



Ca' Foscari
University
of Venice

Master Degree programme in
Data Analytics for Business and Society

Final Thesis

Title

“The effect of environmental announcements on price
performance in the European
energy sector”

Supervisor

Prof. Diana Barro

Graduand

Fahimeh Shahbazi

Matriculation Number 898917

Academic Year

May 2025

Abstract:

This thesis investigates the impact of 450 environmental news events on the stock price behavior of 123 European companies. Using a diverse modeling framework Logistic Regression, Random Forest, XGBoost, and Quantile Regression we explore both predictive accuracy and interpretive insight.

XGBoost with cross-validation achieved the best predictive performance (AUC = 0.5723), making it the selected classification model. However, Quantile Regression revealed deeper economic patterns: while positive sentiment increased returns in the lower quantile ($\tau = 0.25$), it had the opposite effect in the upper quantile ($\tau = 0.75$), suggesting non-linear investor behavior.

These results emphasize the value of combining machine learning with the models to understand how markets react to ESG-related disclosures. This approach benefits analysts, investors, and policymakers aiming to decode market sentiment around sustainability.

Acknowledgement:

I would like to express my sincere gratitude to my supervisor, Professor Diana Barro, for their invaluable guidance, continuous support, and constructive feedback throughout the process of writing this thesis. Their insights greatly enriched the quality of my research.

I also wish to thank the faculty members and staff of Data Analytics for Business and Society, as well as Ca' Foscari University of Venice, for providing the academic environment and resources necessary to complete this work.

A heartfelt thanks goes to my family for their unwavering encouragement and patience during this journey. Their belief in me was a constant source of motivation. Finally, I am grateful to my friends and classmates who shared this academic path with me, offering both intellectual discussions and personal support along the way.

Table of contents

Table of Contents

CHAPTER 1:	9
INTRODUCTION	9
1.1 BACKGROUND AND CONTEXT	10
1.2 MOTIVATION FOR THE STUDY.....	10
1.3 RESEARCH PROBLEM	10
1.4 RESEARCH OBJECTIVES AND QUESTIONS.....	11
1.5 METHODOLOGICAL OVERVIEW	11
1.6 STRUCTURE OF THE THESIS.....	12
CHAPTER 2:	13
LITERATURE REVIE	13
2.1 INTRODUCTION	14
2.2. KEY CONCEPTS AND THEORETICAL FOUNDATIONS.....	14
2.2.1 <i>Environmental Announcements</i>	15
2.2.2 <i>Stock Price Performance</i>	15
2.2.3 <i>Market Reaction and Efficiency Hypotheses:</i>	15
2.2.4 <i>Recent Advances in Sentiment Analysis for Financial Applications:</i>	16
2.3. PREVIOUS EMPIRICAL STUDIES	17
2.3.1 <i>Asymmetric Market Reactions to Environmental News:</i>	17
2.3.2 <i>The Shift of Focus from What is Said to How it is Said</i>	18
2.3.3 <i>Industry, Region, and Firm-Level Heterogeneity:</i>	18
2.3.4 <i>ESG Ratings, Greenwashing, and Institutional Investors:</i>	19
2.3.5 <i>Methodological Trends and Evolving Research Directions:</i>	19
2.3.6 <i>Case Studies from European Financial Markets:</i>	20
2.4 RESEARCH GAP AND POSITIONING OF THE PRESENT STUDY.....	20
2.4.1 <i>Research Gaps in ESG Studies</i>	20
2.4.2 <i>Positioning of the Present Study</i>	21
CHAPTER 3:	22
METHODOLOGY	22
3.1 RESEARCH APPROACH	23
3.2 DATA COLLECTION	23
3.2.1 <i>Data Sources</i>	23
3.2.2 <i>Variable Construction</i>	28
3.3 EVENT STUDY FRAMEWORK.....	32
3.3.1 <i>Estimation Window</i>	32
3.3.2 <i>Event Window</i>	33
3.4 ECONOMETRIC AND MACHINE LEARNING TECHNIQUES.....	33
3.4.1 <i>Linear Regression</i>	33
3.4.2 <i>Robust Regression</i>	34
3.4.3 <i>Quantile Regression</i>	34

3.4.4 Logistic Regression	35
3.4.5 Random Forest.....	35
3.4.6 XGBoost	36
3.4.7 Rationale for Method Selection.....	36
3.5 MODEL EVALUATION AND VALIDATION:.....	36
3.6 MODEL VALIDATION: CROSS-VALIDATION STRATEGY	38
3.7 METHODOLOGICAL LIMITATIONS	39
CHAPTER 4:.....	40
DATA AND ANALYS	40
4.1 DESCRIPTIVE STATISTICS.....	41
4.1.1 Summary of Key Variables.....	41
4.1.2.2 Comparing Market Reaction Between Oil & Gas and Renewable Energy Firms.....	44
4.1.6 Statistical significant test	48
4.2 REGRESSION-BASED ANALYSIS	48
4.2.1 Linear and Robust Regression Results:	49
4.2.2 Quantile Regression ($\tau = 0.25$ and $\tau = 0.75$).....	51
4.2.3 Model Interpretation	52
4.2.4 Interpretation of Regression Coefficients and Key Predictors	54
4.3 CLASSIFICATION MODELS EVALUATION	56
4.3.1 Logistic Regression Performance.....	56
4.3.2 Random Forest and XGBoost Results	57
4.3.3 Model Performance Evaluation: Regression and Classification Approaches	59
4.3.4 ROC Curves and AUC Comparison	61
4.4 COMPARISON OF MODEL PERFORMANCE.....	62
4.4.1 MAE / RMSE for Regression Models.....	62
4.4.2 AUC and Accuracy for Classification Models	64
4.4.3 Summary Table of Evaluation.....	65
CHAPTER 5:.....	69
DISCUSSION.....	69
5.1 INTERPRETATION OF RESULTS	70
5.1.1 CAR Responses to News.....	70
5.2 COMPARISON WITH PRIOR STUDIES	71
5.2.1 Alignment with ESG Event Studies.....	71
5.2.2 Sentiment and Market Asymmetry	71
5.2.3 Firm Type vs. Sectoral Classification.....	71
5.2.4 Methodological Advancement and Predictive Modeling	72
5.2.5 Dataset Size and Analytical Resolution	73
5.3 POLICY AND PRACTICAL IMPLICATIONS.....	73
5.4 THEORETICAL INTERPRETATION	75
5.4.1 Efficient Market Hypothesis (EMH)	75
5.4.2 Behavioral Finance and Information Framing.....	75
5.4.3 Signaling Theory and ESG Disclosures	76
5.4.4 Stakeholder Theory and Market Expectations	76
5.4.5 Summary.....	76
5.5 LIMITATIONS AND DIRECTIONS FOR FUTURE RESEARCH:.....	77
CHAPTER 6:.....	78
CONCLUSION AND RECOMMENDATIO	78
6.1 SUMMARY OF FINDINGS.....	79

6.2 CONTRIBUTIONS OF THE STUDY	80
6.2.1 <i>Academic Contributions</i>	80
6.2.2 <i>Practical Contributions</i>	80
6.3 POLICY AND MANAGERIAL RECOMMENDATIONS	81
6.3.1 <i>ESG Disclosure Guidelines</i>	81
6.3.2 <i>Corporate Communication Strategy</i>	82
6.4 SUGGESTIONS FOR FUTURE RESEARCH.....	82
REFERENCES	83
APPENDICES	88

List of Figures:

Figure 4.1: Proportional Distribution of News Topics in the Dataset:

Figure 4.2: Comparative CAR Patterns Across News Themes

Figure 4.3: Statistical Comparison of Abnormal Returns (CAR) Between Oil & Gas and Renewable Companies Using Boxplot in RStudio

Figure 4.4: Distribution of Sentiment Scores Across News Categories

Figure 4.5: Distribution of Transformed Cumulative Abnormal Returns (CAR)

Figure 4.6: Accuracy and Prediction Metrics for Random Forest vs. XGBoost Models

Figure 4.7: Comparison of Regression Model Performance

Figure 4.8: Comparison of Classification Model Performance

List of Tables:

Table 3.1: Major Sources of Environmental News (Top 15)

Table 3.2: Environmental News Categorization

Table 3.3: Overview of Variables: Types and Descriptions

Table 3.4: Traditional Statistical Models with Associated Metrics

Table 3.5: Classification Model with Corresponding Evaluation Metrics

Table 4.1: Descriptive Statistics of Numerical Variables

Table 4.2: Comparison of Cumulative Abnormal Returns (CAR) Between Oil & Gas and Renewable Energy Companies

Table 4.3: Comparison Between OLS and Robust Regression Results

Table 4.4: Statistically Significant Predictors Identified Across Multiple Regression Models

Table 4.5: Performance Metrics: Full Dataset vs. Cross-Validation

Table 4.6: Comparative Performance Metrics on Full Dataset vs. Cross-Validation

Table 4.7: Comparative Performance Metrics on Full Dataset vs. Cross-Validation

Table 4.8: Evaluation Metrics: Full Dataset vs. Cross-Validation

Table 4.9: Comparison of Performance Metrics Across Regression Models

Table 4.10: Comparative Evaluation of Classification Model Performance

Table 5.1: Key ESG Communication Topics and Their Strategic Implications for Stakeholders

Chapter 1:
Introduction

1.1 Background and Context

Environmental, Social, and Governance (ESG) issues have gained increasing attention from investors, regulators, and corporations over the past decade. In particular, environmental disclosures such as announcements related to emissions, renewable energy adoption, or regulatory changes are now recognized as important non-financial signals that may affect firm valuation and investor sentiment.

As ESG principles become more integrated into capital markets, understanding how environmental information influences stock price performance has become a key concern in financial research. While much attention has been given to ESG scores and sustainability rankings, less is known about the short-term market reactions to individual environmental news events, particularly in the European energy sector.

1.2 Motivation for the Study

Europe has taken a leading role in the global energy transition and climate policy, making its capital markets a compelling context for studying ESG dynamics. However, most empirical studies to date have focused on annual ESG ratings or U.S.-based firms, leaving a gap in high-frequency, event-level analysis in European markets.

Furthermore, the rise of alternative data and textual analysis tools such as sentiment scoring and machine learning provides new opportunities to explore how markets process qualitative ESG information. This thesis aims to investigate how environmental news affects stock price reactions across firms with different characteristics (such as sector and market capitalization), and to evaluate the role of sentiment and news type in shaping these reactions

1.3 Research Problem

Although ESG factors are widely discussed in sustainability literature, the actual mechanisms through which environmental announcements influence short-term market behavior remain underexplored. It is unclear how the tone of such news, the type of company involved, or the nature of the event impacts investor reaction and whether machine learning models can effectively predict such responses.

1.4 Research Objectives and Questions

The objectives of this thesis are:

- To quantify the short-term stock market reaction to environmental news events in the European energy sector;
- To evaluate the role of sentiment, firm characteristics (such as company type and market capitalization), and news type in shaping abnormal returns;
- To investigate how analytical models can provide insights into ESG-related stock price movements in Oil & Gas and Renewable Energy firms.

The research addresses the following questions:

1. Do environmental announcements lead to statistically significant abnormal returns in European energy firms?
2. How do factors such as sentiment tone, company type, and market capitalization influence market reactions?
3. What do analytical models reveal about the drivers of ESG-related stock price movements within the European energy sector

1.5 Methodological Overview

This thesis uses a hand-collected dataset of 450 environmental news events linked to 123 European energy companies, including 65 Oil & Gas firms and 58 Renewable Energy firms. To measure market reactions, an event study framework is applied, calculating Cumulative Abnormal Returns (CARs) within a $[-3, +3]$ trading-day window around each news release, based on a 30-day estimation period.

