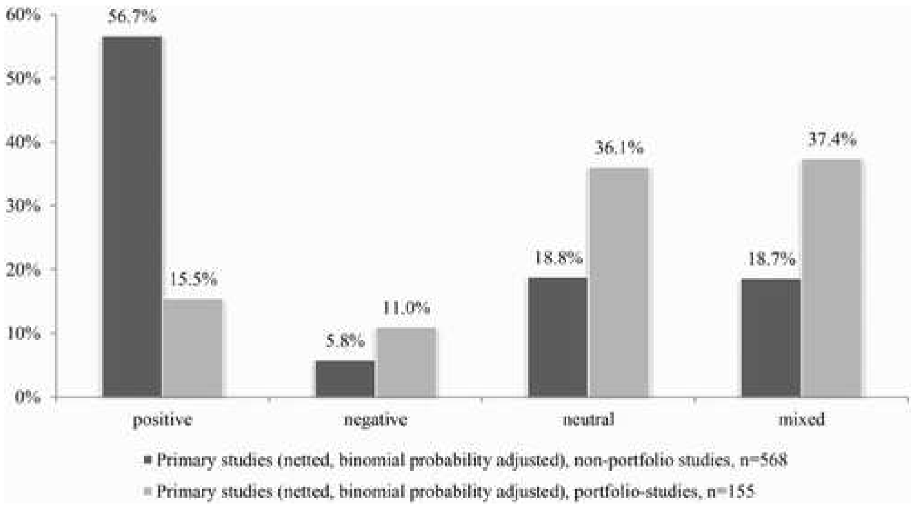

Figure 3.1: ESG–CFP relation portfolio and non-portfolio sample Friede et al. (2015)

The paper further explores whether the ESG–CFP relationship changes over time, particulary, they investigate whether ESG-related returns diminish as markets become more efficient at pricing ESG information. The authors find that the positive relationship appears relatively stable over time, with no clear evidence of systematic deterioration.

In interpreting these findings, they frame ESG as potentially relevant information that can influence financial performance through multiple channels, including improved risk management, stronger stakeholder relationships, operational efficiencies, and better governance practices. Importantly, they argue that the perception that "ESG does not pay" is often driven by evidence from portfolio-level studies, whereas firm-level evidence is more consistently positive.

The paper provides two important insights for this project. First, it establishes that the overall ESG–CFP relationship is generally nonnegative and often positive, implying that incorporating ESG is unlikely to impose a systematic return penalty. Second, it cautions that portfolio-based ESG strategies often yield weaker or more mixed outcomes than firm-level analyses.

This distinction directly motivates further investigation into how ESG is integrated into portfolios, including whether an ESG-only strategy differs materially from a multifactor strategy.

3.2 ESG Integration and the Investment Management Process

The paper of Van Duuren et al. (2016) analyses the social responsible investing (SRI), but it does in a different way with respect to the earlier empirical studies that examined whether SRI strategies produce superior or inferior risk-adjusted returns. The authors note that this performance-centred approach dominates the field and argue that it leaves an important gap: there is limited understanding of how professional asset managers actually interpret and implement ESG criteria within their investment frameworks.

By focusing on the investment process rather than only on performance indicators, the study aims to increase the understanding of ESG integration. Conceptually, the paper identifies ESG integration as the explicit inclusion of ESG risks and opportunities in financial analysis and investment decisions through a systematic process.

The authors argue that ESG investing is often presented as a distinct strategy, and they show that fundamental investing emphasises only company-specific information, long-term performance, management quality, and strategic positioning. The authors suggest that governance quality, environmental strategy, and stakeholder management can be interpreted as extensions of these traditional analytical dimensions. This frame is central for the paper's contribution, as it proposes that ESG investing is a refinement or a broadening of fundamental investment. The idea of the paper is used in this project to build the first portfolio, where the ESG score is used as an additional indicator and not as the only metric.

Methodologically, the study relies on an international survey of portfolio managers. The questionnaire was distributed in 2013 to 251 fund managers and after the follow-ups, 126 responses were collected, roughly 50% of response rate. One important note, the survey targeted conventional managers and not ESG-labelled or ethical funds. This choice allows the authors to assess whether ESG practices were diffused into asset management or if they still were a niche strategies.

The survey measures the degree of ESG integration using a self-reported ESG score ranging from no integration to full integration with explicit quantitative and qualitative ESG information. Additional questions examine how ESG information is used in portfolio construction, and whether managers perceive ESG integration as enhancing performance or reducing risk. The last distinction the questionnaire does is about the types of ESG information used, distinguishing between raw data, ratings, company-level versus sector-level analysis.

The findings reveal that many managers have incorporated elements of ESG into their investment processes, a large majority suggest that ESG considerations are no longer marginal in finance. ESG information is often used for risk management and red flagging of controversial firms, while exclusionary screening plays a minor role. Another important result is that managers prefer processed ESG ratings and company level analysis rather than raw ESG data. Moreover, governance factors receive more attention than environmental and social aspects, highlighting that corporate governance is central for investment analysis.

An additional contribution of the study lies in its national comparison. The authors find huge differences between U.S. and European managers. U.S. managers appear more sceptical about ESG's financial benefits, whereas European managers are more optimistic and likely to integrate ESG.

The paper concludes that ESG integration is less about activism and more about expanding traditional investment analysis. Rather than an ethical movement, ESG is represented as a pragmatic extension of fundamental investing that incorporates new risk and value information. The authors argue that ESG investors perceive their approach as consistent with business practice rather than as a trade-off between financial returns and social impact.

In terms of limitations, the study relies on self-reported survey data, which can be subject to social desirability bias. The sample is international but it is linked to a specific institutional network which may limit generalisability. Nevertheless, the paper makes a significant contribution by shifting the focus from whether ESG pays to how ESG is used in investment management.

3.3 Aggregate Confusion: The Divergence of ESG Ratings

Berg et al. (2022) investigate the disagreement across different ESG rating agencies, in fact, despite the growing importance, ratings from different providers exhibit only moderate correlation, raising questions about what ESG ratings actually measure and how reliable they are. It is important to mention this paper in the literature review because it is related to one of the main limitations of this study, which is the use of a single ESG score.

The paper is not the first one in this topic, prior research had shown that ESG ratings diverge and suggested that differences emerge in what agencies choose to measure and how they measure it. Berg et al. (2022) advance the literature by analyzing six ESG rating providers (KLD, Sustainalytics, Moody's ESG, S&P Global, Refinitiv, and MSCI) and all 709 underlying indicators used to build the ratings. This data analysis allows them to move beyond documenting disagreement to decomposing its causes.

The starting point of the analysis is the empirical confirmation of this divergence and the authors demonstrate that firms rated highly by one agency may receive average or even below the average scores by other agencies.

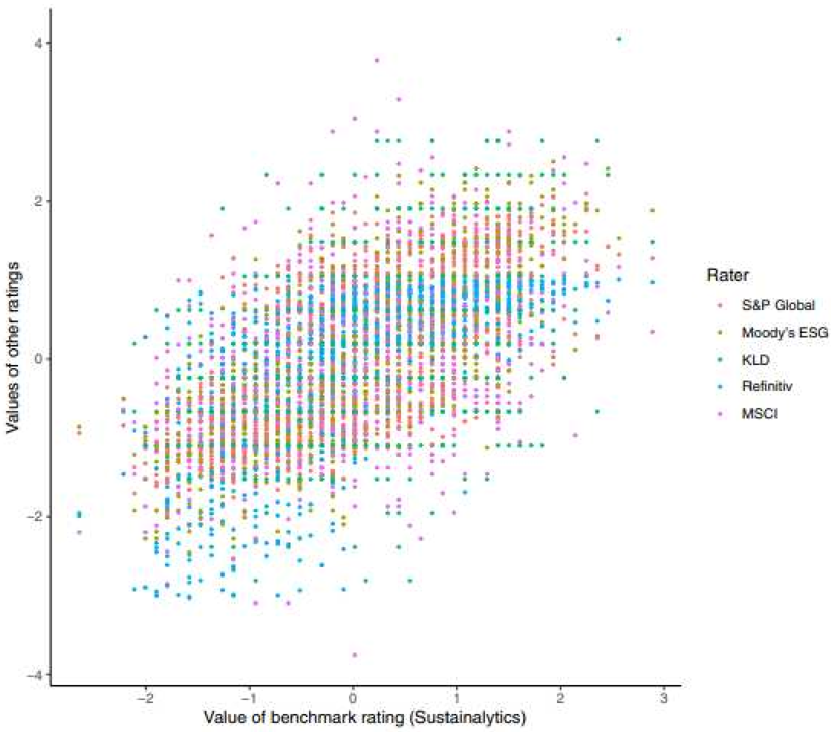

Figure 3.2: ESG rating disagreement Berg et al. (2022)

To explain this divergence, the authors decompose ESG ratings into three components: scope, measurement and weight. Scope refers to which ESG attributes are included in the rating. Measurement refers to how an attribute is quantified. Weight refers to how attributes are aggregated into one final score. Moreover, the authors map all the 709 indicators into 64 comparable categories, such as water, supply chain, biodiversity, etc. Using this structure, they calculate the scores and they demonstrate through regressions that they can replicate agencies' aggregation rules (The R^2 is often above 0.9).[1]

The most important contribution of the paper is the arithmetic decomposition of the rating divergence. On average, 56% is made by measurement differences, 38% to scope differences, and only 6% to weight differences. In other words, the disagreement is not that agencies prioritise different issues or aggregate them differently, but that they assign different values to the same ESG category, they disagree in how they score firms. This result has an important implication, it shows that if the goal is the harmonization, this process would require a standardisation in the underlying data collection and evaluation practices. The problem is not a philosophical disagreement about ESG, it is an empirical disagreement about firm performance.

The authors further investigate the nature of measurement divergence and they identify a "rater effect". They discovered a pattern, when a firm is viewed favourably by one rater in one dimension is more likely to be rated favourably in other dimensions by the same rater. This pattern indicates that disagreement is not a random noise but it follows systematic patterns linked

[1] The paper also replicates the analysis using an alternative taxonomy and non-linear aggregation methods such as neural networks and random forest. Across these variations, the results remain consistent.

to agency methodologies.

To conclude, the first contribution of the paper is that asset pricing tests of ESG risk premia or analyses of ESG-related alpha and studies of ESG-performance relationships may all depend on the chosen ESG rating.

Secondly, from a policy perspective, the paper suggests that harmonising ESG disclosure standards could reduce divergence and improve the comparability. It also highlights the potential value of a standardised taxonomy to make methodological differences easier to understand and more transparent. The importance of this paper is driven by the fact that they shifted the debate from whether ESG ratings disagree to why they disagree.

3.4 The ESG Efficient Frontier

Pedersen et al. (2021) develop a theoretical and empirical framework to analyse how ESG considerations affect portfolio choice, asset pricing and equilibrium returns. The paper aims to address a central issue in the ESG world: whether incorporating ESG into portfolio construction reduces expected returns or improves performance. The core of the paper is to extend the Markowitz framework[2] by introducing ESG scores as an additional portfolio characteristic. ESG plays two distinct roles: first, the ESG may contain information about firm fundamentals, affecting the expected cash flows and the risk; second, investors can derive direct utility from holding an ESG portfolio, reflecting non monetary preferences. This dual role allows the model to have information about the returns and about the preferences of investors.

Before deepening the analysis of the model, it is important to explain the three types of investors that exist in the paper:

- ESG-unaware investors ignore ESG information.
- ESG-aware investors incorporate ESG information in the construction of the portfolio but they do not have ESG preferences.
- ESG-motivated investors both use ESG information and derive utility from holding an ESG portfolio.

The main innovation of the paper is the concept of ESG-efficient frontier. Instead of the usual mean-variance frontier, the authors derive a frontier that characterises the highest obtainable Sharpe ratio for different level of ESG score.

[2]The Markowitz framework is a mean–variance approach to portfolio construction, it shows how investors can construct efficient portfolios that maximise expected return for a given level of risk (variance), or equivalently minimise risk for a given expected return.

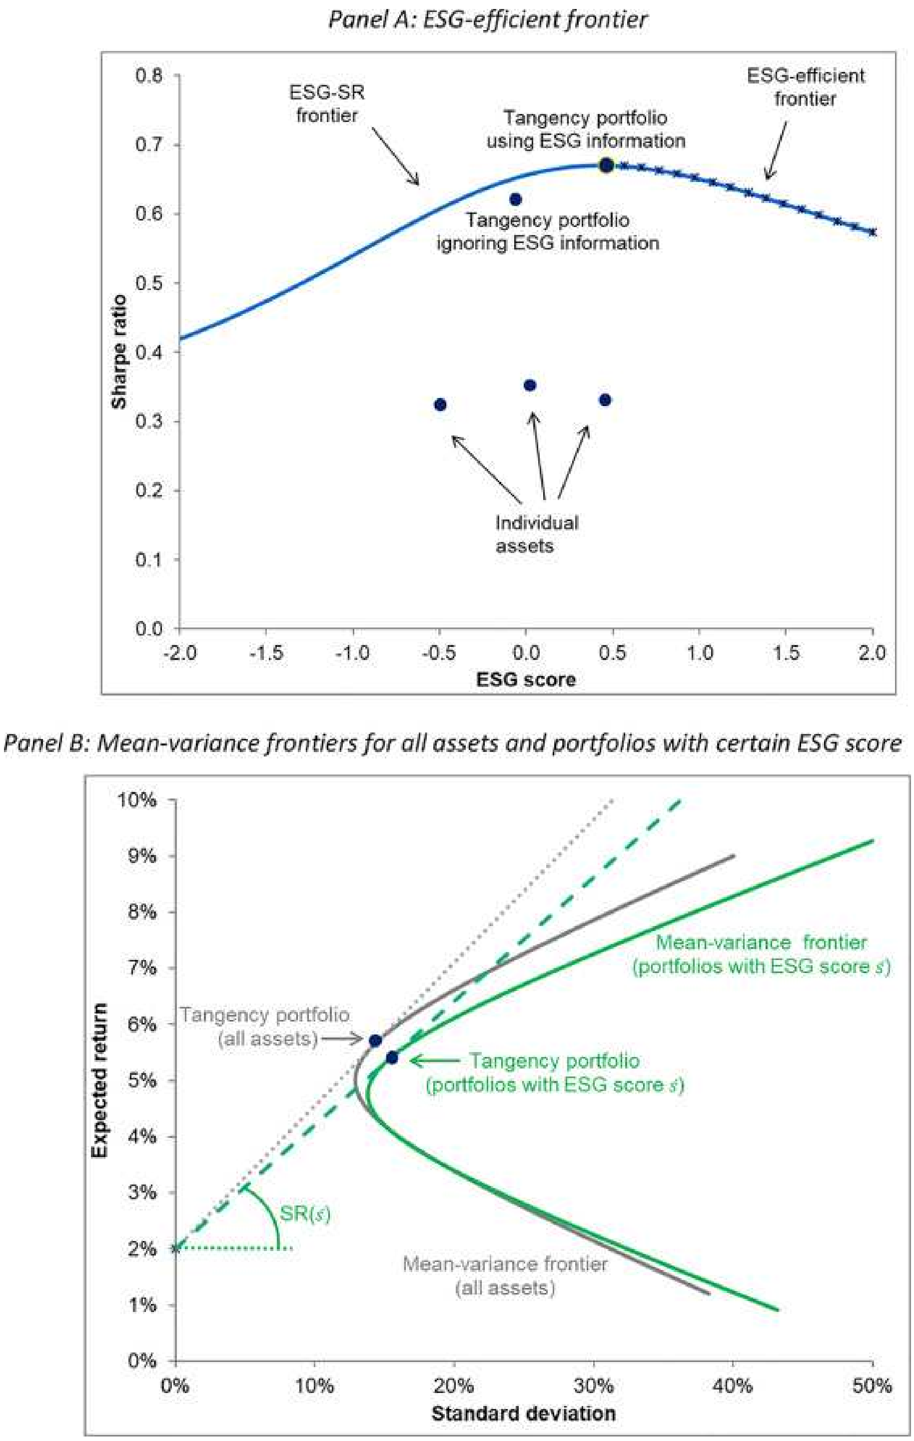

Figure 3.3: Efficient frontier Pedersen et al. (2021)

In the first panel is possible to see the ESG-efficient frontier, which shows the highest level of Sharpe ratio obtainable for each level of ESG exposure. The second panel shows the differences between the standard efficient frontier using all assets and the efficient frontier when portfolios are restricted to have a given ESG score. One of the most important results from the ESG-efficient frontier is that the choice of the portfolio can be reduced to a trade-off between Sharpe ratio and ESG preference.

The optimal ESG level maximises a function that combines the square of the Sharpe ratio and the investor's ESG preference function. This separation allows to first construct the ESG-efficient frontier using asset characteristics, then select the preferred portfolio based on investors preferences.

The other main contribution of the paper tests if expected returns depend not only on market beta but also on ESG exposure (ESG-adjusted CAPM). Empirically, for governance, incorporating the pillar improves the maximum Sharpe ratio by 12% relative to ignoring it. In contrast, environmental shows little predictive power for returns. In addition, the paper also evaluates screening strategies. The paper demonstrates that screening can reduce a portfolio's overall ESG score relative to an unconstrained strategy, since in the second one the investors can short low-ESG stocks to finance long positions.

To conclude, the paper makes three major contributions. It generalises portfolio theory to include ESG preferences, it quantifies the costs and benefits through the ESG-efficient frontiers and it estimates the ESG premia.

3.5 Behavioural Portfolio Theory

The last paper of this literature review shows a different aspect to take into consideration, that investors do not select portfolios based only on rational behaviours. Shefrin and Statman (2000) develop a Behavioural Portfolio Theory (BPT) which proposes an analysis of how investors actually build portfolios, they support the idea that investors have multiple goals (avoid poverty, have a possibility to become rich, non financial goals and lifecycle goals) and they treat these goals separately, which leads to portfolios that can look inefficient from Markowitz perspective.

The paper is based on two behavioural decision theories, Lopes and Oden (1999) and Levy (1992), and with these two models the portfolios are then constructed. The first theory explains that risky choices are driven by a dual process in which individuals balance security (avoiding losses) and aspiration levels (achieving desired outcomes), rather than simply maximising expected utility. The second paper develops the Prospect theory, a theory that explains why individuals often make choices that deviate from traditional rational models such as expected utility theory.

Shefrin and Statman (2000) represent the investor's portfolio in layers and each layer is linked to a distinct goal. The lower layer is designed for capital preservation, while the upper layer is designed for aspirations and high-payoff opportunities.

The results show that the BPT-construction leads to inefficient portfolios in comparison to the standard mean-variance perspective, in fact they look like a combination of safe assets and very risky one ("lottery tickets" to look for high upside).

To conclude, BPT implies that investors consider other aspects rather than only variance and expected returns, they also consider psychological goals such as security, aspiration and status. As a result, portfolios may deviate from mean-variance efficiency without a rational explanation and this conclusion can be directly linked to ESG. Under BPT, investors can accept lower Sharpe ratios if ESG satisfies a different goal (aspirational layer) and this idea can answer and give another support to the ESG integration in portfolio construction even if they deliver lower returns.

4 Data and Methodology

The chapter describes the methodology adopted in this thesis. It begins by explaining the data sources, with a focus on the selection criteria and the construction of the ESG score by Refinitiv. It then analyses the processing steps like data cleaning, data transformation and the procedures to avoid the look-ahead bias. Finally, the chapter explains the portfolio construction strategies and how the returns are computed.

4.1 Data

This thesis analyses European stocks, with data sourced from Refinitiv Workspace. To ensure reproducibility of the analysis, it is important to clearly define how the data are collected and processed. All data are retrieved either through direct API calls or via Datastream functions; the subsequent data manipulation is performed in Python.

ESG scores are also obtained from Refinitiv Workspace, covering the period from 2015 to 2026. This score ranges from 0 to 100. In qualitative terms, scores above 75 are classified in the highest category (A rating). A more detailed explanation of the score is given in the next section.

The risk-free rate used in the computation of the Sharpe ratio is sourced from Fama-French database. The same source is used also to obtain the Fama-French five factors.

A key challenge in the data preparation is the integration of datasets with different identifiers. In particular, Datastream identifies companies by name, while the API relies on Refinitiv Instrument Codes (RICs). To address this issue, a mapping dictionary is constructed in Python to link company names and their RICs.

4.1.1 ESG Score Construction

Before moving on, it is fundamental to define the characteristics of the ESG score that was used. The ESG score provided by LSEG is a relative, data-driven measure designed to assess a company's environmental, social, and governance performance in comparison with its peers. The score is constructed using publicly available, company-reported information and follows a transparent, multi-step methodology that incorporates industry materiality, disclosure practices, and ESG-related controversies.

The process begins with the collection of ESG data. LSEG gathers more than 870 ESG indicators per company from sources such as annual reports, sustainability reports, regulatory filings, company websites, NGO publications, and media coverage. From this large dataset, a subset of approximately 186 indicators that are considered as the most material and comparable across firms is selected for scoring purposes. These indicators include both boolean variables, such as the existence of specific ESG policies, and numeric variables, such as emissions levels or training hours.

In the next step, these indicators are aggregated into ten ESG categories that together form the three main pillars. For each indicator, companies are evaluated using a percentile ranking approach. Environmental and social indicators are benchmarked against firms within the same industry group, as defined by the LSEG Business Classification (TRBC), while governance indicators are benchmarked against firms within the same country of incorporation. This relative ranking approach ensures that scores reflect peer performance rather than absolute thresholds. The percentile scores of individual indicators are then combined to obtain a score for each ESG category.

Because ESG issues do not have the same importance across all industries, LSEG applies an industry-specific materiality adjustment using what is known as the ESG magnitude, or materiality matrix. This matrix assigns weights to each ESG category on a scale from one to ten, reflecting the relative importance of that category for a given industry. For environmental and social categories, these weights are derived from industry medians of relevant indicators and from the degree of disclosure across firms, while governance category weights are kept consistent across industries. The category weights are then normalised so that they sum to one within each industry, ensuring comparability across firms.

The category scores are subsequently aggregated, using these materiality-based weights, to compute the three pillar scores. These pillar scores are then combined into a single overall ESG score, which is normalised to range between 0 and 100. For ease of interpretation, this numerical score can also be expressed as a letter grade, ranging from D– for the weakest relative performance to A+ for the strongest relative performance.

In addition to the ESG score, LSEG calculates an ESG Combined score, known as the ESGC score, which adjusts the ESG score to account for significant ESG-related controversies. This adjustment is based on the occurrence of negative events reported in the media across 23 predefined controversy topics, such as environmental damage, corruption, labour issues, or product safety incidents. Controversies are weighted by severity and by company size to mitigate the bias that larger firms may face due to greater media exposure. When a company is involved in ESG controversies, the ESGC score is calculated as the average of the ESG score and the controversy score; when no such controversies are present, the ESGC score is equal to the ESG score.

Overall, the LSEG ESG scoring framework is explicitly relative in nature. It does not aim to define an absolute standard of what constitutes "good" ESG performance, but instead evaluates how well a company performs compared to its peers, taking into account industry-specific ESG risks, materiality, and transparency. This relative and industry-adjusted approach makes the LSEG ESG score particularly suitable for empirical research, portfolio construction, and comparative analysis in quantitative investment strategies.

4.2 Data Processing

In this part the technical processes are explained, with an analysis on how the data were prepared and how the portfolios were actually built.

4.2.1 Data Preparation

The raw datasets arrived in different formats and frequencies, requiring a series of transformations before they could be combined into a single panel. The stock price data is stored in wide format, with each column representing a different ticker and each row a monthly date. To compute per-stock statistics and to merge with ESG data, the dataset is reshaped into long format, so each row corresponds to a unique month observation with columns for price and returns. The two long-format datasets are then merged on the combination of ticker and month. The Fama-French risk-free rate is also appended through a date index join.