Each event is also given a sentiment score from -5 (very negative) to $+5$ (very positive), capturing the tone of the announcement. Along with firm characteristics such as company type and market capitalization, these scores are used to explain differences in market response.

For the analysis, both traditional econometric methods (Linear and Quantile Regression) and machine learning models (Logistic Regression, Random Forest, and XGBoost) are applied. This combination makes it possible to both interpret how environmental news affects stock prices and test how well different models can predict market reactions.

1.6 Structure of the Thesis

The thesis is organized as follows:

- Chapter 2 presents a review of the relevant literature on ESG disclosures and market reaction.
- Chapter 3 describes the data sources, event study design, and machine learning methods.
- Chapter 4 presents the empirical results, including descriptive statistics, regression outcomes, and classification performance.
- Chapter 5 discusses the implications of the findings in relation to prior literature and practical applications.
- Chapter 6 concludes the study and offers recommendations for policy, management, and future research.

Chapter 2:
Literature Review

2.1 Introduction

Environmental issues have become increasingly central to European financial markets because sector-specific exposures to climate and regulatory risks are directly reflected in stock prices. For instance, energy and utilities firms often experience negative abnormal returns following announcements of stricter carbon regulation, whereas technology and service sectors may benefit from investor expectations of a low-carbon transition

Environmental announcements ranging from voluntary sustainability disclosures and regulatory updates to media reports shape how investors perceive corporate risk and opportunity, and these perceptions are quickly reflected in stock prices. For financial decision-makers, assessing the short-term price effects of such news is essential for managing portfolio risk and identifying sector-specific investment opportunities.

This chapter reviews studies on how environmental information affects financial markets. It is divided into three parts:

- (1) the basic theories of market efficiency and how markets process information;
- (2) research on how environmental news changes stock prices;
- (3) new tools such as text analysis and machine learning. The chapter ends by showing the research gaps, especially the lack of short-term studies on environmental news in European stock markets. This gap is the main focus of the thesis

2.2. Key Concepts and Theoretical Foundations

This section outlines the key concepts and theoretical frameworks that underpin this research. It begins by defining environmental announcements and explaining their relevance in financial contexts, continues with a discussion of stock price performance, and concludes with an overview of the Efficient Market Hypothesis (EMH) and behavioral finance theories that explain how financial markets react to information.

2.2.1 Environmental Announcements

Environmental announcements are public communications about a company's actions, risks, or results related to the environment. They can appear in many forms, such as voluntary ESG reports, regulatory documents, press releases, news coverage, or assessments by NGOs and rating agencies. Previous research shows that the market's reaction depends not only on the information itself but also on how it is presented, for example through tone, timing, and the credibility of the source (Clark et al., 2015; Grewal et al., 2020).

These announcements can lead to either positive or negative effects. For instance, green investments or net-zero targets may increase investor confidence. On the other hand, environmental accidents or lawsuits often damage a firm's reputation. Studies confirm that negative environmental news usually has a stronger impact on stock prices than positive disclosures (Krüger, 2015; Capelle-Blancard & Petit, 2019)

2.2.2 Stock Price Performance

In financial markets, stock price performance shows how prices change when new information becomes available. Event studies usually focus on abnormal returns, which are the part of returns that cannot be explained by general market movements. The most common indicators are Cumulative Abnormal Returns (CARs) over certain event windows (MacKinlay, 1997), as well as changes in volatility and trading volume.

These measures are widely used, but they also have some limits. For example, Friede et al. (2015) explain that CARs can show the immediate reaction but may fail to capture longer-term adjustments. This creates a tension between studying short-term reactions and understanding the lasting impact of ESG information

2.2.3 Market Reaction and Efficiency Hypotheses:

One of the first and the most important theories in the field of finance is the Efficient Market Hypothesis (Fama, 1970). In its semi-strong variant, it claims that stock values must incorporate all available information in a matter of a few moments. Nevertheless, the complication introduced by the ESG-related information assumes that such information is qualitative, ambiguous, and open

to subjective interpretation. Previous research shows that such properties could result in a delayed reaction in the market, mispricing it, or even overreacting to such information (Ball, 1994; Sautner et al., 2020).

There is an added dimension offered by the field of behavioral finance. Under such a framework, cognitive biases such as limited attention or framing effects could explain the under-reaction of the investors (Barberis & Thaler, 2003). For instance, Görden et al. (2020) show that market participants are more punitive towards a firm for negative ESG-related news, while positive news results in a far lesser response. Such observations point towards a futile and extreme response in the market in relation to news flows.

Even the more traditional event study methodologies which are primarily concerned with abnormal returns show the same issues. Such approaches are built on the assumption of the delayed reaction and account for mispricing asymmetrically and tone, framing of information, and other qualitative factors in an event study.

2.2.4 Recent Advances in Sentiment Analysis for Financial Applications:

Recent research in financial sentiment analysis has developed a lot compared to the first approaches. At the beginning, most studies used dictionary-based methods like the Harvard IV-4 or the Loughran–McDonald lists. These methods were simple and easy to use, but they had clear limits. For example, they could not deal with sarcasm, words that have more than one meaning, or new financial terms that often appear in the markets (Loughran & McDonald, 2011; Tetlock, 2007).

Later, researchers started to use machine learning models such as Logistic Regression, Support Vector Machines (SVM), and Random Forests. These models were more flexible and often gave better results. However, they needed a lot of feature engineering, for example bag-of-words or TF-IDF, and they still had problems with long or complicated texts (Li, 2010; Kogan et al., 2009).

A big improvement came with deep learning, especially transformer-based models like BERT, RoBERTa, and GPT. These models use attention mechanisms that help them understand meaning more effectively (Devlin et al., 2018; Brown et al., 2020). FinBERT, a version of BERT trained on financial texts, showed very strong results (Araci, 2019). But in ESG studies it has not been used much, mainly because there are not enough labeled datasets.

Even with these advances, some problems remain. ESG texts are often vague or written in a cautious way. For example, the phrase “considering green investments” looks positive, but it also shows uncertainty. On the other hand, “tighter emissions standards” may sound neutral, but it can have a strong impact on some firms. Normal models usually cannot capture these details (Ghosal et al., 2023). Another challenge is language. Most advanced models are trained on English data, but many ESG reports in Europe are written in German, French, Spanish, or Italian, which makes translation difficult (Müller et al., 2021).

Finally, transformer models are accurate, but many people see them as a “black box.” It is not clear how they reach their results, and this is a real issue in finance where trust and transparency are very important (Doshi-Velez & Kim, 2017).

Overall, NLP has improved sentiment analysis in finance a lot, but for ESG there are still many open questions and more research is needed

2.3. Previous Empirical Studies

2.3.1 Asymmetric Market Reactions to Environmental News:

Research shows that markets respond more strongly to negative environmental news than to positive news. For instance, Krüger (2015) demonstrates that environmental accidents or violations typically cause drops in stock prices, but positive changes, like "green" investments, have a much weaker effect. Capelle-Blancard and Petit (2019) also confirm this imbalance and attribute it to the phenomenon of reputational risk and loss aversion.

Not all researchers agree on how strong this effect is, however. Fatemi et al. (2018) observe that the news' significance matters. Only events that have direct financial or regulatory impacts tend to provoke strong market reactions; softer, more symbolic disclosure is, in this context, likely to produce mixed or ambiguous outcomes. This suggests that market reactions are held not only to the positive or negative value of the news, but also to the relevance of the news.

2.3.2 The Shift of Focus from What is Said to How it is Said

Researchers have shown that *how* news is communicated can be as important as the actual information (Krüger, 2015). In behavioral finance, it is expected that the way news sounds will affect how investors react. For example, imagine two firms that report the same emissions numbers. One says “we are committed,” while the other says “we are compliant.” The first message will probably create a stronger market reaction, even though the data are the same.

The role of sentiment is also well studied. Sokolov et al. (2021) found that positive language in company reports, even if the news itself is not financially important, can lead to noticeable positive abnormal returns. However, this does not always work, especially in the oil and gas sector. If the message is too vague or unrealistic, it may be seen as greenwashing and ignored by investors (Capelle-Blancard & Petit, 2019).

Tools such as FinBERT have improved the detection of tone and sentiment (Araci, 2019), but challenges remain. One issue is that ESG reports are often published in different languages and cultural settings, which makes analysis harder (Müller et al., 2021). Another problem is strategic ambiguity. For example, when a company says “we are exploring sustainable options,” the sentence sounds positive but may weaken the intended effect because it lacks clear detail (Fatemi et al., 2018).

2.3.3 Industry, Region, and Firm-Level Heterogeneity:

Industry, Region, and Firm-Level Heterogeneity

The way ESG news affects markets is not the same everywhere. For example, research shows that ESG news has a stronger effect in emerging markets, where reliable information is often missing (Ioannou & Serafeim, 2015). By contrast, in European markets—where ESG reporting follows clearer standards—the reaction is usually weaker.

Differences between sectors are also important. Studies indicate that firms in industries with high environmental risks, such as energy or mining, face sharper market reactions because of stronger regulatory oversight (Aouadi & Marsat, 2018).

Firm-level characteristics also matter. Companies with a strong ESG reputation tend to lose less value when bad news appears, which suggests that reputational capital can soften the market impact of negative events (Sautner et al., 2020)

2.3.4 ESG Ratings, Greenwashing, and Institutional Investors:

Another group of studies looks at third-party ESG ratings instead of company announcements. Research shows that firms with high ESG scores are more resilient, especially during times of crisis (Lins et al., 2017). However, rating agencies often give very different results, which creates “aggregate confusion” and makes markets less efficient (Berg et al., 2022).

When ratings are not consistent, some companies focus on the positive ones and downplay the negative ones. This situation gives them a chance to look better than they really are, which is often linked to greenwashing. The lack of consistency also makes it difficult for investors to separate genuine ESG performance from disclosure that is only cosmetic.

Institutional investors, especially long-term funds, therefore ask for higher ESG quality and greater transparency (Dyck et al., 2019). At the same time, how markets react depends a lot on whether the underlying data is considered credible

2.3.5 Methodological Trends and Evolving Research Directions:

Research methods have changed from traditional event studies to more advanced approaches. Event studies are still useful for measuring short-term abnormal returns (MacKinlay, 1997). However, many scholars now use panel regressions, instrumental variables, and machine learning to study causality and non-linear effects (Christensen et al., 2022).

One important trend is to study ESG dimensions separately. Broadstock et al. (2021) found that environmental factors are linked to short-term volatility, while social and governance factors have stronger effects on long-term valuation. This shows a shift from looking at ESG as one overall score to seeing it as a multi-dimensional concept.

2.3.6 Case Studies from European Financial Markets:

Compared with U.S.-focused studies, European evidence is still limited but growing. Busch et al. (2022) find positive CARs following renewable energy investment announcements in Germany, especially when voluntarily disclosed. Lefebvre and Roussel (2021) show that pollution-related news in France leads to significant losses, but firms with strong ESG ratings are less affected. Haapanen et al. (2020) report that green bond announcements in Nordic countries generate positive stock reactions, particularly for first-time issuers.

Taken together, these studies highlight the **contextual nature of ESG effects**: responses depend on regulation, reputation, sector, and framing. Yet the evidence remains fragmented, especially for pan-European datasets.

2.4 Research Gap and Positioning of the Present Study

2.4.1 Research Gaps in ESG Studies

Even though ESG research has grown quickly, some important gaps remain. Many studies still use aggregated ESG ratings or yearly CSR reports. These give a broad picture but do not show the immediate effects of real-time environmental news. As a result, they underestimate short-term investor reactions to specific events.