To avoid look-ahead bias, the ESG score is shifted forward by twelve months, that because Refinitiv publishes ESG score only after the year has ended, and so the score for 2018 is available only in 2019. By applying a twelve month shift, the portfolio formed in January 2020 relies on the 2018 ESG score, the most recent one that an investor could have.

The second adjustment regards all the risk and return metrics used for portfolio ranking (Sharpe ratio, Sortino ratio, Maximum Drawdown and expected shortfall) that are lagged by one month. This means the portfolio formed at the beginning of month t is based on information available through the end of month t-1, preventing any contemporaneous data leakage.

Third, the first three years of data are used as a warm up period, ensuring that the expanding window [1] has a meaningful baseline by the time the first portfolio is created in January 2018. The minimum number of observations required to produce a valid metric is set equal to the pre-2018 history of the stock or twelve months for stocks entering the index after 2015.

As previously mentioned, a set of risk and return metrics is computed for each month and all metrics rely on an expanding window. The scope of the metrics is the following:

- The Sharpe ratio captures the risk-adjusted return per unit of total volatility.
- The Sortino ratio replaces the total volatility of the Sharpe ratio with the downside deviation, which considers only negative excess returns.
- Maximum Drawdown quantifies the largest cumulative loss from a peak to a subsequent trough. This metric informs an investor about the worst-case scenario the investor could

[1] A monthly expanding window is a rolling approach in which the sample used for calculation includes all past observations available up to that month. As time progresses, the window expands and the sample incorporates the new data without discarding earlier observation.

have experienced.

- Expected Shortfall or Conditional value at risk (CVaR) at 95% confidence level measures the expected loss in the worst 5% of monthly return outcomes. This metric is computed as the mean of all returns falling at or below that threshold.

4.2.2 Portfolio Construction

Three portfolio strategies are constructed, each representing a different way of integrating ESG into the investment process. All three use equal weighting across their selected stocks.

The first portfolio combines the financial and risk metrics with ESG scores through a composite ranking. Each month, every stock is assigned to a percentile rank (between 0 and 1) on five dimensions: the annualized Sharpe ratio, the annualized Sortino ratio, the Maximum Drawdown, the CVaR, and the shifted ESG score. The ranking is designed so that higher percentile values reflect stronger performance. To ensure this, the two risk measures are ranked in descending order, as less negative values indicate better results.

These five percentile ranks are then aggregated into a composite score. The Sharpe and Sortino ranks are averaged to form a return score, while the Maximum Drawdown and CVaR ranks are averaged to form a risk score. The final score is a weighted combination and the stocks that were selected are the one that falls in the top 5%. The weights of the composite score are decided with a sensitivity analysis explained in the next chapter. Because portfolio 1 is the only strategy that relies on hyper-parameters, an additional out-of-sample window is taken to test these weights on unseen data. Portfolios 2 and 3 do not need this extra period because they do not depend on tunable parameters. To conclude the first portfolio, it is the only one rebalanced monthly.

The second portfolio selects stocks exclusively on their ESG score. Each year, all ESG scores are ranked and stocks in the top 10% are included. This strategy serves as a pure ESG screening selection, isolating the effect of a sustainability based selection. The portfolio is rebalanced annually with change in the ESG score.

The third portfolio introduces a dynamic ESG perspective by considering the level of improvement of a company's ESG score. This strategy reflects the regulatory principle that companies demonstrating ESG progress, regardless of their starting point or sector, deserve recognition. For each stock, the year on year percentage change is computed, capturing the momentum. The top 10% of stocks are chosen to see the differences between a dynamic screening and a static one as portfolio 2. The portfolio is rebalanced yearly as the second one.

4.2.3 Computation of Results

With the portfolios determined, the next step is to compute the actual returns generated by each strategy and account for the costs of trading. To compute portfolio returns efficiently, the long format data is pivoted back into a wide format, producing two matrices: a return matrix containing monthly returns and a position matrix containing binary indicators. The portfolio return is obtained by a single vectorized operation and the row wise sum of this product is divided by the number of active positions in that month, finding the return of the equal weighted portfolio return.

Transaction costs depend on the number of trades executed each month or year. To identify these trades, the position matrix is subtracted across consecutive months and the total number of trades in a given month is the sum of all position changes across tickers.

5 Empirical results

The performance of each portfolio is assessed along three dimensions. The first dimension is the return and risk statistics, for each portfolio and the benchmark, the annualized mean return, annualized volatility and annualized Sharpe ratio are reported.

The second dimension is the drawdown, where a time series of net return is constructed to establish the metrics. The first step is to identify the peak and then compute the percentage decline from that peak. The key results from this analysis include the maximum drawdown value, the date at which the trough occurred, the recovery date at which the cumulative wealth surpassed the prior peak and the duration of the drawdown episode.

The last dimension is the decomposition of the returns into systematic and idiosyncratic components. The first model is the Capital Asset Pricing Model (CAPM), which regresses the net return on the market excess return:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i \left(R_{m,t} - R_{f,t}\right) + \varepsilon_{i,t} \tag{5.1}$$

The intercept α measures the average return earned on top of the market exposure.

The second model is the Fama-French five-factor, which extends the CAPM with 4 additional factors:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i \left(R_{m,t} - R_{f,t}\right) + s_i\, SMB_t + h_i\, HML_t + r_i\, RMW_t + c_i\, CMA_t + \varepsilon_{i,t} \tag{5.2}$$

Here, SMB captures the size effect (small minus big), HML the value effect (high book-to-market minus low), RMW the profitability effect (robust minus weak), and CMA the investment effect (conservative minus aggressive). The alpha explains the return component that is not explained by these five-factor.

Both models are estimated with the ordinary least squares with Newey-West heteroskedasticity and autocorrelation consistent standard errors, using a lag truncation of three months. This correction takes into account the fact that monthly return series may show correlation and non-constant variance, ensuring the statistical significance of the coefficients (Mei, 1994).

5.1 Portfolio 1: Multifactor ESG Integration

Portfolio 1 is the only portfolio first analysed in an in-sample period; this step is essential to address the research question, because this period shows how the portfolio return and volatility change when the ESG score is integrated with different weights.

The in-sample analysis allows the portfolio construction parameters to be calibrated through a sensitivity analysis, identifying the combination of return, risk, and ESG scores that maximizes performance. In this sense, the in-sample period provides the upper bound of the strategy's performance, as it reflects results obtained under optimised conditions.

The insights obtained in this phase serve as a reference point for the subsequent out-of-sample analysis.

5.1.1 In-Sample Analysis

As previously explained, portfolio 1 selects the top 5% of stocks according to the composite score of risk, return and ESG. The weight assigned to each score is decided with a sensitivity analysis, and the optimal solution is:

$$\text{Final\_Score} = 0.1 \cdot \text{return\_metric} + 0.8 \cdot \text{risk\_metric} + 0.1 \cdot \text{ESG\_score} \tag{5.3}$$

Over the period from January 2018 to December 2024, the portfolio holds on average 27 stocks per month. The results of the strategy compared to the benchmark are the following:

Metric	**Portfolio 1 (Gross)**	**Portfolio 1 (Net)**	**Benchmark (STOXX 600)**
Cumulative Return	169.21%	140.13%	31.26%
Annualized Mean Return	18.19%	16.54%	5.11%
Annualized Volatility	28.51%	28.47%	15.19%
Sharpe Ratio	0.64	0.58	0.34

Table 5.1: Performance comparison: Portfolio 1 in-sample

Examining the monthly return distribution, the best month generated a net return of 38.27% while the worst one produced a loss of 28.91%. The modest dimension of the portfolio amplifies the idiosyncratic risk but the positive months are larger in magnitude than the negative ones, producing a strong cumulative result. The distribution exhibits positive skewness, meaning that extreme positive returns occur more times than negative ones. However, the distribution is a fat-tailed one, showing the limitation of a concentrated portfolio.

The maximum drawdown of the portfolio is -34.37% with the trough occurring in April 2020 at the peak of the COVID-19. The portfolio was able to recover only in December of the same year, implying a duration of 8 months. By comparison, the benchmark experienced a shallower maximum drawdown of -24.05% recovered during a similar window. The fact that the maximum drawdown is bigger than the benchmark is not surprising given the portfolio's higher volatility. The positive aspect is the relatively fast recovery, which reflects the quality of the underlying holdings. A portfolio formed on Maximum Drawdown and CVaR criteria tends to hold stocks that have demonstrated resilience through past drawdowns, and this characteristic appears to have translated into a faster than average recovery during the 2020 episode.

Regressing the net return on the market excess return from the Fama-French dataset yields a CAPM α of 1.64% per month, though this estimate is significant only at 10% (p-value= 0.074). The market beta is negative, suggesting that the portfolio's returns move, on average, in the opposite direction of the market but it is statistically insignificant (p-value= 0.135). The R-squared is low, indicating that market movements explain almost none of the variation in the portfolio's returns. The full regression is reported in the appendix.

Extending the regression to the five-factor model yields a higher α of 1.80%, statistically significant at the 5% level (p-value= 0.022). The market beta remains negative, while none of the four additional factors are individually significant.

The central finding from this factor analysis is that the five-factor fail to explain the returns of the portfolio. This suggests that the composite scoring methodology captures a return dimension that is distinct from the standard Fama-French factors.

5.1.2 Sensitivity Analysis

The composite score that creates the portfolio is the weighted average of three components and the optimal weights were found with a sensitivity analysis. All possible triplets were tested at increments of 5 percentage points. The result of the analysis reveals a clear pattern, portfolio performance is mainly driven by the risk score.

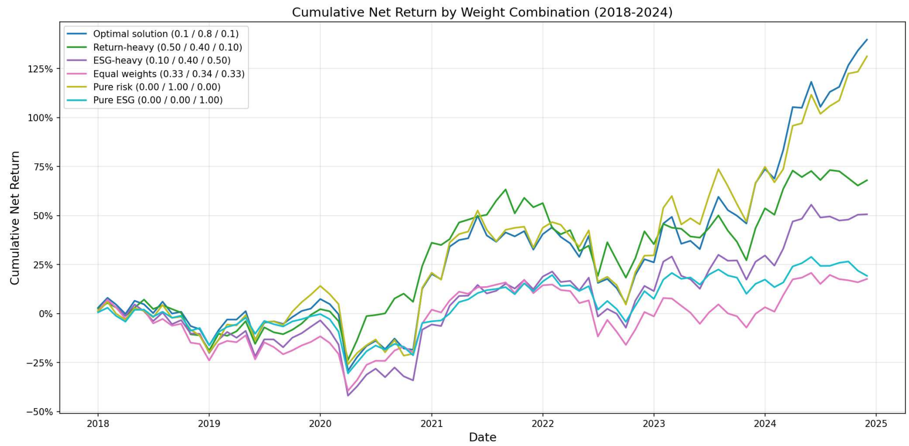

Figure 5.1: Sensitivity cumulative paths

To isolate the contribution of each component, three "pure" strategies were examined, each allocating 100% to a single dimension. The pure risk strategy is the only one that achieves a similar result to the optimal one, but nevertheless, it shows lower cumulative return and a bigger maximum drawdown.

The heatmap of cumulative returns also confirms that the risk metric is the one that drives returns.