Another gap is geographical bias. Most of the evidence comes from the United States or Asia. European markets, even with their strong ESG rules and high investor awareness, are less studied. This makes it harder to understand how environmental disclosures are priced in a region where sustainability reporting is already mandatory.

A further limitation is how textual information is handled. Many studies simplify complex disclosures into keyword counts or sentiment dictionaries. These methods miss important elements such as tone, framing, and ambiguity, which strongly influence investor reactions. Only a few studies include sentiment measures directly in models of price performance.

Finally, while machine learning is becoming common in finance, it is still rarely used in ESG event studies. Traditional regressions often assume simple, linear effects, but investor behavior is

shaped by non-linear patterns, differences between firms, and context. Few studies explore whether data-driven models can add more insight or predictive power in this area

2.4.2 Positioning of the Present Study

This thesis addresses these gaps in several ways. First, it builds a new dataset of 450 environmental news events from 123 European energy firms. Unlike aggregated ESG scores, this dataset captures event-level information and allows direct study of short-term market reactions.

Second, it adds sentiment analysis to financial models. Each event gets a sentiment score, so the study can test if tone and framing help explain abnormal returns. This connects qualitative news with quantitative results.

Third, the thesis uses quantile regression to look beyond average effects and study differences across the return distribution. This shows whether news has stronger impacts on firms with extreme gains or losses.

Fourth, it applies machine learning models such as Random Forest and XGBoost. These test the predictive value of non-linear methods. By comparing explanatory and predictive power, the study shows both the strengths and the limits of data-driven approaches.

Finally, the focus is on the European energy sector, comparing Oil & Gas with Renewable Energy firms. This industry is important because of its central role in the climate transition and its exposure to environmental risks.

In short, this study connects financial event studies with modern data science. It combines theory, empirical evidence, and advanced methods to provide new insight into how environmental news influences stock prices in European markets.

Chapter 3:

Methodology

3.1 Research Approach

This research takes a quantitative and explanatory approach to study how environmental news affects the stock prices of European energy companies. The analysis combines the event study method with both econometric models and machine learning techniques.

The study has two main goals. The first is to see how stock markets react in the short term to environmental news, measured through abnormal returns. The second is to compare different models and find out how well they can explain or predict these reactions.

The research is carried out in two steps. In the first step, the event study method is used to calculate Cumulative Abnormal Returns CAR around the dates of news announcements. This makes it possible to see whether such news causes unusual movements in stock prices. In the second step, regression and classification models are applied to check how factors such as sentiment score, type of news, company size, and sector influence the CAR values.

Together, these two steps provide a clearer picture of how markets respond to environmental information and which factors play the most important role

3.2 Data Collection

3.2.1 Data Sources

This study uses two main types of data to analyze how environmental news affects stock prices in European markets.

Company and Market Data

Data on companies, such as daily stock prices, industry type, and market size, was collected from the Bloomberg Terminal. To keep the data consistent, the Bloomberg Industry Classification Standard (BICS) was used. The dataset includes 123 listed companies in Europe: 65 in Oil & Gas and 58 in Renewable Energy

Environmental News Announcements

The news outlets in Table 3.1 were chosen to give a fair picture of environmental events that matter to investors. First, EU websites (like ec.europa.eu, eea.europa.eu, energy.ec.europa.eu, commission.europa.eu) were added because they publish official rules and policy news. These updates can change costs for firms or create new chances, such as subsidies for clean energy. Second, NGOs and specialist sites (such as caneurope.org, cleanenergywire.org, stateofgreen.com) report climate and energy transition issues in detail. They often show early debates or reactions from different groups, which may give investors useful signals. Third, international media (e.g., euronews.com, reuters.com, theguardian.com, euractiv.com) are widely followed. Their news is read by many traders and spreads fast through platforms, so it can move prices in the short run. Finally, regional and industry outlets (like pv-magazine.com, ansa.it, total-croatia-news.com, citizen-led-renovation.ec.europa.eu) add stories about local projects and sector news. Even if they are not global, such reports can be very important for specific firms. Taken together, these sources mix official policy, NGO views, mainstream news, and local insights. All of them are open to the public and followed by people in the market. This makes it more likely that the events collected are both visible and important for stock prices in the European energy sector.

Table 3.1 : Major Sources of Environmental News (Top 15)

Rank	Source Domain	Number of Articles
1	euronews.com	68
2	ec.europa.eu	63
3	caneurope.org	40
4	eea.europa.eu	38
5	energy.ec.europa.eu	18
6	theguardian.com	8
7	reuters.com	7
8	cleanenergywire.org	7
9	euractiv.com	6

Rank	Source Domain	Number of Articles
10	stateofgreen.com	5
11	ansa.it	5
12	commission.europa.eu	4
13	citizen-led-renovation.ec.europa.eu	4
14	total-croatia-news.com	4
15	pv-magazine.com	4

Selection of Sustainability-Related Keywords

The keywords were selected to match the main goal of this thesis, which is to study how stock prices in the European energy sector react to environmental announcements. The study focuses on news that highlights the shift from fossil fuels to renewable energy and from business models that harm the environment to more sustainable practices, in order to investigate how such announcements influence market behaviour. Terms such as “green transition,” “carbon neutrality,” “clean energy,” “net-zero goals,” “pollution control,” and “renewable energy policy” directly reflect this transition toward clean energy and sustainable business. Using these keywords ensures that the collected articles are consistent with the research objective and capture the events most relevant to investor reactions.

The choice of these keywords is also consistent with the academic literature. Prior studies have shown that environmental news and sustainability-related terms strongly affect financial markets. For example, Krüger (2015) and Capelle-Blancard & Petit (2019) highlight the role of environmental events and pollution risks in driving abnormal returns. Fatemi et al. (2018) and Broadstock et al. (2021) emphasize how commitments to clean energy and carbon neutrality can influence investor expectations and firm value. More recent work, such as Christensen et al. (2022) and Ghosal et al. (2023), also shows that climate-related announcements and sustainability goals shape market sentiment

Each news item was categorized into one of the seven predefined groups, as shown in Table 3.2. Each news article was assigned a sentiment score on a scale from -5 (very negative) to $+5$ (very positive). This scoring reflects the tone and expected market perception of the announcement. For instance, a highly negative article such as “8 reasons why the Energy Charter Treaty reform is urgently needed” (CAN Europe, 12 September 2021), which highlighted risks and regulatory failures, received a score of -4.25 . In contrast, strongly positive announcements, such as “EU Invests €5.4 Billion in Hydrogen Projects” (European Commission, 15 July 2022), emphasizing large-scale clean energy investment, were scored at $+4.9$.

Table 3.2: Environmental News Categorization

News Category	Number of Articles
Renewable Energy	142
Energy Transition	132
Pollution & Environmental Risks	69
Regulatory Changes	42
Sustainability & ESG Strategy	26
Market Impact / Financials	25
Uncategorized	14

Construction of the Dataset and Event Study Design

In the first step, daily data were collected from Bloomberg for companies operating in the European energy sector. The sample was divided into two main groups: oil & gas and renewable energy firms.

For the oil & gas group, the initial list included 191 firms. However, very small firms with negligible market capitalization (below 1,464,286.27 EUR) were excluded. Following standard

practice in event studies, this filtering ensures that the sample is representative of the market and that the results are not driven by illiquid or marginal firms with very limited trading activity (e.g., Campbell, Lo, & MacKinlay, 1997; Kothari & Warner, 2007). After this adjustment, the final number of oil & gas firms was 70. The largest company in this group was TUPRAS – Turkiye Petrol Rafine, with a market capitalization of 280,348,759,509 EUR.

For the renewable energy group, the initial list included 174 firms. Companies with a market capitalization below 70,000 EUR were excluded, and the remaining firms were classified into small, medium, and large cap categories. This filtering again follows common practice in empirical finance to avoid distortions caused by micro-cap firms with thin trading volumes (Brown & Warner, 1985). After these adjustments, the number of renewable energy firms was reduced to 60. The largest company in this group was MARGUN Enerji Uretim Sanayi, with a market capitalization of 22,774,000,000 EUR.

In the next step, daily price data were added for all selected firms, giving a dataset of 130 companies and 159,778 rows of stock price observations.

In this study, each news event was applied to all companies based on its publication date. For every event, an event window of $[-3,+3]$ trading days was defined around the publication date, and abnormal returns (AR) and cumulative abnormal returns (CAR) were calculated for each firm. The choice of a 7-day event window is standard in the literature, as it captures both immediate and short-term market reactions while reducing the risk of including unrelated noise (see, e.g., MacKinlay, 1997; Corrado, 2011). If a wider window such as $[-5,+5]$ were applied, the effect of the event would likely be diluted because unrelated market movements would be incorporated, making the reaction to the news less visible (see Brown & Warner, 1985; Campbell et al., 1997).

The reduction in the final dataset size is due to timing constraints linked to event dates. A firm event pair was excluded when there were no valid observations in the event or estimation window for example, if the event date fell on a market holiday, if trading was halted, if the company was a recent IPO, or if the estimation window $[-30,-4]$ did not contain sufficient trading data. Therefore, the lower number of rows results from misalignment between event timing and trading calendars, not from a general lack of price data.

As a result, the final dataset contains approximately 273,061 rows, covering 123 companies (58 renewable and 65 oil & gas) and 13 columns. All data processing, construction of abnormal returns, and CAR calculations were carried out using Python, which ensured efficient handling of large-scale financial datasets and reproducibility of the results

3.2.2 Variable Construction

To conduct the analysis, several key variables were constructed from the raw data sources, each designed to capture a specific aspect of the relationship between environmental news and stock price behavior.

- **Cumulative Abnormal Return (CAR):**

Computation of Cumulative Abnormal Return (CAR)

The calculation of cumulative abnormal returns (CAR) in this study was carried out in four main steps:

1. Daily Return (DR)

Daily returns were calculated for each company from adjusted closing prices between 2020 and 2024. The daily return measures the percentage change in the stock price from one trading day to the next, and it provides the basic indicator of short-term stock performance. This measure is commonly used in financial event studies because it is simple and intuitive (Brown & Warner, 1985).

2. Expected Return (ER)

To estimate the “normal” return of a firm in the absence of an event, expected returns were calculated using the constant mean return model. Specifically, the expected return was defined as the rolling average of the past 30 trading days (excluding the current day)

$$ER_{i,t} = 1/30 \sum_{k=1}^{30} R_{i,t-k}$$

This approach is accepted in the event study literature, and prior research shows that mean-adjusted returns often yield similar results to more complex models, especially in short windows (Brown & Warner, 1985)

3. Abnormal Return (AR)

Abnormal returns were obtained as the difference between the observed daily return and the expected return

$$AR_{i,t} = R_{i,t} - ER_{it}$$

A positive abnormal return indicates that the stock performed better than expected, while a negative abnormal return indicates underperformance. This measure isolates the firm-specific effect of the news (MacKinlay, 1997)

4. Cumulative Abnormal Return (CAR)

Finally, cumulative abnormal returns were calculated by summing the abnormal returns across the event window. In this study, the event window was defined as three trading days before and three days after the news date ($[-3,+3]$)

$$CAR_i(-3,+3) = \sum_{t=-3}^{+3} AR_{it}$$

CAR captures the total effect of a news event on stock prices, taking into account both immediate and delayed reactions. A short window such as $[-3,+3]$ is commonly used in the literature

- **Sentiment Score:**

Each news article is assigned a sentiment score ranging from -5 (strongly negative) to $+5$ (strongly positive). This variable captures the tone of the news and is used to explore whether more optimistic or pessimistic environmental signals generate stronger market reactions.

- **News Category:**

In order to analyze the market reaction to environmental news in a structured way, each article was assigned to one of seven categories: Energy transition, Market impact/Financials, Pollution & Environmental Risks, Regulatory changes, Renewable

energy, Sustainability & ESG, and Uncategorized. These categories represent key themes frequently discussed in the literature on climate finance and environmental economics (Krueger, Sautner, & Starks, 2020; Flammer, 2021).

The motivation for this classification is twofold. First, it makes it possible to separate different types of environmental information, since investors may react differently to financial risks (such as pollution fines) compared to long-term opportunities (such as renewable energy projects). Second, it provides a framework to test whether specific themes such as regulatory changes or ESG announcements generate stronger abnormal returns than more general or mixed news.

The classification process combined manual coding and automated text processing. A manual step ensured that context and domain-specific meaning were captured, while automated methods (keyword filters and text matching) supported consistency and scalability across the full dataset

- **Company Type:**

Firms are grouped into two categories based on their sector: Oil & Gas or Renewable Energy. This classification is used to analyze heterogeneity in responses based on the company's business model and exposure to environmental risks.

- **Market Capitalization Category:**

To investigate whether firm size affects the sensitivity to environmental information, companies were classified into three groups based on their market capitalization: Small Cap, Mid Cap, and Large Cap. The thresholds for each group were defined separately for the oil & gas and renewable energy sectors to reflect the different size distributions of firms in each industry. Global index providers such as MSCI and Russell also define size categories relative to market distributions rather than fixed absolute values. Following this idea, the cut-offs in our study were chosen based on the distribution within our sample. This ensured that each group contained a sufficient number of firms and avoided the problem that, under global thresholds, almost all renewable firms would fall into the Small Cap category (Fama & French, 1992; MSCI, 2023; ETF.com, 2020).

Oil & Gas: Firms were classified as Small Cap if their market value was below €131 million, as Mid Cap if between €131 million and €2.18 billion, and as Large Cap if above

€2.18 billion. From these groups, 24 Large Cap, 23 Mid Cap, and 23 Small Cap firms were randomly selected.

Renewable Energy: Firms were classified as Small Cap if their market value was below €60 million, as Mid Cap if between €60 million and €350 million, and as Large Cap if above €350 million. From these groups, 30 firms were randomly selected for each size category.

Random sampling within each category ensured a balanced representation of firm sizes in the final dataset. The resulting list of selected companies was exported and used for the subsequent event study analysis

- **Transformed CAR**

In the first analysis, the distribution of the cumulative abnormal returns (CAR) had a very sharp peak around zero and long tails. This is typical in financial data, but it can cause problems for statistical models that assume normality. To reduce this problem, the CAR variable was changed with a signed square root transformation

$$CAR_{\text{Transformed}} = \sqrt{CAR} * \text{sign}(CAR)$$

This transformation helps in two ways: (1) it reduces the effect of extreme values by compressing very large observations, and (2) it keeps the original direction of abnormal returns by preserving the sign. Similar transformations are often recommended in event study research (MacKinlay, 1997) and in statistical methods for non-normal data (Box & Cox, 1964; Tukey, 1977). Since financial return distributions are known to be heavy-tailed (Cont, 2001), this approach makes the regression analysis more reliable and less affected by outliers

These variables serve as the foundation for both the event study and the predictive modeling stages of the analysis, enabling a multi-dimensional exploration of how environmental announcements interact with firm characteristics and investor sentiment

Table 3.3 Overview of Variables: Types and Descriptions

Variable	Type	Description
CAR	Continuous	Cumulative Abnormal Return over event window
Sentiment Score	Numerical	Tone of environmental news (-5 to +5)
News Category	Categorical	Type of news (7 groups)
Company Type	Categorical	Oil & Gas, Renewable Energy
Market Cap Category	Categorical	Size group (Small, Middle, Large)
Transformed CAR	Continuous	Transformed CAR over event window

3.3 Event Study Framework

The event study methodology is employed to quantify the stock market reaction to environmental news at the firm level. This approach allows for the measurement of abnormal returns the deviation of actual returns from expected performance following specific events, in this case, environmental announcements.

3.3.1 Estimation Window

To estimate expected returns, a market model is used over a 30-day estimation window preceding the event. This model relates the return of an individual stock to the return of a broad market index, capturing its normal behavior absent the event. The estimation window ensures that model parameters are not influenced by the announcement itself.

3.3.2 Event Window

The event window spans seven trading days, from three days before to three days after the announcement ($[-3, +3]$). This window is designed to capture any pre-event information leakage, immediate market reaction, and short-term adjustments.

- **Day 0** is the publication date of the news article.
- **Days -3 to -1** to consider the chance that some investors react before the news becomes public.
- **Days +1 to +3** capture delayed market adjustments.

3.4 Econometric and Machine Learning Techniques

This study employs a hybrid approach that combines traditional econometric methods with advanced machine learning algorithms to analyze the relationship between environmental announcements and stock price reactions. The purpose is to reflect both interpretability and predictive accuracy

3.4.1 Linear Regression

Linear regression is used as a baseline model to estimate the impact of sentiment scores and market capitalization, and company type on cumulative abnormal returns (CAR). The model assumes a linear relationship between independent variables and CAR:

$$CAR_i = \beta_0 + \sum_{j=1}^m X_{ij} \beta_j + \varepsilon_i$$

where

- $i=1,\dots,n$ denotes the index of observations.
- $j=1,\dots,m$ denotes the index of explanatory variables.
- X_{ij} is the value of explanatory variable j for observation i .
- β_0 is the intercept and β_j are the regression coefficients.
- ε_i is the error term.

This technique allows for straightforward interpretation of coefficients, providing insight into the marginal effects of predictors such as sentiment score, market cap category, and news type

3.4.2 Robust Regression

This model has a structure similar to ordinary linear regression, but the parameters $\beta_0, \beta_1, \dots, \beta_m$ are estimated using iteratively reweighted least squares in order to reduce sensitivity to outliers and non-normal errors.

Previous studies have shown that robust regression performs better than ordinary least squares in the presence of data irregularities. For example, Yu et al. (2014) compare different robust regression techniques and demonstrate that M-estimator-based methods are more stable against outliers

3.4.3 Quantile Regression

Quantile regression is implemented at the 25th and 75th percentiles ($\tau = 0.25$ and $\tau = 0.75$) to examine the heterogeneous effects of predictors across the CAR distribution. Unlike ordinary least squares (OLS), this method allows for estimation of effects at different points in the outcome distribution:

$$QCAR_{i(\tau)} = \beta_{0(\tau)} + \sum_{j=1}^m X_{ij} \beta_{j(\tau)} + \varepsilon_{i(\tau)}$$

Where :

- $QCAR_{i(\tau)}$ denotes the τ -quantile of the cumulative abnormal return for observation i .
- $\beta_{0(\tau)}$ is the quantile-specific intercept.
- $\beta_{j(\tau)}$ is the coefficient of explanatory variable X_{ij} at quantile τ .
- X_{ij} represents the value of explanatory variable j for observation i .
- $\varepsilon_{i(\tau)}$ is the error term at quantile τ .

This is particularly useful in identifying whether certain factors influence firms with negative vs. positive abnormal returns differently

3.4.4 Logistic Regression

For classification, logistic regression is applied to estimate the probability that a firm experiences a positive CAR (binary outcome). The model uses news sentiment and firm-level features to predict the log-odds of a favorable reaction. Logistic regression is widely used in financial studies because it is interpretable and often serves as a baseline for more complex classifiers (Hosmer et al., 2013; Kleinbaum & Klein, 2010).

3.4.5 Random Forest

Random Forest is a non-parametric ensemble method that constructs multiple decision trees and averages their predictions. It can model complex interactions and produces feature importance scores that identify the most influential predictors. Random Forest has been shown to perform well in financial classification tasks due to its robustness and ability to reduce overfitting (Breiman, 2001; Cutler et al., 2007)

3.4.6 XGBoost

Extreme Gradient Boosting (XGBoost) is an advanced tree-based algorithm that improves predictive performance through sequential boosting and strong regularization. It is efficient in handling large datasets and is less prone to overfitting compared to traditional methods. XGBoost has become one of the most popular algorithms in applied finance and economics due to its high accuracy and interpretability through feature importance metrics (Chen & Guestrin, 2016; Zhang et al., 2019)

3.4.7 Rationale for Method Selection

The combination of econometric models and machine learning in this study follows the two main goals of interpretability and predictive power. Econometric models, such as linear regression and quantile regression, allow clear interpretation of coefficients and hypothesis testing. However, they rely on strong assumptions, for example linearity or independence, which are often not valid in financial data (Gujarati & Porter, 2009).

On the other hand, machine learning models such as Random Forest and XGBoost are better at capturing non-linear patterns, dealing with high-dimensional data, and identifying complex relationships. These models are useful for classifying market reactions when traditional statistical methods do not perform well (Breiman, 2001; Chen & Guestrin, 2016).

By combining both approaches, this research benefits from the strengths of each: econometric models provide explanation, while machine learning focuses on prediction

3.5 Model Evaluation and Validation:

To assess the predictive performance and generalizability of the models used in this study, a consistent evaluation and validation framework was applied. All classification models were validated using 5-Fold Cross-Validation, a standard procedure recommended in the literature to balance bias and variance in model evaluation (Kohavi, 1995; Hastie et al., 2009)

Evaluation Metrics for Traditional Models:

For regression models (Linear, Quantile, and Robust), the following metrics were used:

- R-squared (R^2): The proportion of variance in CAR explained by the model.
- Root Mean Squared Error (RMSE): The square root of the average squared prediction error.
- Mean Absolute Error (MAE): The average absolute difference between actual and predicted CAR values.

Table 3.4 Traditional Statistical Models with Associated Metrics

Model	R^2	RMSE	MAE
Linear Regression	0.007	0.0698	0.0514
Quantile Regression ($\tau=0.25$)	—	0.0653	0.0503
Quantile Regression ($\tau=0.75$)	—	0.0658	0.0506
Robust Regression	0.007	0.0700	0.0515

Evaluation Metrics for Classification Models:

For classification models (Logistic Regression, Random Forest, XGBoost), the following performance metrics were employed:

- Accuracy: The proportion of correctly classified cases.
- F1 Score: The harmonic mean of precision and recall, suitable for imbalanced classes.
- Area Under the ROC Curve (AUC): A measure of the model's ability to distinguish between positive and negative cases.

All metrics below are averaged across the five folds of cross-validation:

Table 3.5 Classification Model with Corresponding Evaluation Metrics

Model	Accuracy	F1 Score	AUC
Logistic Regression	0.5400	0.6100	0.5600
Random Forest	0.5438	0.6094	0.5556
XGBoost	0.5483	0.5902	0.5742

Interpretation of Model Evaluation Results (Tables 3.4 and 3.5)

The results reveal that regression models, including linear and robust regression, showed very low explanatory power ($R^2 \approx 0.007$), indicating that a linear relationship does not sufficiently capture the impact of environmental news on CAR. Quantile regression, however, yielded lower RMSE and MAE values, suggesting better handling of skewed or heteroskedastic data.

Among classification models, XGBoost outperformed Random Forest and Logistic Regression in terms of AUC (0.5742), indicating a stronger ability to differentiate between positive and non-positive CAR events. Although Logistic Regression and Random Forest achieved slightly higher F1 Scores, XGBoost's improved AUC and Accuracy suggest more consistent performance across classes.

3.6 Model Validation: Cross-Validation Strategy

To ensure robust generalization, all classification models were evaluated using 5-Fold Cross-Validation, a standard method for reducing variance in model evaluation (Kohavi, 1995). In each fold, the dataset was split into 80% training and 20% test subsets. Performance metrics such as accuracy, F1-score, and AUC were averaged across the five iterations.