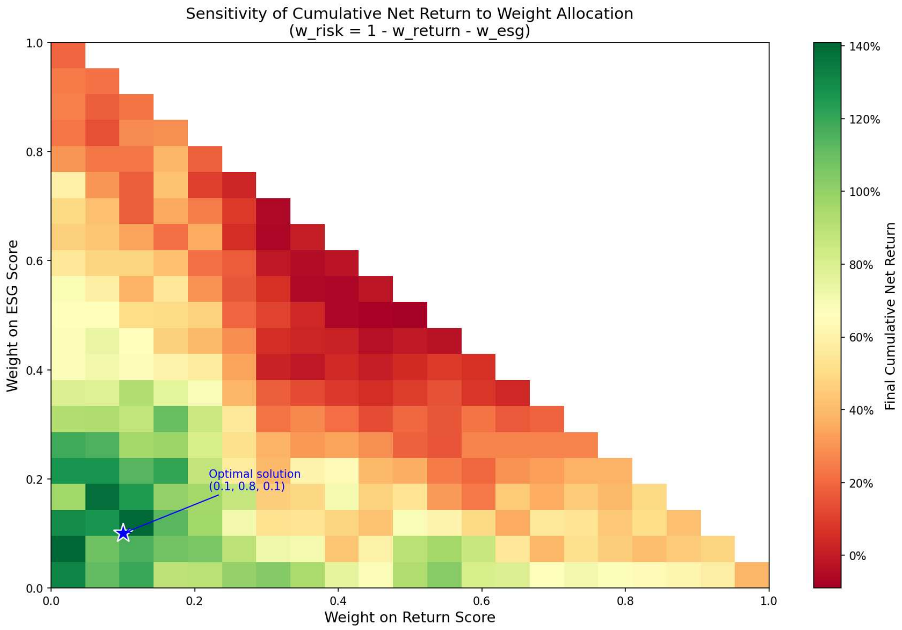

Figure 5.2: Sensitivity heatmap

The colour runs from red in the upper right, corresponding to high return and ESG weights, to deep green in the lower left region, where the risk metric dominates. The shape of the optimal region is roughly triangular, meaning that the strategy is insensitive to how the non-risk weight is split between return and ESG metric, as long as the risk weight remains high (75% of the total).

5.1.3 Out-of-Sample Analysis

The results delivered by portfolio 1 in the in-sample period are the core of this thesis, because they show the potential of the strategy, this out-of-sample window is an additional period intended to examine how the strategy behaves with unseen data. The weights identified by the sensitivity analysis are used in the out-of-sample period and given that the ESG score data is not available for all the stocks for this period, the ESG scores available in December 2024 are carried forward in the out-of-sample period.

The portfolio holds an average of 26 stocks and the results in the out-of-sample are the following:

Metric	**Portfolio 1 (Gross)**	**Portfolio 1 (Net)**	**Benchmark (STOXX 600)**
Cumulative Return	18.68%	17.38%	18.83%
Annualized Mean Return	14.33%	13.45%	14.60%
Annualized Volatility	10.89%	10.99%	12.35%
Sharpe Ratio	1.04	0.95	0.94

Table 5.2: Performance comparison: Portfolio 1 out-of-sample

The portfolio follows the benchmark in cumulative return but it delivers a lower volatility (10.99% versus the 12.35%). The higher Sharpe ratio observed in the in-sample survives the transition to the unseen data (0.95 versus 0.94).

Examining the monthly distribution, the best month generates a net return of 6.44% while the worst produces a loss of -4.85%.

The maximum drawdown of the portfolio is -5.56% with a four-month recovery, by comparison, the benchmark experiences a bigger maximum drawdown of -7.65% and a longer recovery time.

The CAPM delivers also a positive alpha of 0.51% per month but it is not significant (p-value= 0.699). Moreover, this regression delivers a positive market beta of 0.30 but it is not significant (p-value= 0.498).

The Fama-French regression delivers a positive alpha of 1.26% and as the alpha in the CAPM is not significant (p-value= 0.3). In this case the market beta returns negative at -0.20 but it is not distinguishable from zero (p-value=0.663). Among the four additional factors, HML is 0.57 and is significant (p-value=0.032), suggesting that the strategy concentrates on value stocks.

It is important to highlight that the 13-month sample gives the test very low statistical power, so the non significance should be read as an absence of evidence rather than an absence of alpha.

5.2 Portfolio 2: ESG Screening Strategy

Having established that risk metric is the dominant driver of portfolio 1's performance, a natural question is what happens when all the financial metrics are removed and the stocks are selected only by the ESG score. Portfolio 2 addresses this question, and it holds the top 10% of the stocks in the index for their shifted ESG score. This strategy represents a pure sustainability screen and serves as a benchmark for isolating the value of the ESG information.

Over the evaluation period, portfolio 2 holds on average 56 stocks, varying between 52 and 59. The results of the strategy compared to the benchmark are the following:

Metric	**Portfolio 2 (Gross)**	**Portfolio 2 (Net)**	**Benchmark (STOXX 600)**
Cumulative Return	31.58%	19.27%	31.26%
Annualized Mean Return	5.64%	4.24%	5.11%
Annualized Volatility	18.43%	18.51%	15.19%
Sharpe Ratio	0.31	0.23	0.34

Table 5.3: Performance comparison: Portfolio 2

The situation of portfolio 2 is completely different from the one of portfolio 1, in fact it underperforms the passive benchmark and this gap can be labelled as the cost of a pure ESG screening.

The monthly net distribution shows that the best month generated a net return of 20.58%, while the worst month shows a loss of 21.85%. The strategy has more positive months than negative ones but the magnitude of the positive is smaller than the magnitude of the negative ones. The tighter distribution reflects the benefit of the higher diversification, on average 56 stocks versus the 27 of portfolio 1, as well as the nature of the ESG screen itself. High ESG companies tend to be large, mature and they exhibit lower volatility, and this is a desirable property for a risk-averse investor, but it also limits the upside potential that drives performance.

The maximum drawdown is -28.95%, a better result compared to portfolio 1 that showed a -34.37%. This shallower drawdown reflects the broader diversification with 56 stocks, however the recovery period takes twelve months, so four months more than portfolio 1. This slower recovery is consistent with the absence of risk based measurement to screen the stocks to select.

The CAPM regression yields an alpha of 0.62% which is not significant (p-value= 0.255) and a market beta of -0.26 that is significant (p-value= 0.004). The negative beta supports the negative beta found in portfolio 1, in fact the ESG portfolio moves, on average, against the market direction. This attitude may reflect the defensive nature of high ESG companies.

The Fama-French regression confirms the results delivered by the CAPM, the alpha is 0.60% and it is still insignificant (p-value= 0.221). The market beta strengthens to -0.28 (p-value= 0.002). The four additional factors are statistically insignificant with small coefficients. This second regression confirms that pure ESG screen does not generate abnormal returns.

5.3 Portfolio 3: ESG Momentum Strategy

Portfolio 3 is an evolution of portfolio 2, rather than screening stocks on the ESG score, the strategy rewards the trajectory of their improvement, their ESG momentum.

The results of the strategy compared to the benchmark are the following:

Metric	**Portfolio 3 (Gross)**	**Portfolio 3 (Net)**	**Benchmark (STOXX 600)**
Cumulative Return	126.09%	66.85%	31.26%
Annualized Mean Return	13.47%	9.20%	5.11%
Annualized Volatility	18.76%	19.28%	15.19%
Sharpe Ratio	0.72	0.48	0.34

Table 5.4: Performance comparison: Portfolio 3

These figures place portfolio 3 in the middle between the benchmark and portfolio 1. The momentum signal, given its dynamic construction, generates better performance than the simpler static screen. The portfolio holds on average 58 stocks per year, and this is the reason of the similar volatility with portfolio 2, they hold the same number of stocks.

Also, The monthly net return distribution is similar in shape to portfolio 2, the best month generated a net return of 21.9% while the worst month showed a loss of 24.2%. As in the other portfolios, positive months are more than negative ones.

The maximum drawdown reached -27.49% with the peak occurring in September 2021 and the trough in October 2022. The portfolio recovered only in June 2024, so approximately two years and 9 months. The recovery period is substantially longer than the other two portfolios.

The magnitude of the drawdown reflects the portfolio's intermediate characteristics: larger and more diversified than portfolio 1, but not as broadly diversified as the benchmark itself.

The regressions delivered similar results of the other portfolios. The CAPM regression yields an alpha of 0.98% which is significant at the 10% confidence level (p-value= 0.087). The market beta is always negative and is -0.2094 and it is significant (p-value= 0.031). The negative beta mirrors the findings of the other two portfolios, reinforcing the interpretation that ESG portfolios in this universe (STOXX 600) tend to be negatively related to the market.

The Fama-French regression produces a comparable alpha, but no longer significant (p-value= 0.113) and strengthens the market beta to -0.2281, always significant (p-value= 0.028). Among the four additional factors, none reaches the statistical significance. Therefore, the ESG momentum signal generates abnormal returns only in the CAPM regression with a significance level of 10%.

5.4 Comparative Analysis

To conclude efficiently the empirical part, all the results obtained are reviewed and compared across strategies, to give a full picture. Across the three portfolios, in the same period, a clear hierarchy emerges:

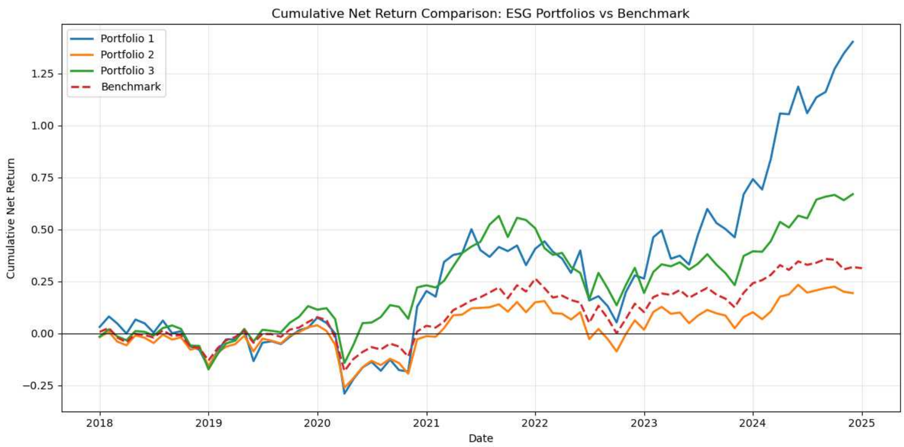

Figure 5.3: Cumulative return comparison

All four series share the same April 2020 trough, but diverge in the recovery. Portfolio 1 and 3 separate from portfolio 2 and the benchmark from mid-2020 onward, and they continue to widen this gap.

Regarding the maximum drawdown, three out of four portfolios reach the trough in April 2020, the graph shows the overall picture:

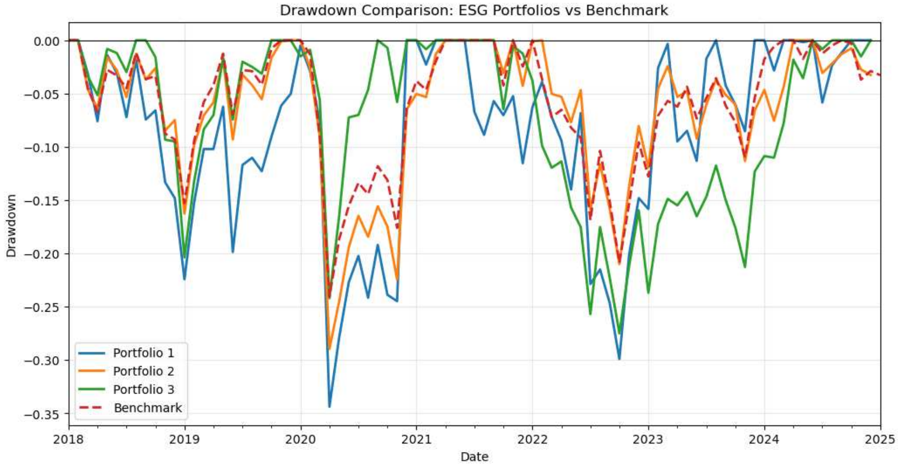

Figure 5.4: Maximum drawdown comparison

The magnitude and the recovery time differ for each portfolio, showing that the fastest portfolio to recover is portfolio 1, which is also the one with the largest loss.