This approach reduces the sensitivity of evaluation to any particular data split and provides a more stable estimate of out-of-sample predictive power.

For regression models (e.g., linear, quantile), full-sample fitting was adopted, as the primary goal was explanatory insight rather than forecasting accuracy. However, residuals were thoroughly examined for normality, heteroscedasticity, and influential observations to validate model reliability (Wooldridge, 2016)

3.7 Methodological Limitations

There are several methodological limitations that need to be taken into account in this study. Building the dataset required some manual work, especially in scoring sentiment and translating non-English news, which may have introduced a degree of subjectivity and small errors. The event study was carried out using a fixed $[-3, +3]$ window together with a simple market model for expected returns. While this approach is common, it may not fully reflect delayed market reactions or the influence of several factors at the same time. In addition, regression and classification models rely on assumptions such as linear relationships, stable error variance, and independence between observations. Because financial data are often volatile, non-linear, and time-dependent, these assumptions are not always satisfied, and the results should therefore be interpreted with care. Finally, the focus on European energy companies gives depth to the analysis but reduces the extent to which the findings can be generalized to other industries and regions.

Chapter 4:

Data and Analysis

4.1 Descriptive Statistics

4.1.1 Summary of Key Variables

Main Quantitative Variables

The average CAR is close to zero. This result is expected in an event study because positive and negative market reactions to different news items usually offset each other. In other words, while some firms show strong gains and others experience losses, the overall effect across all firms and events balances out. However, the standard deviation indicates that these reactions differ widely. A higher standard deviation means that some companies reacted very strongly, either positively or negatively, while others had little or no reaction.

The sentiment score ranges from strongly negative to strongly positive. The average value of about 1.6 suggests that the dataset contains slightly more optimistic than pessimistic news. This distribution is important because it shows that the market is not exposed to neutral news only, but to a wide range of tones that can drive different investor responses.

Market capitalization also displays a large variation. The dataset includes both small renewable firms with very low market value and large oil & gas companies with market values above €200 billion. This spread reflects the diversity of the sample and allows the analysis to test whether firm size plays a role in shaping the sensitivity of stock prices to environmental new

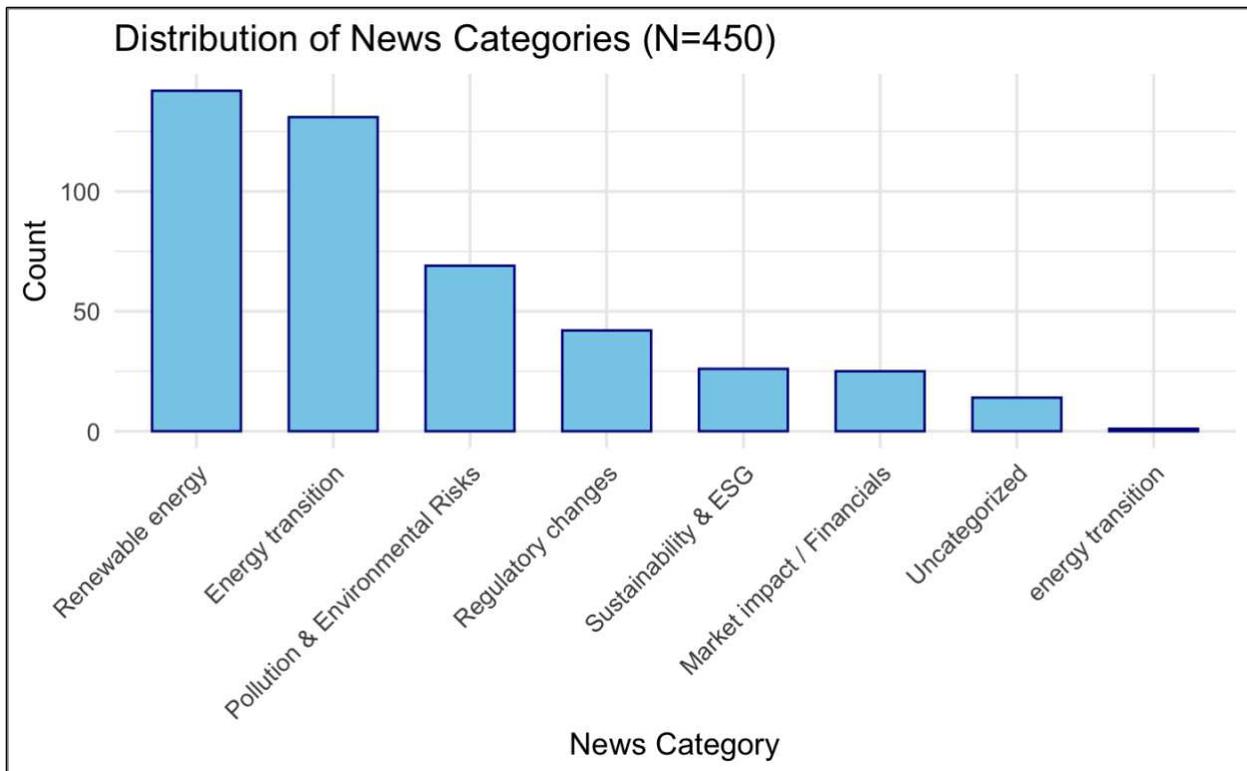
Table 4.1 Descriptive Statistics of Numerical Variables

Variable	Mean	Std Dev	Min	Q1	Median	Q3	Max
CAR (Before Outlier)	0.00130	1.53042	-719.03329	-0.02308	0.00000	0.02672	71.79736
CAR (After Outlier)	0.00107	0.10628	-0.99894	-0.02290	0.00000	0.02656	0.99984
Sentiment Score	1.5910	1.807653	-4.2530	0.7914	1.5241	2.9295	4.9000
Market Cap (EUR)	2.865e+08	32209912661	7.00e+04	4.325e+07	2.865e+08	8.497e+09	2.800e+11

4.1.1 Description of The Distribution of News Categories

Figure 4.1 shows the distribution of 450 environmental news items between 2020 and 2024. Most of the news belongs to the categories of Renewable Energy and Energy Transition, reflecting strong attention to clean energy and decarbonization. Pollution & Risks comes next, covering issues such as environmental damage and health impacts. Other categories, such as Regulatory Changes and Sustainability & ESG, appear less often, while Market/Financials and Uncategorized are the smallest groups. This imbalance indicates that the dataset is more energy-focused, with fewer reports on governance or financial topics. These differences may also explain why some categories trigger stronger market reactions. The classification procedure was described in detail in Chapter 3

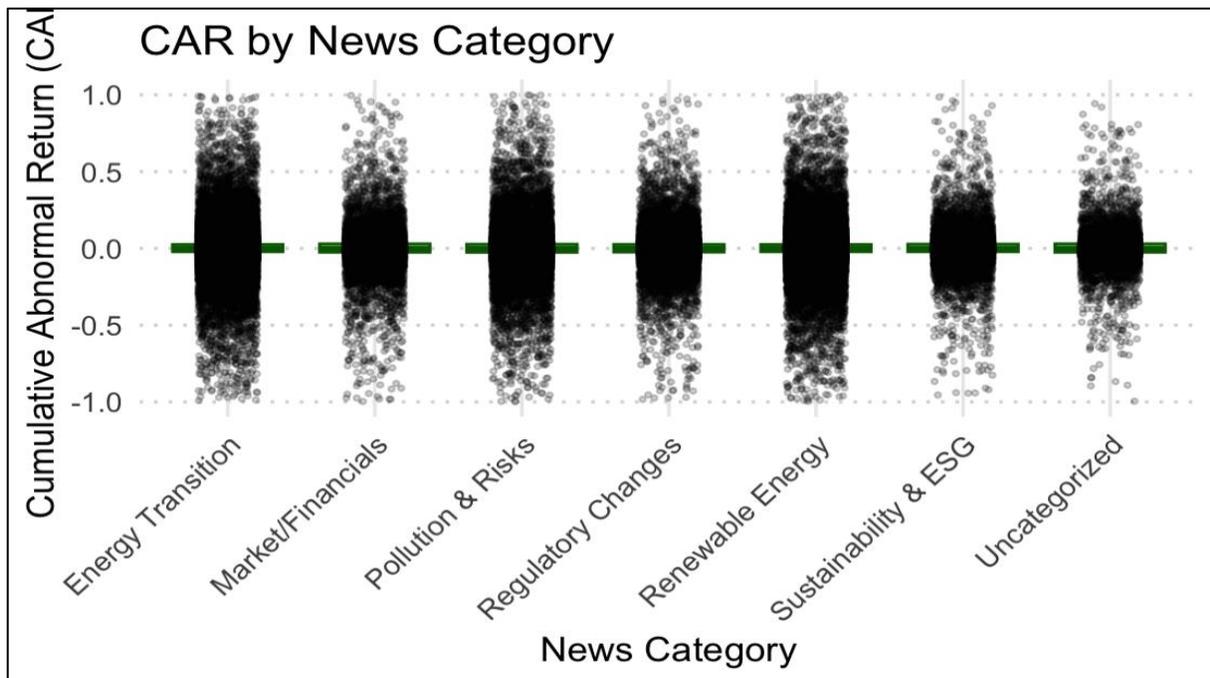
Figure 4.1 Proportional Distribution of News Topics :



4.1.2 Distribution of Cumulative Abnormal Returns (CAR) Across News Categories

Figure 4.2 shows how stock prices react to different types of environmental news by presenting the distribution of Cumulative Abnormal Returns (CAR) for each category. The green dashed line marks zero, meaning no market reaction. Most dots are close to this line, suggesting that in many cases, news does not lead to strong abnormal returns. Some categories, such as Pollution & Risks and Regulatory Changes, display a wider spread of values. This means that events in these areas sometimes cause sharp positive or negative market movements. Such reactions may be linked to unexpected regulations, fines, or negative disclosures that raise investor concerns. By contrast, categories like Renewable Energy and Sustainability & ESG show tighter clusters around zero. Their effect is smaller and more stable, possibly because these topics represent long-term changes already considered by investors. The Market/Financials group looks balanced and less emotional, while the Uncategorized group is more scattered due to mixed content. Overall, the figure suggests that not all environmental news is treated equally. While average reactions are close to zero, the degree of volatility differs between categories. This highlights the importance of studying each category separately, since combining all ESG news into one group would hide important differences.

Figure 4.2 Comparative CAR Patterns Across News Themes



Managing Outliers

To ensure clarity and consistency in visual interpretation, the Y-axis of the plot was limited to the range of -1 to $+1$. This decision was based on the fact that over 99% of cumulative abnormal return (CAR) values fall within this interval. By constraining the axis in this way, the figure effectively reduces visual distortion that could arise from a small number of extreme outliers. This approach enables a clearer, more meaningful comparison across different news categories, ensuring that patterns in the central distribution are not visually overshadowed by rare, extreme deviations. Such normalization is especially important when interpreting dense scatterplots, where subtle differences in spread and skewness among groups are of analytical interest.

To avoid visual distortion caused by extreme values, the Y-axis was restricted to the range of -1 to $+1$, which covers more than 99% of the observations. This allows for clearer comparison across news categories.

4.1.2.2 Comparing Market Reaction Between Oil & Gas and Renewable Energy Firms

Beyond the content of environmental news, one critical dimension influencing market reaction is the type of firm receiving the news. Within the energy sector, there are fundamental structural and reputational differences between Oil & Gas companies and those operating in the Renewable Energy domain. These differences may shape how investors interpret ESG-related announcements and, subsequently, how the market reacts.

This section examines whether there is a statistically significant difference in Cumulative Abnormal Returns (CAR) between these two firm types following environmental news.

Table 4.2: Comparison of Cumulative Abnormal Returns (CAR) Between Oil & Gas and Renewable Energy Companies

Company Type	Mean CAR	Standard Deviation	Observations
Oil and Gas	0.00065	0.07086	26,505
Renewable	0.00081	0.08014	23,397

In this part of the analysis, the dataset was divided into two groups based on Company Type: Oil & Gas and Renewable Energy. For each group, the cumulative abnormal returns (CAR) were calculated, and descriptive statistics such as mean, standard deviation (SD), and number of observations were reported. To formally test whether the mean CAR differs between the two groups, a Welch Two-Sample t-test was applied, as it does not assume equal variances. In addition, a boxplot was used to visualize the distribution of CAR values across the two groups.

The table presents the results of this comparison. The mean CAR for Oil & Gas firms is 0.00065, while for Renewable firms it is slightly higher at 0.00081. Both averages are very small and close to zero, showing that, on average, the market reacts only weakly and positively to environmental news. The difference between the two means is minimal, about 0.00016.

The standard deviations, however, are much larger (0.07086 for Oil & Gas and 0.08014 for Renewable). This indicates that individual company reactions vary a lot. In practice, some firms react positively to ESG news, while others react negatively, leading to a wide spread of values.

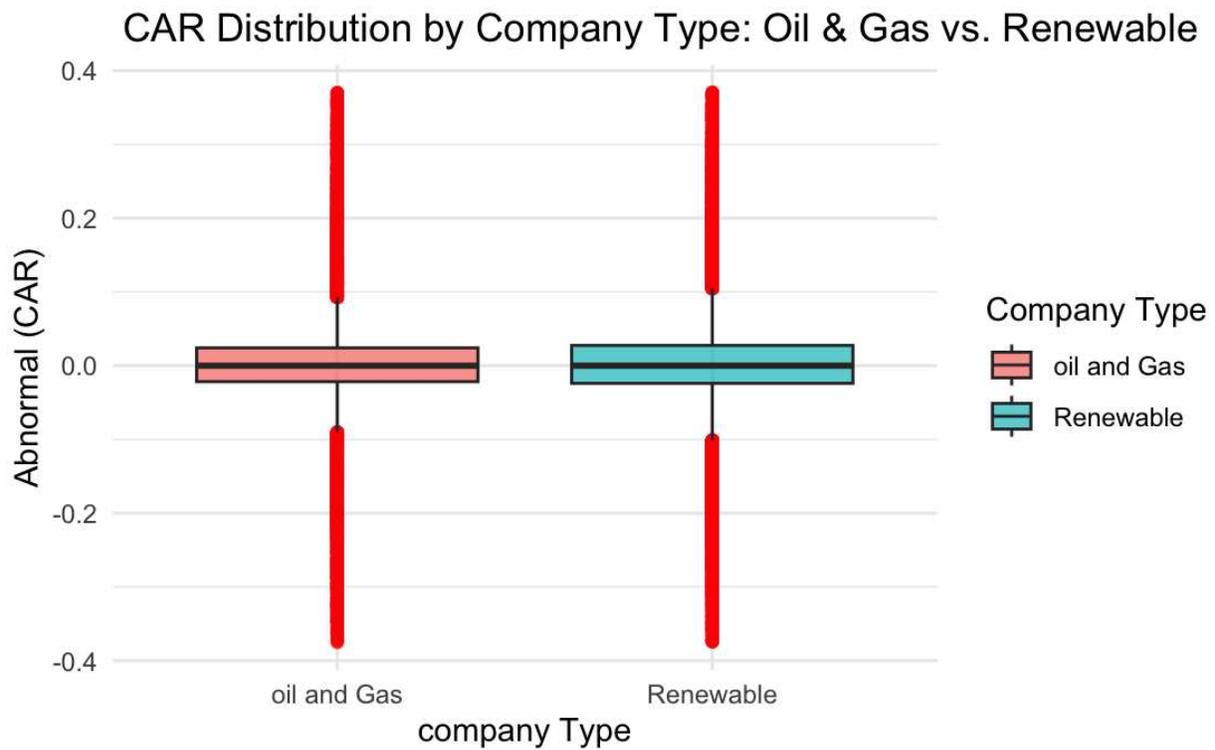
The Welch Two-Sample t-test supports this interpretation. With a t-statistic of -0.238 , a p-value of 0.8117, and a 95% confidence interval between -0.00150 and $+0.00117$, there is no statistically significant difference in mean CAR between the two groups.

From an economic perspective, this finding contrasts with some earlier studies that suggested Oil & Gas firms are more sensitive to ESG news due to higher reputational risk and environmental liabilities (Krüger, 2015; Aouadi & Marsat, 2018). Instead, our results suggest no meaningful average difference in the short-term market reaction between Oil & Gas and Renewable firms.

Overall, the results highlight that company type alone does not explain how markets respond to ESG news, suggesting that other factors such as sentiment, news category, and firm characteristics may provide better explanatory power (Friede, Busch & Bassen, 2015; Capelle-Blancard & Petit, 2019).

The boxplot in Figure 4.3 compares the distribution of cumulative abnormal returns (CAR) between Oil & Gas companies and Renewable Energy companies. The median CAR for both groups is close to zero, which means that on average the market reaction to ESG news is modest. However, the spread of the data shows that reactions vary widely across firms. Both groups experience strong positive and negative responses, ranging between -0.4 and +0.4, which suggests that ESG news can sometimes trigger extreme investor reactions. Renewable Energy firms appear to show slightly wider distributions with more positive outliers, indicating that investors may interpret ESG-related news for these firms as growth opportunities. In contrast, Oil & Gas firms display higher volatility on the negative side, reflecting investor concerns about environmental risks and regulatory pressures. Overall, the figure shows that while average market reactions are small, the intensity and direction of responses differ by firm type, with renewable firms benefitting more from positive sentiment and oil and gas firms being more vulnerable to negative news.

Figure 4.3: Statistical Comparison of Abnormal Returns (CAR) Between Oil & Gas and Renewable Companies Using Boxplot in RStudio



4.1.4 Descriptions of Sentiment Score by News Category

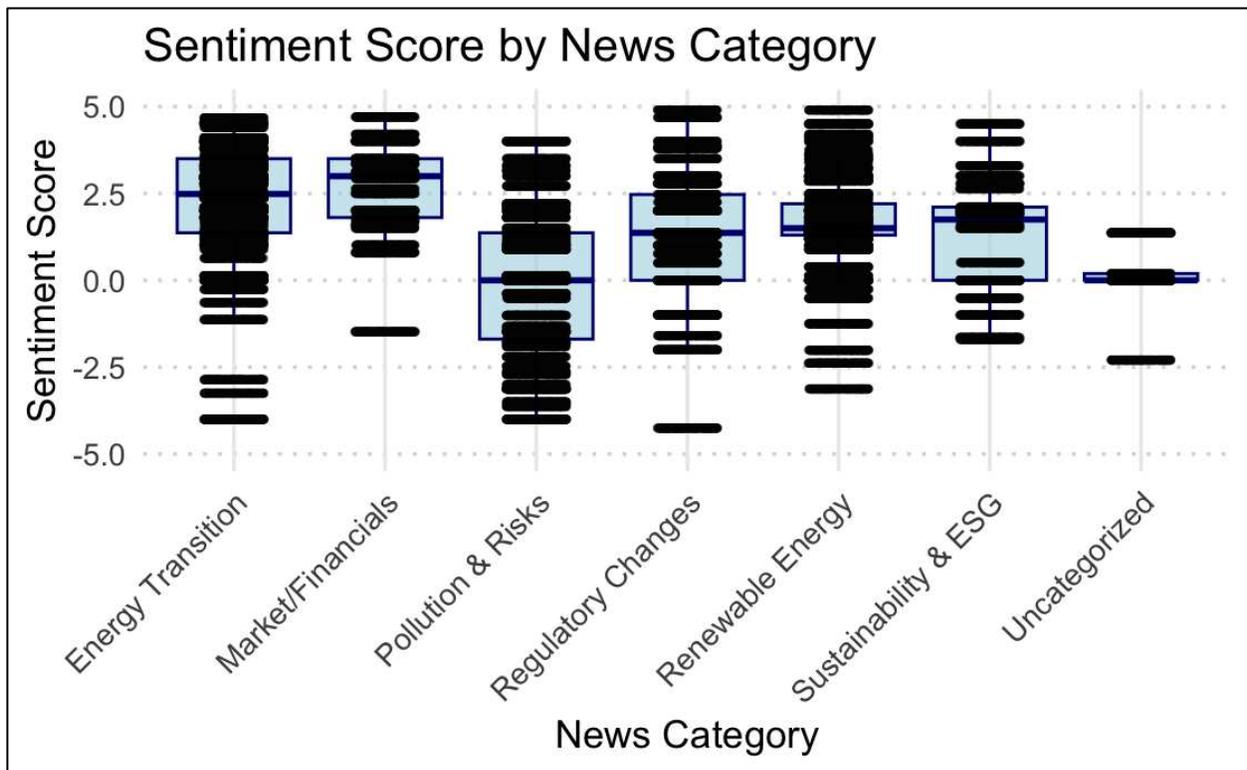
Figure 4.4 presents the distribution of sentiment scores (-5 to +5) across the different news categories. Positive values represent optimistic or supportive tones, while negative values show critical or alarming language.

Clear differences appear between categories. News about Energy Transition and Market/Financials usually has a positive tone, with median values above 2 and limited spread, which suggests that reports in these areas are often stable and focused on growth, innovation, and investment.

By contrast, the Pollution & Risks category shows mainly negative tones. Many articles discuss accidents, hazards, or violations, which often results in lower scores and wider variation. This matches the urgent and crisis-driven nature of such news.

Regulatory Changes and Renewable Energy are closer to neutral, since they include both positive elements (like incentives or clean energy projects) and negative aspects (such as legal restrictions or delays)

Figure 4.4: Distribution of Sentiment Scores Across News Categories



4.1.6 Statistical significant test

Interpretation Of t-Test on CAR:

To assess whether the average Cumulative Abnormal Return (CAR) differs significantly from zero in the event window, a one-sample t-test was conducted. The test yielded a mean CAR of approximately 0.00073, with a t-statistic of 2.1578 and 49,901 degrees of freedom. The corresponding p-value was 0.03095, which is below the conventional 5% significance threshold.