The last comparison concerns the regression models and they reveal a consistent pattern across all three portfolios, a negative market beta (only the market beta in the CAPM of the out-of-sample in portfolio 1 is positive but not significant) and low explanatory power from the standard risk factors. The distinctive part of the regressions is the alpha of portfolio 1 (in-sample) and 3 that is statistically significant, respectively at 5% and 10% significance levels.

6 Discussion

This chapter interprets the results presented in Chapter 5 and compares them with the results showed in the literature review of Chapter 3. The aim is to extract the economic meaning and to examine whether they confirm or contradict the evidence on ESG investing.

The discussion is structured in nine sections. The first theme, that is contained in the first three sections, is dedicated to the analysis of returns. The second macro argument, developed in the following three sections, explains the risk characteristics of the three strategies. The third macro argument, always covered in three sections, develops a broader interpretation of the results with a behavioural perspective and concludes by setting out the limitations and the future research of the study.

6.1 Main Findings

The empirical analysis delivers a clear hierarchy. Portfolio 1, which integrates ESG as a marginal component inside a composite score, outperforms the benchmark with a wide margin in the in-sample period, producing a net cumulative return of 140.13% against the 31.26% of the benchmark and generating an alpha of 1.80% per month that is statistically significant at the 5% level. Moreover, also in the out-of-sample period the portfolio keeps a better risk-adjusted performance and a better drawdown behaviour than the benchmark so the results delivered in the in-sample period are confirmed in the out-of-sample.

The situation of the other two portfolios, which select stocks only on the basis of the ESG score, is more differentiated. Portfolio 2, the static screening, underperforms the benchmark with a net cumulative return of 19.27%. Portfolio 3, the dynamic screening, moves in a different direction in fact it produces a net cumulative return of 66.85% and an alpha in the CAPM that is significant at the 10% significance level. It is important to note that, when the regression is extended to the Fama-French model, the magnitude of the alpha is preserved but its significance is lost.

The central message of these results is that ESG metrics do not pay in a uniform way regardless of how they are implemented. A static screen produces the cost observed in portfolio 2, while a dynamic signal built on the momentum of the score delivers a cumulative return that exceeds the benchmark. This message is important because it reframes the question of "does ESG pay?" into a deeper question that is "how should ESG be incorporated in portfolio construction?". The empirical evidence taken in this thesis suggests that anchoring stock selection on risk metrics and adding ESG as a secondary criterion can preserve performance and deliver returns, whereas a process that puts ESG as a primary filter in a static way pays a huge cost and that an ESG momentum signal sits in between, in terms of cumulative return and risk-adjusted performance.

A point that deserves emphasis in the first portfolio is that the best performing strategy among all possibilities is not an ESG-free strategy. Portfolio 1 achieves the highest cumulative return when ESG is inside the construction of the portfolio, this result therefore states that ESG information is not irrelevant to portfolio construction. On the contrary, the empirical evidence is consistent with the interpretation that ESG metrics add information to the portfolio, and that this information has an economic value.

The key discovery is that ESG information is valuable as a complement to financial and risk metrics but not as a replacement for them. The failure of portfolio 2 is not a failure of the ESG signal, it is a failure of the architectural choice to rely only on a static signal alone. The same ESG score that underperforms when used in a static way becomes a useful contributor to performance when it is included into a risk and return composite score or when its momentum is studied.

6.2 Interpretation of Alpha

The most interesting result is the Fama-French five-factor alpha of the first portfolio in the in-sample period, which is statistically significant at the 5% level. This positive alpha is then confirmed in the out-of-sample period but in this case is not statistically significant. The natural temptation is to interpret this return as an ESG alpha, however this is not the meaning of it.

The sensitivity grid demonstrates that the performance of the portfolio is mainly driven by the risk score, every high performing combination gives at least 75% of its weight in the risk dimension. The pure risk strategy, which gives 100% of the weight on the risk score, produces a cumulative return and a Sharpe ratio that are only marginally below the optimal solution.

The conclusion is clear, the ESG component contributes to portfolio 1 at the margin but is not the mechanism that generates the alpha. What the regression fails to explain is not an ESG premium but a tail-risk selection premium.

Portfolio 3 also produces an alpha of 0.98% statistically significant at the 10% level, while the Fama-French regression makes this alpha disappear, nevertheless two elements deserve attention. First, the magnitude of the alpha is stable across the two regressions, suggesting that the abnormal return is not due to a misspecification of the CAPM. Second, none of the additional factors of the Fama-French is significant, so the alpha is not absorbed by a single factor.

This interpretation provides a bridge with the literature review. Pedersen et al. (2021) argue that ESG information can generate abnormal returns when the market fails to price the consequences of ESG exposure. Their empirical estimates show that governance information improves the maximum Sharpe ratio by roughly 12% relative to ignoring it. The ESG score used in this thesis is built also with the governance pillar and this design choice means that the ESG component is partially a governance signal. This result is therefore consistent with Pedersen et al. (2021), but it is not a direct replication because, as previously said, the ESG contribution is only a marginal factor for portfolio 1 and the ESG-momentum of portfolio 3 is not robust to the Fama-French adjustment.

The second link is with Van Duuren et al. (2016), their survey shows that ESG integration takes the form of an extension of fundamental investment rather than a standalone strategy. In their framework, ESG information is treated as an additional risk and value signal. Portfolio 1 follows exactly this idea, where the ESG score sits between the return and risk metrics inside a composite ranking, and it enters the selection decision without replacing the standard signal but complementing them. The fact that Portfolio 1 generates abnormal returns while portfolio 3 generates an alpha only in the CAPM and the portfolio 2 generates no alpha at all, provides empirical support for this perspective.

6.3 The Cost of Static ESG Screening

The underperformance of portfolio 2 is the mirror image of portfolio 1's success and these results can be linked with the literature review in three ways.

The first connection is with Friede et al. (2015). Their second order meta analysis concludes that the relationship between ESG and corporate financial performance is positive or nonnegative in 90% of cases at the firm level, but they distinguish between firm-level and portfolio-level studies. The portfolio level studies, such as those analysing ESG mutual funds or ESG indices, tend to report more neutral and mixed results. Friede et al. (2015) attribute this gap to different things that they summarize as a "wash out" effect: the firm level ESG advantage is diluted by noise, by overlapping in conventional risk factors, by the positive and negative screening approaches, by the fees and transaction costs and by diversification. Portfolio 2 is a clear example of this wash out mechanism. By selecting only the top decile of the stocks ranked by the ESG score, the strategy creates a diversified portfolio of large, liquid and relatively homogeneous stocks. Under these conditions, any firm-level ESG advantage is averaged out and becomes part of the portfolio noise. The 12 percentage points by which the portfolio lags behind the benchmark can be interpreted as the empirical cost of the wash out effect described by Friede et al. (2015).

The second reconciliation is with Pedersen et al. (2021). Their ESG-efficient frontier framework shows that, for any given level of ESG score, there is a maximum obtainable Sharpe ratio, and this maximum is typically lower than the Sharpe ratio of the unconstrained mean-variance frontier. The act of restricting the possibilities of investment to only high ESG firms shifts the investor to a lower point on the Sharpe axis, and the size of this shift depends on how binding the ESG constraint is. Portfolio 2 is an implementation of this constraint, it eliminates 90% of the stocks only based on the ESG score, and its observed Sharpe ratio drops from 0.34 (the benchmark Sharpe ratio) to 0.23. This drop is a direct quantification of the cost that Pedersen et al. (2021) have predicted.

The thesis therefore contributes to quantifying the sacrifice in Sharpe ratio imposed by a pure top decile ESG screen in the European universe.

The last link is with Van Duuren et al. (2016). Their survey states that ESG integration should be an extension of fundamental analysis, and the deduction is that a portfolio that uses ESG as its only selection criterion departs from the philosophy of practitioners. Its underperformance is therefore not a failure of ESG information but the expected consequence of using ESG in a way that the literature has already flagged as non adequate.

Reading all together, the results of portfolio 2 are fully consistent with the three papers.

6.4 The Systematic Negative Beta

A common feature of all the portfolios, when the market beta is statistically significant, is the negative coefficient. All three portfolios show a beta of -0.20 or more in the CAPM and of -0.22 or more in the Fama-French regression, with statistically significant p-values at least at the 5% level. The consistency of this finding across three strategies with different selection rules, different turnovers and different ESG integrations is noticeable and it deserves an interpretation.

The first interpretation relates to the composition of the selected stocks. High ESG firms within the STOXX Europe 600 tend to be large, mature companies operating in relatively defensive sectors such as utilities, healthcare and consumer staples. These sectors have exhibited lower sensitivity to market fluctuations. In this sense the stocks selection make all the three portfolios converge toward a defensive allocation, whether driven by ESG screening or a composite score.

This sectoral and firm characteristic bias provides a first explanation for the negative market beta.

The second explanation is linked to Pedersen et al. (2021). In their equilibrium framework, the coexistence of ESG unaware, ESG aware and ESG motivated investors produces a change in the demand that moves the valuations of high ESG and low ESG stocks. As ESG motivated investors put their capital on high ESG firms, the prices of these firms go up and their expected returns go down, while the prices of low ESG firms go down and their expected returns rise. As a consequence, a portfolio with only high ESG stocks moves away from the high beta, the high expected return segment of the market and goes to the low beta, the low expected return segment. The direct empirical result is the negative market beta observed in the three portfolios of the thesis.

After the theory behind the result is explained, it is important to focus on the practical implication of this finding. A portfolio with a negative beta, on average, moves in the opposite direction of the market and from the perspective of an investor who holds an exposure with the market, the three ESG portfolios can be used as partial hedges against market downturns. Even the worst of the three portfolios can contribute positively to the creation of a diversified portfolio given its negative correlation with market returns that reduces the total volatility. This perspective changes the idea of "cost" of pure ESG screening, in fact the 12 percentage point gap with the benchmark can be interpreted as the price paid for a defensive, negative beta exposure.

6.5 Drawdown Dynamics and Portfolio 1 Alpha

The drawdown analysis provides a link between the alpha of portfolio 1 in the in-sample period and the mechanism that generates it. The four series do not share a single maximum drawdown episode, portfolio 3 is the only one that does not reach the trough in April 2020 in the COVID shock but it experiences its maximum drawdown in October 2022. Portfolio 1 experiences the worst drawdown, reaches -34.37%, but it recovers in 8 months. By contrast, the other two portfolios have shallower drawdowns of around 30%, with portfolio 2 recovering in 12 months and portfolio 3 recovering in 33 months.

This pattern gives an insight, the portfolio with the deepest drawdown is also the portfolio with the fastest recovery. This asymmetry is not random, because portfolio 1 selects stocks also based on historical maximum drawdown and historical expected shortfall, both of which summarize tail loss behaviour. A stock that receives a high rank on these measures is a stock that has shown resilience in tail shocks and it has recovered quickly. The 2020 episode is a case of tail shock, and the portfolio's stocks responded to the shock as the selection criterion predicted: with a big initial drop, followed by a swift rebound. The deeper initial drop is explained by the lower diversification of the portfolio, which on average holds 27 stocks and therefore carries high idiosyncratic risk, while the faster recovery is explained by the quality of the holdings, which had already demonstrated resilience in prior drawdown episodes and did again in 2020.

By contrast, portfolios 2 and 3 select stocks on a criterion that is indifferent from tail risk management. Their shallower drawdowns reflect the lower volatility of a large portfolio, but their slower recoveries reflect the absence of a mechanism of selection that would favour the selection of stocks with strong rebound characteristics.

This interpretation has a valuable insight for the reading of the portfolio 1 alpha. The statistically significant alpha is partly generated in this period, when the stocks selected for tail resilience start to outperform the others, the figure below shows the paths of the 4 portfolios during the covid:

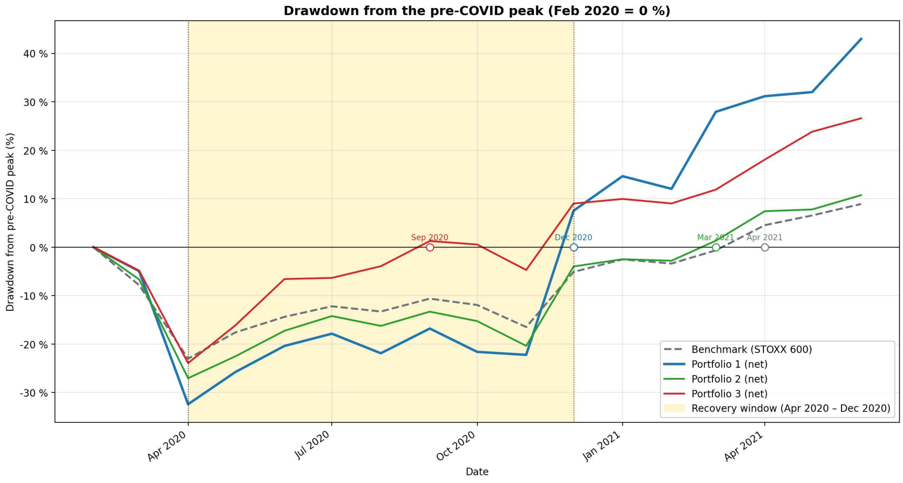

Figure 6.1: Drawdown during Covid

An important part of the cumulative outperformance of portfolio 1 is earned in the 8 months of recovery, during which the portfolio's stocks rebounded more than the others (In December 2020 the portfolio created a gap of 12% with the benchmark). The alpha is therefore concentrated in time rather than evenly distributed across the whole period, and this concentration raises a legitimate question of whether the alpha persists even without a shock. The thesis does not claim that the strategy would generate an equivalent alpha in a different environment, and the 2020/2021 crisis falls within the scope of the limitations.

6.6 Sensitivity Analysis

The sensitivity analysis presented in Chapter 5 tested 231 distinct weight combinations across the return, risk and the ESG score. Before discussing what the grid shows, it is important to state the purpose of the sensitivity analysis. This analysis is an ex-post analysis of the full weight space. The exercise does not estimate the performance that an ex-ante investor could have earned by choosing those weights in this period, and it would be incorrect to read it as a forward recommendation. The purpose of the sensitivity analysis is to see the shape of the performance surface, and the structural relationship between the three dimensions and the realised performance of the strategy. The shape of the surface is a property of the data and it does not suffer from look-ahead bias that affects the best combination and in this sense the sensitivity reveals how the best portfolio behaves.

The first finding is the dominance of the risk weight. Every combination that achieves a Sharpe ratio close to the optimal solution puts at least 75% of its weight on the risk score, and also the pure risk strategy delivers a performance that is marginally below the optimum. In the composite score, the return and ESG dimensions contribute at the margin to improve the work of the risk metrics.

The second insight is the indifference in the exchange between the return weight and the

ESG weight within the high performance region. Once the risk weight is held around 75%, the remaining weight can be redistributed between the return and ESG dimensions keeping a high Sharpe ratio. This flatness reflects the fact that the return and ESG rankings carry overlapping information and so after the risk anchor has done its job, the rest can be fulfilled by one of the two.

The third and most important insight, follows directly from the second, and is about the existence of a broad region over which the ESG weight can be changed without a cost in the performance, this can be described as a region of indifference. An investor after putting the risk weight as dominant, can then set the ESG weight anywhere between 0% and 25% without affecting the performance. Recalling the regulatory framework, investment firms are required to match financial products with the sustainability preference of the investors, and a selection rule that allows the ESG intensity to change without compromising the performance is operationally attractive.

The sensitivity grid also supports the theory of Pedersen et al. (2021). Their ESG-efficient frontier implies that, for any given risk anchored portfolio, there is a range in which the Sharpe ratio is insensitive to the ESG level, and a range where the Sharpe ratio deteriorates. The triangular high performing region in the sensitivity grid is an empirical visualisation of the flat segment of the ESG efficient frontier for the strategy of this thesis.

6.7 Behavioural Interpretation of ESG Strategies

The economic interpretation made in sections 6.3 and 6.4 treats the underperformance of portfolio 2 as a cost due to the structure of the selection method. This interpretation is correct in a framework that takes the return as the only goal, but it is incomplete in a framework that takes investor preferences as heterogeneous and multi-dimensional.

The paper of Shefrin and Statman (2000) offers an alternative reading that links the underperformance with the existence of investor demand for pure ESG strategies. The paper argues that investors do not construct portfolios according to a single mean-variance or return objective but according to a layered structure in which each layer corresponds to a distinct goal. Recalling the two opposite layers, the lower one is designed for capital preservation while the upper is for aspiration and high payoff outcomes. Portfolios constructed under this idea are mean-variance inefficient, this is because they respond to the multiple goals that investors have. The authors suggest that investors can rationally accept a lower Sharpe ratio if the portfolio satisfies a goal that the Sharpe ratio does not measure.

Applied to the thesis, an investor with an aspirational ESG goal can rationally accept the roughly 12% of cumulative underperformance observed in portfolio 2. The gap with the benchmark is not a failure of the strategy but is the price paid for satisfying the aspirational layer, and the strategy is a rational choice for an investor who values more this layer than returns.

This behavioural interpretation interacts with the negative market beta finding discussed in section 6.4. The portfolio simultaneously satisfies the aspirational layer, with the ESG exposure, and the security layer, with the negative correlation with the market. A single strategy therefore contributes to two different layers that Shefrin and Statman have identified, and under this light, the underperformance is more than a simple cost. The strategy gives up returns but delivers both an aspirational goal and a downside protection, and an investor who values these dimensions may rank these strategies above the benchmark despite the lower cumulative return.

Portfolio 3 requires a different behavioural reading because it does not present an underperformance but instead it shows the slowest recovery from the maximum drawdown. Within Shefrin and Statman (2000)'s framework, portfolio 3 can be classified as a strategy that combines an aspirational layer with a good return profile but that imposes long periods of losses. The behavioural acceptance of portfolio 3 is therefore about losses tolerance and willingness to bear this exposure. For investors whose aspirational layer is driven by ESG objectives and who do not have strict control on the intermediate losses, portfolio 3 offers an attractive combination.

This section does not invalidate the financial interpretation given in the previous section; it complements it. The rational interpretation explains why portfolio 2 underperforms and why portfolio 3 exhibits a particular drawdown, while the behavioural interpretation explains why these strategies are still attractive.

6.8 Limitations of the Study

The empirical results discussed should be read in light of a number of limitations, and each of these opens the possibility for future research.

6.8.1 Reliance on a Single Rating Provider

The ESG score used in the thesis is provided by LSEG (Refinitiv). Berg et al. (2022) show the disagreement across six ESG ratings, and their core result is that firms highly rated by one agency can receive average or below the average score by other agencies, and the consequence is that any result obtained with a single rating provider is conditional on the methodology of the single provider.

This limitation touches the three portfolios differently. Portfolio 1 is the most robust because the ESG score accounts for only 10% of the composite selection, so a replication of the portfolio with another rating should produce a similar portfolio. By contrast, the other two portfolios are very sensitive to the choice of the rating agency because their methodology is based only on the ESG score. The table below taken by Berg et al. (2022) shows the relative contribution of scope, measurement, and weight to the ESG rating divergence:

Panel A: Rater pairs				
		Scope	**Measurement**	**Weight**
KLD	Sustainalytics	18%	69%	13%
KLD	Moody's ESG	31%	59%	10%
KLD	S&P Global	20%	68%	11%
KLD	Refinitiv	22%	63%	15%
KLD	MSCI	81%	17%	3%
Sustainalytics	Moody's ESG	20%	64%	16%
Sustainalytics	S&P Global	22%	70%	8%
Sustainalytics	Refinitiv	12%	66%	22%
Sustainalytics	MSCI	68%	30%	2%
Moody's ESG	S&P Global	41%	56%	3%
Moody's ESG	Refinitiv	19%	79%	2%
Moody's ESG	MSCI	66%	41%	-6%
S&P Global	Refinitiv	23%	74%	3%
S&P Global	MSCI	59%	52%	-10%
Refinitiv	MSCI	68%	38%	-7%
Average		38%	56%	6%
Panel B: Rater averages				
		Scope	**Measurement**	**Weight**
KLD		34%	55%	10%
Sustainalytics		28%	60%	12%
Moody's ESG		35%	60%	5%
S&P Global		33%	64%	3%
Refinitiv		29%	64%	7%
MSCI		68%	36%	-4%

Therefore, the results of portfolio 2 and 3 are the results of the LSEG methodology and they can achieve different results with different rating provider.

6.8.2 Sample Period and Time Horizon

The sample period covers seven years, from January 2018 to December 2024, and it is marked by the COVID-19 shock. As section 6.5 has already explained, a share of the portfolio 1 alpha is earned during the recovery phase, and the robustness of the result over a different time span cannot be guaranteed. However, the choice of the time period is not free; it is constrained by three considerations that are difficult to overcome.

The first constraint is the role of COVID in the risk selection rule. The maximum drawdown and CVaR informativeness depends on the presence of at least one episode of tail stress in the sample. Shortening the period to exclude 2020, or changing the in-sample period and using as an out-of-sample period the 2018-2024, would remove the single most informative event that the selection rule uses.

The second constraint is the availability and reliability of ESG data before 2018. Recalling the regulations, the NFRD was adopted in the 2014 and applied from 2017 and before this regulation there was not a unified EU reporting framework. Therefore, any results on a pre 2017 sample would reflect the evolution of the rating methodology as much as the actual ESG performance of the firms, and would not be directly comparable with the results of the 2018-2024 window.

The third constraint is a common evaluation period across the three portfolios. The main objective of the thesis is to make a portfolio comparison, and this comparison is interpretable only if the three strategies are evaluated on the same window.

A shorter window for portfolio 1, for example a split with an in-sample period between 2018-2021 and out-of-sample 2022-2024, would have produced a result that is not a clean comparison. The decision to calibrate the composite score on the full 2018-2024 window is therefore the only decision that gives consistency and interpretability.

The cost of these constraints is that the out-of-sample period for portfolio 1 is limited to thirteen months after December 2024. This period is sufficient to define the risk-adjusted and drawdown behaviour, but it is not long enough to reject or accept the statistical evidence.

6.8.3 Equal-Weight Portfolio Construction

All three portfolios use an equal weight strategy across the selected stocks. This design is useful to compare the three strategies, isolating the effect of the selection rather than the effect of the weighting scheme. The cost of this limitation is that the portfolios do not benefit from the possible variance reduction if the weight would have been chosen with a mean-variance optimisation process, as a consequence the reported Sharpe ratios can be improved changing the weight allocate to each stock.

6.8.4 ESG Momentum Signal

The LSEG score is bounded between 0 and 100, while the momentum filter relies exclusively on the year-on-year improvement, ignoring the level reached by a firm. A firm that improves from a score of 80 to 100 is selected only once, but in the following year it cannot improve any further and it will go out from the portfolio. The same logic applies to all firms whose score has already reached a very high level, they only have a marginal room for improvement. The design of this strategy has 2 implications. First, the strategy generates a turnover that is driven by the bounded characteristic of the variable and not by the signal. Second, firms that have reached very high ratings are excluded from the portfolio, which contradicts the ESG oriented strategy.

At the same time, the momentum strategy focuses on changes rather than levels, and so the strategy allocates capital to firms that are improving their ESG profile. This exposes investors to the potential upside associated with companies that are experiencing the transition. In this sense, ESG momentum can be interpreted as a proxy of corporate quality. Therefore, while the bounded nature of the score creates mechanical turnover and exclusion effects, the signal also provides exposure to companies that are undergoing improvement opportunities.

6.8.5 Transaction Cost Assumption

The thesis assumed as transaction cost a 5 basis points per trade, which is the commission cost faced by an institutional investor that is trading large companies. This assumption does not consider the bid and ask spread, which is typically of the same order of magnitude as the commission, and the market impact of large trades. Portfolio 1, which is rebalanced monthly, is the most exposed to these frictions, and the real application of the model would compromise its net performance.