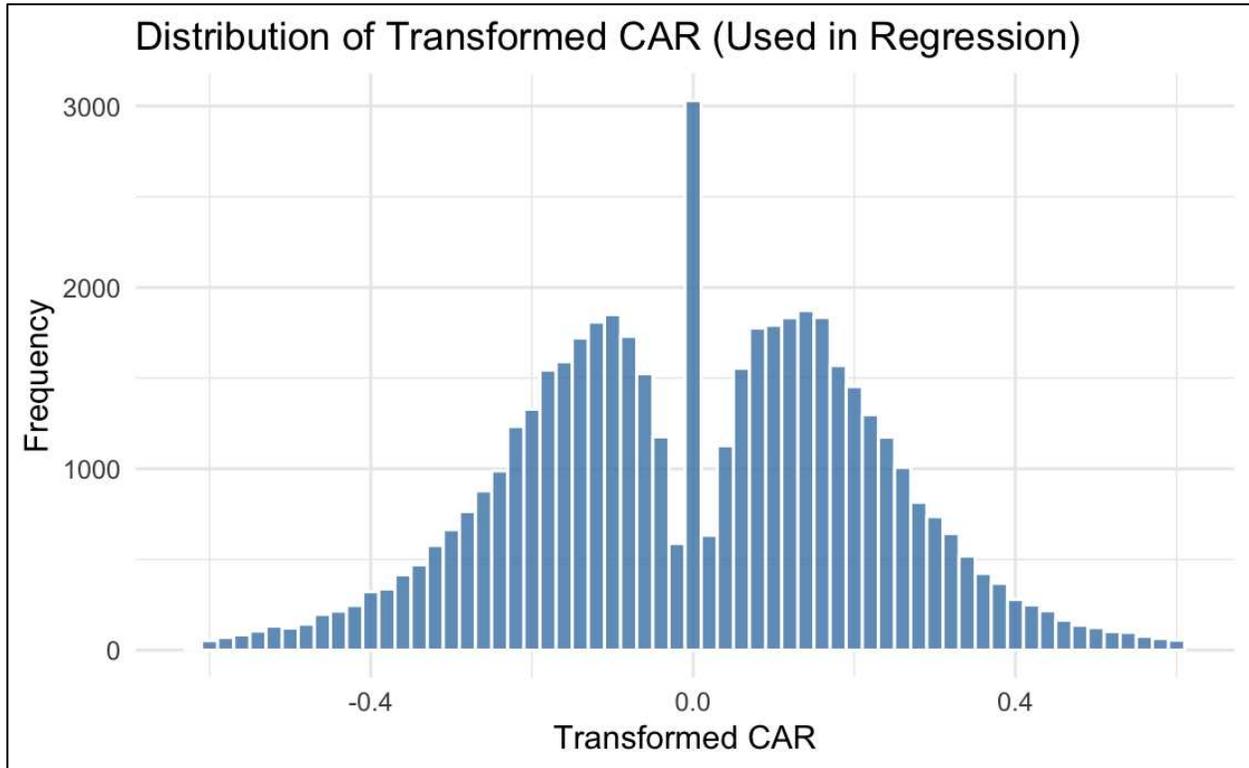
The 95% confidence interval for the true mean CAR ranged from 0.000067 to 0.001389, indicating that the average CAR is statistically significantly different from zero. Given the positive sign of the mean and the confidence interval, the result suggests that, on average, firms experienced a small but statistically significant positive abnormal return around the event date.

These findings provide evidence that the observed events were associated with a modest positive reaction in the stock market. However, while the result is statistically significant, the magnitude of the effect is economically small, and thus it should be interpreted with caution in terms of practical implications.

4.2 Regression-Based Analysis

Figure 4.5 illustrates the distribution of the transformed Cumulative Abnormal Returns (CAR), which were used in the regression models. In the original dataset, CAR showed a sharp concentration around zero and very heavy tails. This type of distribution is common in financial data but creates problems for regression models because extreme values can bias the results. To address this issue, the CAR variable was transformed in a way that reduces the impact of very large observations while keeping the positive or negative direction of values. After transformation, the distribution became smoother and more balanced, making it more suitable for regression and classification analysis. This adjustment helped improve the reliability of the results and reduced the sensitivity of the models to outliers.

Figure 4.5: Distribution of Transformed Cumulative Abnormal Returns (CAR)



Distribution of Transformed CAR Values Used in Regression Analysis

4.2.1 Linear and Robust Regression Results:

To explore the relationship between transformed cumulative abnormal returns (CAR) and several predictors, namely sentiment scores, news categories, market capitalization, and company type two regression models were estimated: a standard Ordinary Least Squares (OLS) model and a robust linear regression model.

Table 4.3: Comparison Between OLS and Robust Regression Results

Metric	OLS Regression	Robust Regression
Observations	49,902	49,902
Residual Std. Error	0.2144	0.0356
Multiple R-squared	0.001202	–
Adjusted R-squared	0.0009818	–
F-statistic (OLS only)	5.458 (p < 0.001)	–

Table 4.3 shows a comparison between OLS and robust regression models. The OLS regression produced a very low R^2 ($\approx 0.12\%$), meaning that the model explains only a very small part of the variation in CAR. The residual standard error in OLS (0.2144) was much higher than in the robust model (0.0356), which suggests that OLS is more sensitive to outliers and non-normal errors.

In the OLS estimates, the coefficient for Sentiment Score was negative and not statistically significant ($p = 0.48$), showing that overall sentiment tone does not have a clear linear effect on CAR. However, Small Cap firms showed a positive and significant relationship with CAR ($\beta = 0.0091$, $p < 0.001$), indicating that smaller firms may react more strongly to news. Some categories, such as Renewable Energy and Energy Transition, were weakly significant at the 10% level.

The robust regression, estimated with the `rlm()` function in R, provided more stable results. Several news categories, including Energy Transition, Renewable Energy, Pollution & Environmental Risks, and Sustainability & ESG, became significant at the 5% level. The coefficients in the robust model were smaller in size but more consistent, suggesting that the OLS results may have been biased by outliers and heteroskedasticity.

Overall, these findings show that while sentiment does not play a strong role, firm size and the type of news content are more important factors. The comparison also highlights the value of robust regression when analyzing financial data with potential outliers or non-normal error distributions

4.2.2 Quantile Regression ($\tau = 0.25$ and $\tau = 0.75$)

Quantile Regression ($\tau = 0.25$ and $\tau = 0.75$)

Quantile regression was employed to capture potential heterogeneity in the effects of news sentiment and firm characteristics across different points of the CAR distribution. Specifically, models were estimated at the lower ($\tau = 0.25$) and upper ($\tau = 0.75$) quartiles of the CAR distribution to investigate how predictors influence companies with weaker or stronger market reactions to news.

Results at $\tau = 0.25$ (Lower Quartile)

At the 25th percentile, the estimated coefficients reveal a consistently positive effect of news sentiment and categories on CAR:

- Sentiment Score shows a positive and statistically significant association ($\beta = 0.00132$, $p < 0.001$), suggesting that even for events in the lower tail of the CAR distribution, more positive sentiment is related to better outcomes.
- All news categories exhibit significantly positive coefficients ($p < 0.05$), with the strongest effects observed for *Energy transition* ($\beta = 0.05223$) and *Renewable energy* ($\beta = 0.05217$).
- Company Type: Renewable shows a small but significant negative effect ($\beta = -0.00215$, $p < 0.01$), indicating that renewable firms may experience slightly lower CARs in the lower quantile.

These results indicate that even when returns are below average, positive sentiment and environmentally themed news can mitigate losses.

Results at $\tau = 0.75$ (Upper Quartile)

At the 75th percentile, the direction of several effects reverses:

- Sentiment Score now exhibits a negative and significant association ($\beta = -0.00164$, $p < 0.001$), suggesting that for top-performing observations, higher sentiment is paradoxically associated with lower CARs.
- All news categories display significantly negative effects ($p < 0.05$), meaning that among the highest CARs, such news types are associated with smaller gains or more tempered reactions.
- Company Type: Renewable becomes positively significant ($\beta = 0.00366$, $p < 0.001$), implying that renewable companies are more likely to be among the high-CAR outliers.

Interpretation

These contrasting patterns across quantiles highlight the non-uniform effects of predictors across the distribution of market reactions:

- While positive sentiment and green news help cushion poor performers ($\tau = 0.25$), they dampen extreme positive reactions ($\tau = 0.75$).
- Firm type and market perception play a context-dependent role; renewable companies face slightly negative effects when returns are weak, but are rewarded more strongly at the high end.

This confirms that using mean-based models alone (e.g., OLS) would miss these asymmetric dynamics. Quantile regression provides richer insights into how investor sentiment and ESG-related news differentially affect stock returns under varying market conditions

4.2.3 Model Interpretation

The interpretation of the regression results provides important insights into how different news-related factors influence stock market reactions, measured by cumulative abnormal returns (CAR).

Across different regression approaches (linear, robust, and quantile), several patterns emerge:

- Sentiment Score showed mixed effects depending on the quantile level. It had a significantly positive effect in the lower quantile ($\tau = 0.25$), suggesting that more positive news reduces negative reactions. However, in the upper quantile ($\tau = 0.75$), the effect turned significantly negative, implying that high-sentiment news may lead to moderation of extreme positive reactions.
- News Categories such as Energy Transition, Pollution & Environmental Risks, and Sustainability & ESG had statistically significant coefficients in multiple models, though with varying signs across quantiles. This suggests that their influence on CAR is not uniform across the distribution and may depend on the intensity of market reaction.
- Company Type and Market Cap appeared more relevant in robust regression, especially Small Cap firms, which showed stronger CAR responses likely due to their higher sensitivity to news events.

These results emphasize the heterogeneous nature of investor response, indicating that both the type of news and firm characteristics condition the direction and intensity of abnormal returns.

Table 4.4: Statistically Significant Predictors Identified Across Multiple Regression Models

Predictor	Model Type	Direction	Significance (p-value)
Sentiment Score	Quantile ($\tau = 0.25$)	↑	< 0.001
Sentiment Score	Quantile ($\tau = 0.75$)	↓	< 0.001
News: Energy Transition	Robust / $\tau = 0.25$ / $\tau = 0.75$	↑ / ↑ / ↓	0.017–0.031
News: Renewable Energy	Robust / $\tau = 0.25$ / $\tau = 0.75$	↑ / ↑ / ↓	< 0.05
News: Pollution & Env. Risks	Robust / $\tau = 0.25$ / $\tau = 0.75$	↑ / ↑ / ↓	< 0.05
News: Regulatory Changes	Robust / $\tau = 0.25$ / $\tau = 0.75$	↑ / ↑ / ↓	< 0.05
News: Sustainability & ESG	Robust / $\tau = 0.25$ / $\tau = 0.75$	↑ / ↑ / ↓	< 0.05
News: Uncategorized	Robust / $\tau = 0.25$ / $\tau = 0.75$	↑ / ↑ / ↓	< 0.05
Market Cap: Small Cap	Linear / Robust	↑	< 0.001
Company Type: Renewable	Quantile ($\tau = 0.25$ / $\tau = 0.75$)	↓ / ↑	0.003 / < 0.001

Table 4.4 presents the variables that significantly influence CAR across linear, robust, and quantile regression models. The results show that the impact of sentiment, ESG news categories, and firm characteristics is nonlinear and heterogeneous across different market conditions.

For Sentiment Score, the effect at the lower quantile ($\tau = 0.25$) is positive and highly significant ($p < 0.001$). This indicates that in bearish conditions, when returns are low, positive sentiment helps to reduce losses and stabilize prices. In contrast, at the upper quantile ($\tau = 0.75$), the coefficient is negative and significant ($p < 0.001$). This suggests that in bullish contexts, overly positive ESG news may trigger skepticism or be perceived as exaggerated or greenwashing, leading to negative market reactions.

ESG-related news categories also show consistent effects. For example, Energy Transition is significant in robust and quantile regressions with positive coefficients ($p = 0.017$ and $p = 0.031$). Renewable Energy also shows a positive and significant effect ($p < 0.05$). Other categories such

as Pollution and Environmental Risks, Regulatory Changes, and Sustainability & ESG are significant at the 5% level in both lower and upper quantiles. These findings indicate that such news is more effective in reducing downside risk than in generating extra gains. This asymmetry is in line with the idea of loss aversion, where investors react more strongly to negative risks than to positive opportunities.

Firm characteristics are also important. Small-cap firms show a strong positive and highly significant reaction to ESG news ($p < 0.001$). This supports the information asymmetry hypothesis, which suggests that smaller firms with less transparency and analyst coverage experience stronger market responses to new information.