6.9 Future Research

The limitations found in the previous section are not a list of weaknesses but the basis for the future studies.

The natural extension is to extend the out-of-sample period for portfolio 1 and integrate the new ESG score in the strategy. With this approach it is possible to validate the optimal weight found in the sensitivity, and each additional year of returns would then contribute to a true result of the strategy.

A second improvement is the data provider robustness. As discussed in section 6.8.1, any result obtained with a single rating provider is conditional on the methodology of the provider, and the divergence shown by Berg et al. (2022) is too large to be ignored. A next study could replicate the three strategies using different rating providers and after the resulting portfolios can be compared in three main dimensions: the overlap in stocks selected, the correlation of net returns and the presence or not of the alpha. A further idea would combine different rating scores into an aggregated score. If the underperformance of portfolio 2 persists despite this score aggregation, it would suggest the lower returns associated with ESG screening in the European market are robust, if it disappears, it would indicate that the cause of the underperformance is driven by the LSEG rating.

A similar analysis could be conducted for the ESG momentum strategy by replicating portfolio 3 using a different rating providers. This would allow to assess whether the observed outperformance is a general characteristic of dynamic screening or a consequence of the LSEG score.

Another extension of the study is to run again the strategies under alternative weighting schemes. This analysis would quantify how much of the portfolio performance gap is attributable to the selection rule and how much to the weighting choice. The mean-variance rule is interesting because it maps the framework of this thesis in the ESG efficient frontier of Pedersen et al. (2021) and allows a quantitative comparison between empirical results and the theoretical predictions.

A future research motivated by Pedersen et al. (2021) is to use the single pillar decomposition of the ESG score. The governance pillar has predictive power in the paper, and also Van Duuren et al. (2016) confirm that practitioners allocate more attention to governance than to the other two dimensions. A follow up study could replicate the three strategies by replacing the composite score with the single pillar, and a positive result for the governance would reinforce the interpretation of section 6.2.

Taken together, these four extensions relax the limitations that constrain the current results. The current study should be used as a pilot study where the principal outputs provide the inputs for a richer and more rigorous investigation on how ESG information should be integrated in the European market.

7 Conclusion

The thesis aims to answer the following question:

> *“In the EU equity market, how do portfolio returns and volatility differ when ESG scores are used as a primary screening criterion versus as a weighted component within a multifactor framework, and how sensitive are these outcomes to alternative ESG weighting schemes?”*

The empirical analysis provides the answer to this research. First, the results demonstrate that ESG integration does not have a uniform impact on portfolio performance, rather, its impact depends on how it is implemented. When ESG is used as a single screening tool, as in the best in class strategy, it leads to an underperformance relative to the benchmark.

In contrast, when ESG metrics are incorporated within a multifactor framework, they contribute to portfolio construction enhancing the performance. This composite strategy delivers the strongest result, both in terms of cumulative returns and risk-adjusted performance.

The ESG momentum strategy offers an intermediate perspective. By focusing on improvements in ESG scores, it captures part of the upside associated with firms undergoing sustainability transition. This dynamic outperforms both the benchmark and the static ESG screen, suggesting that changes in ESG scores contain valuable information.

A second key finding concerns the role of risk. Across all strategies, portfolios show a negative market beta, indicating a defensive characteristic. This insight implies that ESG portfolios may serve as hedging instruments in broader portfolios.

The sensitivity analysis reinforces the role of risk in portfolio construction. Performance is mainly driven by risk metrics, while ESG weights can vary within a range without affecting the outcomes. This result has a practical implication, investors can adjust the ESG exposure to match sustainability preferences without sacrificing financial performance. Moreover, the sensitivity shows that the best strategy includes the ESG score in the optimal weighting, highlighting the economic value of the information inside the score.

From a broader perspective, the findings contribute to changing the debate question from "does ESG pay?" to "how should ESG be integrated?". The evidence suggests that ESG metrics are valuable when used as an additional information within a composite score, but they are less effective when applied as a rigid constraint.

Finally, the thesis has several limitations, including the reliance on a single ESG rating provider and a particular time period analysed. These constraints set the basis for future research, particularly to test the robustness across different ESG data providers and exploring different time horizons.

In conclusion, ESG information has economic value in portfolio construction, but its contribution is conditional on how it is integrated. The most effective approach is not to replace financial metrics with ESG criteria, but to integrate them. This insight aligns both with empirical evidence and with the practice of asset management, where ESG is increasingly treated as an extension of fundamental analysis rather than a substitute investment philosophy.

Appendix

Appendix A: Market Overview and Supporting Figures

Parameter	**Details**
Market revenue in 2024	USD 7,204.0 million
Market revenue in 2030	USD 20,484.7 million
Growth rate	19.9% (CAGR from 2025 to 2030)
Largest segment	ESG Integration
Fastest growing segment	Green Bonds
Historical data covered	2018–2023
Base year for estimation	2024
Forecast period covered	2025–2030
Quantitative units	Revenue in USD million
Market segmentation	ESG Integration, Impact Investing, Sustainable Funds, Green Bonds
Key market players worldwide	BlackRock Inc; BNP Paribas Act. Cat.A; The Goldman Sachs Group Inc; JPMorgan Chase & Co; Morgan Stanley; State Street Corporation; UBS Group AG

Table 1: Europe Sustainable Investment Market Overview

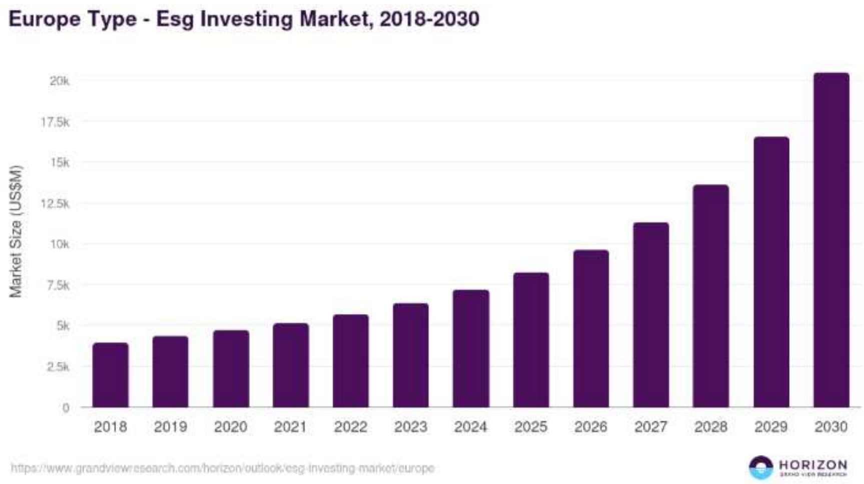

Figure 1: Europe ESG Investing Market Overview

Appendix B: ESG and Regulatory Supporting Tables

Table 2: EU Taxonomy Eligible Industries and Their Description

Industry	**Description**
Forestry	Activities related to the management, conservation, and sustainable use of forests, including timber production, afforestation, and biodiversity preservation.
Environmental protection and restoration activities	Activities aimed at preventing, reducing, and repairing environmental damage, including pollution control, ecosystem restoration, and conservation projects.
Manufacturing	Industrial production of goods (e.g., chemicals, metals, machinery), representing a major source of emissions and a key sector for decarbonisation.
Energy	Generation, transmission, and distribution of energy, including renewable energy sources and transitional systems.
Water supply, sewerage, waste management and remediation	Provision and treatment of water, wastewater management, waste collection, recycling, and environmental remediation.
Transport	Movement of goods and passengers via road, rail, air, and water, including infrastructure, with a key role in emission reduction.
Construction and real estate	Development, construction, and management of buildings, including energy-efficient and sustainable infrastructure.
Information and communication	Digital infrastructure and services such as telecommunications, data processing, and software supporting efficiency across sectors.
Professional, scientific and technical activities	Research, engineering, consulting, and technical services supporting innovation and sustainability.
Financial and insurance activities	Financial services including banking, investment, and insurance, facilitating capital allocation toward sustainable activities.
Education	Teaching, training, and research institutions contributing to human capital and sustainability awareness.

Continued on next page

Industry	Description
Human health and social work activities	Healthcare and social services aimed at improving wellbeing, inclusion, and public health.
Arts, entertainment and recreation	Cultural, creative, and recreational activities contributing to social wellbeing and quality of life.

	Total	Taxonomy-eligible	%	Taxonomy-non-eligible	%
Sales	**68,902**	**8,421**	**12.2**	**60,481**	**87.8**
Of which climate change mitigation environmental objective		8,207	11.9		
Of which pollution prevention and control environmental objective		215	0.3		
Investments (capex)	**6,006**	**1,385**	**23.1**	**4,621**	**76.9**
Of which climate change mitigation environmental objective		1,366	22.7		
Of which pollution prevention and control environmental objective		19	0.3		
Operating expenditures (opex)	**4,645**	**577**	**12.4**	**4,068**	**87.6**
Of which climate change mitigation environmental objective		536	11.5		
Of which pollution prevention and control environmental objective		42	0.9		

Table 3: EU Taxonomy for companies

Table 4: EU Taxonomy for companies

		2023		Substantial contribution criteria					
Economic activities	**Code**	**Sales**	**%**	**Mitigation**	**Adaptation**	**Water**	**Circular**	**Pollution**	**Biodiversity**
A. Taxonomy-eligible activities									
A.1. Taxonomy-aligned									
Manufacture of batteries	CCM 3.4	847	1.2	Y	N/EL	N/EL	N/EL	N/EL	N/EL
Manufacture of energy efficiency equipment	CCM 3.5	32	0.0	Y	N/EL	N/EL	N/EL	N/EL	N/EL
Manufacture of soda ash	CCM 3.12	7	0.0	Y	N/EL	N/EL	N/EL	N/EL	N/EL
Manufacture of organic chemicals	CCM 3.14	200	0.3	Y	N/EL	N/EL	N/EL	N/EL	N/EL
Manufacture of plastics	CCM 3.17	10	0.0	Y	N/EL	N/EL	N/EL	N/EL	N/EL
Aligned revenue		**1,095**	**1.6**						
Of which enabling		879	1.3						
Of which transitional		217	0.3						
A.2. Eligible but not aligned									
Manufacture of hydrogen	CCM 3.10	12	0.0	EL	N/EL	N/EL	N/EL	N/EL	N/EL
Manufacture of chemicals	CCM 3.14	2,157	3.1	EL	N/EL	N/EL	N/EL	N/EL	N/EL
Manufacture of plastics	CCM 3.17	4,632	6.7	EL	N/EL	N/EL	N/EL	N/EL	N/EL
Pharmaceutical ingredients*	PPC 1.1	215	0.3	N/EL	N/EL	N/EL	N/EL	EL	N/EL
Not aligned revenue		**7,326**	**10.6**						
Total eligible		**8,421**	**12.2**						
B. Non-eligible activities									
Non-eligible revenue		60,481	87.8						
Total		**68,902**	**100.0**						

* API = active pharmaceutical ingredients.

Appendix C: Regression Results – Portfolio 1

OLS Regression Results			
Dep. Variable:	net_return_1	R-squared:	0.027
Model:	OLS	Adj. R-squared:	0.015
Method:	Least Squares	F-statistic:	2.278
Date:	Thu, 26 Feb 2026	Prob (F-statistic):	0.135
Time:	09:21:58	Log-Likelihood:	92.353
No. Observations:	84	AIC:	-180.7
Df Residuals:	82	BIC:	-175.8
Df Model:	1		
Covariance Type:	nonrobust		

	coef	std err	t	P> $\|t\|$	[0.025	0.975]
Intercept	0.0164	0.009	1.809	0.074	-0.002	0.034
mkt	-0.2629	0.174	-1.509	0.135	-0.609	0.084

Omnibus:	26.159	Durbin-Watson:	1.968
Prob(Omnibus):	0.000	Jarque-Bera (JB):	108.003
Skew:	0.791	Prob(JB):	3.53e-24
Kurtosis:	8.325	Cond. No.:	19.6

Standard errors are heteroskedasticity and autocorrelation robust (HAC) using 3 lags and without small sample correction.

Table 5: CAPM Regression – Portfolio 1 (In-Sample)

OLS Regression Results			
Dep. Variable:	net_return_1	R-squared:	0.046
Model:	OLS	Adj. R-squared:	-0.015
Method:	Least Squares	F-statistic:	1.273
Date:	Thu, 26 Feb 2026	Prob (F-statistic):	0.284
Time:	09:14:26	Log-Likelihood:	93.176
No. Observations:	84	AIC:	-174.4
Df Residuals:	78	BIC:	-159.8
Df Model:	5		
Covariance Type:	HAC		

	coef	std err	z	P> $\|z\|$	[0.025	0.975]
Intercept	0.0180	0.008	2.297	0.022	0.003	0.033
mkt	-0.2805	0.145	-1.939	0.053	-0.564	0.003
smb	0.2474	0.386	0.640	0.522	-0.510	1.005
hml	-0.1404	0.395	-0.356	0.722	-0.914	0.633
rmw	-0.2723	0.392	-0.694	0.488	-1.041	0.497
cma	0.2002	0.517	0.387	0.698	-0.813	1.213

Omnibus:	21.817	Durbin-Watson:	1.981
Prob(Omnibus):	0.000	Jarque-Bera (JB):	76.573
Skew:	0.668	Prob(JB):	2.36e-17
Kurtosis:	7.483	Cond. No.:	68.4

Standard errors are heteroskedasticity and autocorrelation robust (HAC) using 3 lags and without small sample correction.

Table 6: Fama-French Five-Factor Regression – Portfolio 1 (In-Sample)

OLS Regression Results

Dep. Variable:	net_return	R-squared:	0.057
Model:	OLS	Adj. R-squared:	-0.029
Method:	Least Squares	F-statistic:	0.460
Date:	Thu, 16 Apr 2026	Prob (F-statistic):	0.512
Time:	14:49:37	Log-Likelihood:	29.872
No. Observations:	13	AIC:	-55.74
Df Residuals:	11	BIC:	-54.61
Df Model:	1		
Covariance Type:	HAC		

	coef	std err	z	P> $\|z\|$	[0.025	0.975]
Intercept	0.0051	0.013	0.387	0.699	-0.021	0.031
mkt	0.3048	0.450	0.678	0.498	-0.576	1.186

Omnibus:	0.562	Durbin-Watson:	2.594
Prob(Omnibus):	0.755	Jarque-Bera (JB):	0.233
Skew:	-0.305	Prob(JB):	0.890
Kurtosis:	2.759	Cond. No.:	50.9

Standard errors are heteroskedasticity and autocorrelation robust (HAC) using 3 lags and without small sample correction.

Table 7: CAPM Regression – Portfolio 1 (Out-of-Sample)

OLS Regression Results

Dep. Variable:	net_return	R-squared:	0.332
Model:	OLS	Adj. R-squared:	-0.144
Method:	Least Squares	F-statistic:	5.393
Date:	Thu, 16 Apr 2026	Prob (F-statistic):	0.0237
Time:	14:49:37	Log-Likelihood:	32.115
No. Observations:	13	AIC:	-52.23
Df Residuals:	7	BIC:	-48.84
Df Model:	5		
Covariance Type:	HAC		

	coef	std err	z	P> $\|z\|$	[0.025	0.975]
Intercept	0.0126	0.012	1.037	0.300	-0.011	0.036
mkt	-0.2002	0.460	-0.436	0.663	-1.101	0.700
smb	-0.3133	0.276	-1.136	0.256	-0.854	0.227
hml	0.5655	0.264	2.141	0.032	0.048	1.083
rmw	-0.2804	0.152	-1.840	0.066	-0.579	0.018
cma	-1.6633	0.863	-1.928	0.054	-3.354	0.028

Omnibus:	1.463	Durbin-Watson:	2.389
Prob(Omnibus):	0.481	Jarque-Bera (JB):	1.072
Skew:	0.639	Prob(JB):	0.585
Kurtosis:	2.413	Cond. No.:	172

Table 8: Fama-French Five-Factor Regression – Portfolio 1 (Out-of-Sample)

Standard errors are heteroskedasticity and autocorrelation robust (HAC) using 3 lags and without small sample correction.

Appendix D: Regression Results – Portfolio 2

OLS Regression Results			
Dep. Variable:	net_return_2	R-squared:	0.064
Model:	OLS	Adj. R-squared:	0.053
Method:	Least Squares	F-statistic:	8.163
Date:	Thu, 26 Feb 2026	Prob (F-statistic):	0.00542
Time:	09:16:58	Log-Likelihood:	130.17
No. Observations:	84	AIC:	-256.3
Df Residuals:	82	BIC:	-251.5
Df Model:	1		
Covariance Type:	HAC		

	coef	std err	z	P> $\|z\|$	[0.025	0.975]
Intercept	0.0062	0.005	1.138	0.255	-0.004	0.017
mkt	-0.2634	0.092	-2.857	0.004	-0.444	-0.083

Omnibus:	14.108	Durbin-Watson:	1.966
Prob(Omnibus):	0.001	Jarque-Bera (JB):	50.588
Skew:	0.192	Prob(JB):	1.03e-11
Kurtosis:	6.782	Cond. No.:	19.6

Standard errors are heteroskedasticity and autocorrelation robust (HAC) using 3 lags and without small sample correction.

Table 9: CAPM Regression – Portfolio 2

OLS Regression Results			
Dep. Variable:	net_return_2	R-squared:	0.067
Model:	OLS	Adj. R-squared:	0.007
Method:	Least Squares	F-statistic:	2.334
Date:	Thu, 26 Feb 2026	Prob (F-statistic):	0.0498
Time:	09:16:59	Log-Likelihood:	130.28
No. Observations:	84	AIC:	-248.6
Df Residuals:	78	BIC:	-234.0
Df Model:	5		
Covariance Type:	HAC		

	coef	std err	z	P> $\|z\|$	[0.025	0.975]
Intercept	0.0060	0.005	1.224	0.221	-0.004	0.016
mkt	-0.2792	0.090	-3.092	0.002	-0.456	-0.102
smb	0.1077	0.219	0.493	0.622	-0.321	0.536
hml	-0.0480	0.276	-0.174	0.862	-0.590	0.494
rmw	0.0865	0.213	0.407	0.684	-0.330	0.503
cma	0.0581	0.348	0.167	0.867	-0.624	0.740

Omnibus:	13.450	Durbin-Watson:	1.981
Prob(Omnibus):	0.001	Jarque-Bera (JB):	47.986
Skew:	0.110	Prob(JB):	3.80e-11
Kurtosis:	6.696	Cond. No.:	68.4

Standard errors are heteroskedasticity and autocorrelation robust (HAC) using 3 lags and without small sample correction.

Table 10: Fama-French Five-Factor Regression – Portfolio 2

Appendix E: Regression Results – Portfolio 3

OLS Regression Results			
Dep. Variable:	net_return_3	R-squared:	0.037
Model:	OLS	Adj. R-squared:	0.026
Method:	Least Squares	F-statistic:	4.662
Date:	Mon, 20 Apr 2026	Prob (F-statistic):	0.0338
Time:	10:17:17	Log-Likelihood:	125.56
No. Observations:	84	AIC:	-247.1
Df Residuals:	82	BIC:	-242.3
Df Model:	1		
Covariance Type:	HAC		

	coef	std err	z	P> $\lvert z \rvert$	[0.025	0.975]
Intercept	0.0098	0.006	1.710	0.087	-0.001	0.021
mkt	-0.2094	0.097	-2.159	0.031	-0.399	-0.019

Omnibus:	2.415	Durbin-Watson:	2.006
Prob(Omnibus):	0.299	Jarque-Bera (JB):	1.892
Skew:	-0.146	Prob(JB):	0.388
Kurtosis:	3.675	Cond. No.:	19.6

Standard errors are heteroskedasticity and autocorrelation robust (HAC) using 3 lags and without small sample correction.

Table 11: CAPM Regression – Portfolio 3

OLS Regression Results			
Dep. Variable:	net_return_3	R-squared:	0.053
Model:	OLS	Adj. R-squared:	-0.008
Method:	Least Squares	F-statistic:	1.499
Date:	Mon, 20 Apr 2026	Prob (F-statistic):	0.200
Time:	10:17:18	Log-Likelihood:	126.25
No. Observations:	84	AIC:	-240.5
Df Residuals:	78	BIC:	-225.9
Df Model:	5		
Covariance Type:	HAC		

	coef	std err	z	P> $\lvert z \rvert$	[0.025	0.975]
Intercept	0.0089	0.006	1.585	0.113	-0.002	0.020
mkt	-0.2281	0.104	-2.203	0.028	-0.431	-0.025
smb	0.2310	0.210	1.099	0.272	-0.181	0.643
hml	-0.2405	0.244	-0.985	0.325	-0.719	0.238
rmw	0.2047	0.282	0.726	0.468	-0.348	0.757
cma	0.2186	0.344	0.636	0.525	-0.456	0.893

Omnibus:	4.496	Durbin-Watson:	1.979
Prob(Omnibus):	0.106	Jarque-Bera (JB):	4.214
Skew:	-0.319	Prob(JB):	0.122
Kurtosis:	3.893	Cond. No.:	68.5

Standard errors are heteroskedasticity and autocorrelation robust (HAC) using 3 lags and without small sample correction.

Table 12: Fama-French Five-Factor Regression – Portfolio 3

Bibliography

Berg, F., Kölbel, J. F., and Rigobon, R. (2022). Aggregate confusion: The divergence of esg ratings. *Review of finance*, 26(6):1315–1344.

CFA Institute (2022). Esg investing: A practitioner's guide.

Commission, E. (2018). Action plan on financing sustainable growth. *European Commission, Brussels*.

European Commission (2020). Eu taxonomy regulation.

European Commission (2022). Corporate sustainability reporting directive (csrd).

European Union (2019). Sustainable finance disclosure regulation (sfdr).

Fink, L. (2020). A fundamental reshaping of finance. Technical report, BlackRock.

Friede, G., Busch, T., and Bassen, A. (2015). Esg and financial performance: aggregated evidence from more than 2000 empirical studies. *Journal of sustainable finance & investment*, 5(4):210–233.

Global Sustainable Investment Alliance (2022). Global sustainable investment review.

Horizon, G. V. (2024). Europe esg investing market size & outlook, 2025-2030. Technical report, Grand View Horizon.

Institute, T. C. G. (2021). What is esg? Technical report, The Chartered Governance Institute UK & Ireland.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.

Keeble, B. R. (1988). The brundtland report: 'our common future'. *Medicine and war*, 4(1):17–25.

Levy, J. S. (1992). An introduction to prospect theory. *Political psychology*, pages 171–186.

Lopes, L. L. and Oden, G. C. (1999). The role of aspiration level in risky choice: A comparison of cumulative prospect theory and sp/a theory. *Journal of mathematical psychology*, 43(2):286–313.

LSEG Data & Analytics (2023). Refinitiv esg scores methodology.

Mei, J. (1994). Application of the newey-west matrix for correction of heteroskedasticity and cross-sectional correlation. pages 91–111.

Morningstar (2023). Morningstar esg investing research.

MSCI (2023). Msci esg ratings methodology.

Pedersen, L. H., Fitzgibbons, S., and Pomorski, L. (2021). Responsible investing: The esg-efficient frontier. *Journal of financial economics*, 142(2):572–597.

Shefrin, H. and Statman, M. (2000). Behavioral portfolio theory. *Journal of financial and quantitative analysis*, 35(2):127–151.

United Nations (1992). Rio declaration on environment and development.

United Nations (2023). Principles for responsible investment.

United Nations Global Compact (2004). Who cares wins.

Van Duuren, E., Plantinga, A., and Scholtens, B. (2016). Esg integration and the investment management process: Fundamental investing reinvented. *Journal of business ethics*, 138(3):525–533.