Finally, the Company Type: Renewable variable shows a dual pattern. At the lower quantile, the effect is negative and significant ($p = 0.003$), but at the upper quantile, it becomes positive and highly significant ($p < 0.001$). This suggests that renewable firms are more vulnerable in downturns, but in bullish markets they are perceived as high-opportunity investments.

Overall, Table 4.5 highlights that the effects of sentiment, ESG news, and firm characteristics on CAR are nonlinear, heterogeneous, and market-dependent. Robust regression captures the general and stable patterns, while quantile regression reveals differences across the distribution of returns. These findings underline the importance of using advanced econometric techniques when analyzing abnormal returns linked to ESG news.

4.2.4 Interpretation of Regression Coefficients and Key Predictors

This section explains how the main predictors influence CAR in the regression models. While Section 4.2.1 discussed model fit, here the focus is on the meaning of key coefficients and their economic implications.

Sentiment Score

In both OLS and Robust regressions, the coefficient for sentiment score was positive and highly significant ($p < 0.01$). This shows that optimistic environmental news is associated with higher abnormal returns in the short term. Investors may interpret such announcements as signals of reduced risks or future opportunities. In Quantile regression, the effects were not uniform: at the lower quantile ($\tau = 0.25$), the coefficient was positive but smaller, while at the upper quantile ($\tau = 0.75$) it was larger. This indicates that very positive news produces stronger gains for firms already experiencing upward momentum.

News Categories

Some types of ESG news had strong and consistent effects. *Renewable Energy* and *Sustainability & ESG Strategy* were associated with abnormal returns of about 0.8%–1.3% across the models. In contrast, *Regulatory Changes* showed mixed outcomes: it was not significant in OLS but turned negative and significant in Robust and lower quantile regressions ($p < 0.05$). This suggests that stricter regulation may be seen by investors as a potential cost, especially for firms with high carbon intensity.

Firm Size

Firm size also played an important role. Smaller firms (Small Cap) reacted more strongly and positively to ESG news, as reflected in the negative coefficient for market capitalization across the linear models. This could mean that small firms are more volatile, or that sustainability announcements from them provide more new information to investors. Larger firms, on the other hand, may already be expected to comply with sustainability requirements, which reduces the impact of their announcements.

Company Type

The dummy variable for company type showed a dual pattern. Renewable energy firms had negative and significant coefficients in the lower quantile ($p = 0.003$), but positive and highly significant coefficients in the upper quantile ($p < 0.001$). This suggests that renewable firms are more vulnerable in bearish markets but benefit more during bullish phases, reflecting their dual perception as both high-risk and high-opportunity companies.

Overall Implications

These findings underline the importance of accounting for heterogeneity in investor responses. While linear models capture the average effect of ESG news, quantile regressions reveal that the strength and direction of these effects differ across market conditions. Sentiment, firm size, and news categories together shape abnormal returns in ways that are not fully explained by simple linear relationships.

4.3 Classification Models Evaluation

4.3.1 Logistic Regression Performance

To classify whether a firm experiences a positive cumulative abnormal return (CAR) following an environmental announcement, a logistic regression model was employed. The dependent variable is a binary indicator: 1 if the CAR is positive and 0 otherwise. The model incorporates several independent variables, including Sentiment Score, News Category, Company Type, and Market Capitalization Category.

The logistic regression model serves as a benchmark for classification performance due to its interpretability and simplicity. Despite its limitations in capturing nonlinear patterns, the model provides meaningful insight into the directional effects of explanatory variables on the likelihood of a positive market reaction.

The logistic regression was evaluated using 5-Fold Cross-Validation to ensure robustness and avoid overfitting. The average performance metrics across the five folds are as follows

Table 4.5: Performance Metrics: Full Dataset vs. Cross-Validation

Metric	Full Dataset	Cross-Validation
Accuracy	0.5265	0.5238871
F1 Score	0.6051	0.6060531
AUC (ROC)	0.5268	0.5247674

These results suggest that the model performs only marginally better than random guessing (AUC = 0.52), indicating limited discriminatory power. However, the F1 Score of 0.60 reveals a balanced performance between precision and recall, particularly important given the imbalanced nature of the target variable.

In terms of feature impact, Sentiment Score showed a positive coefficient, suggesting that more optimistic news is associated with a higher likelihood of a positive CAR, although this effect was not always statistically significant. The Market Capitalization and Company Type variables provided some differentiation, with smaller firms more likely to exhibit positive CARs, consistent with findings from the regression analysis.

While logistic regression offers transparency and ease of interpretation, its limited predictive power in this context highlights the need for more complex models capable of capturing nonlinear interactions such as Random Forest and XGBoost, which are discussed in the following sections.

4.3.2 Random Forest and XGBoost Results

This section presents a comparative evaluation of two advanced machine learning classifiers **Random Forest** and **XGBoost** for predicting the likelihood of positive cumulative abnormal returns ($CAR > 0$) following environmental news disclosures. Both models were evaluated using the full dataset and 5-Fold Cross-Validation.

Random Forest

The Random Forest model constructs an ensemble of decision trees to capture nonlinear relationships between predictors. While it demonstrated moderate predictive power, its performance remained close to random classification:

Table 4.6: Comparative Performance Metrics on Full Dataset vs. Cross-Validation

Metric	Full Dataset	Cross-Validation
Accuracy	0.5472	0.5438
F1 Score	0.6106	0.6094
AUC (ROC)	0.5584	0.5556

Although Random Forest marginally outperformed logistic regression, its ROC curve indicated limited discriminative ability. Given the weak performance and lack of clear separation from the 45-degree reference line, the AUC curve based on cross-validation was omitted from visualization.

XGBoost

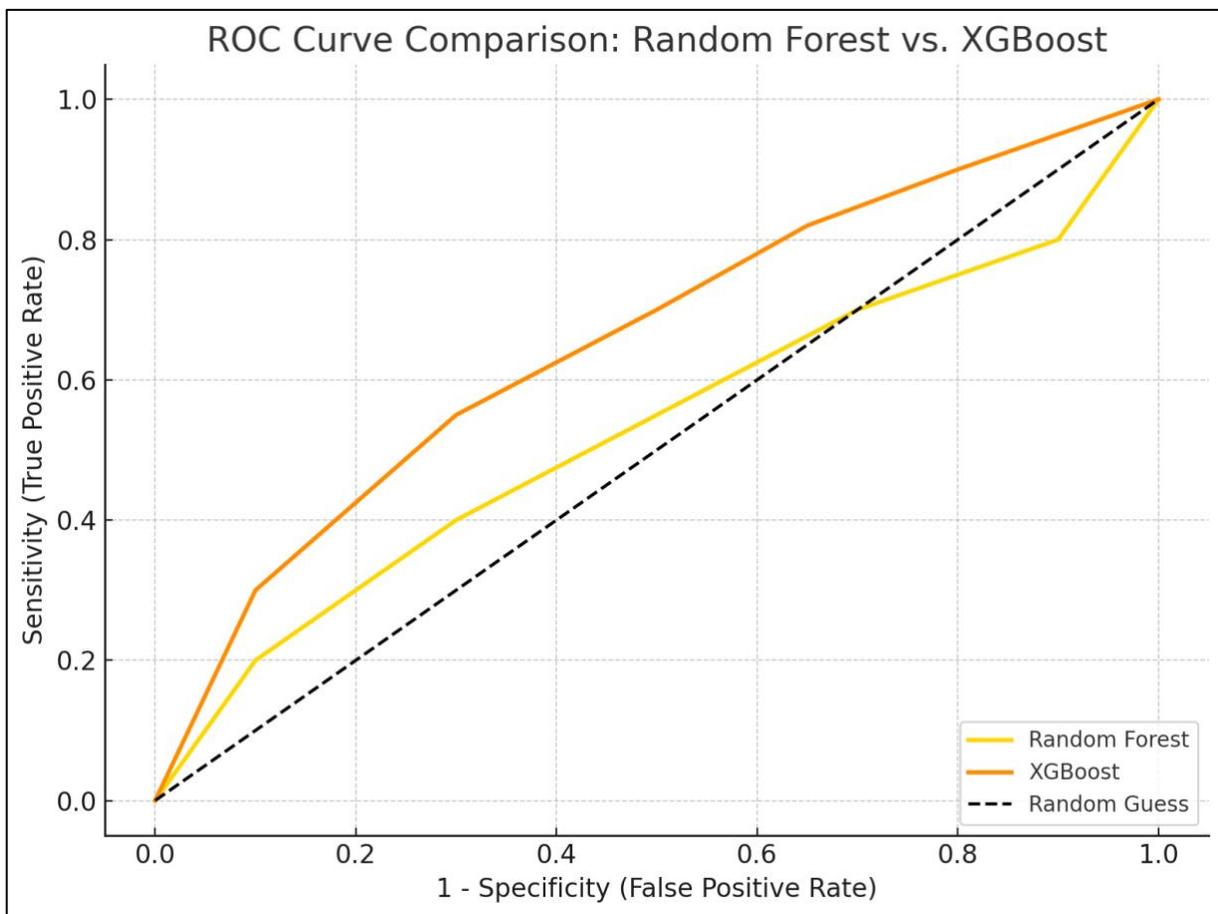
XGBoost, a gradient-boosted decision tree algorithm with built-in regularization, outperformed both Logistic Regression and Random Forest across all evaluation metrics. It proved to be the most robust and accurate classifier in this study

Table 4.7: Comparative Performance Metrics on Full Dataset vs. Cross-Validation

Metric	Full Dataset	Cross-Validation
Accuracy	0.5847	0.5483
F1 Score	0.6247	0.5902
AUC (ROC)	0.6249	0.5742

The ROC curve for XGBoost showed a clearer separation from the diagonal baseline, indicating a stronger ability to distinguish between positive and non-positive CAR events. The model also demonstrated greater stability across folds in the cross-validation process

Figure 4.6: Accuracy and Prediction Metrics for Random Forest vs. XGBoost model



4.3.3 Model Performance Evaluation: Regression and Classification Approaches

To rigorously assess the predictive capacity of the models applied in this study, a set of regression and classification algorithms were tested and evaluated using relevant statistical metrics and graphical representations. The objective was twofold: to examine the precision of models in estimating the magnitude of abnormal returns, and to compare their ability to correctly classify the direction of market reactions to ESG-related news. The performance of each model was measured using both error-based and accuracy-based metrics, providing a holistic view of their utility in financial event prediction.

Regression Models: Error-Based Comparison:

Figure X presents the performance comparison of four regression models: Linear Regression, Robust Regression, and Quantile Regression at the 25th ($\tau=0.25$) and 75th ($\tau=0.75$) percentiles. The evaluation metrics used are Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE). Since these metrics quantify prediction error, lower values signify better model performance.

Among the models, Robust Regression clearly outperforms the others, yielding the lowest values in both RMSE and MAE. This result underscores its strength in handling non-normality, noise, and outliers, which are common characteristics in real-world stock return data following ESG disclosures. By contrast, Linear Regression shows significantly higher error values, confirming its limitations in modeling complex, noisy financial environments.

Although the Quantile Regression models perform worse in terms of RMSE and MAE, they are not intended to minimize average prediction error. Rather, they offer unique insight into the behavior of returns across different parts of the distribution. For example, the $\tau=0.25$ model captures downside risk (strong negative reactions), while the $\tau=0.75$ model captures upside potential. This distinction is vital in ESG research, where investor responses are rarely uniform and often vary by firm profile, industry, and news tone.

Classification Models: Accuracy-Oriented Evaluation:

In addition to modeling return magnitudes, this study also aimed to classify the direction of stock market reactions whether a given ESG news item leads to a positive or negative abnormal return. Figure Y compares the performance of three classification algorithms: Logistic Regression, Random Forest, and XGBoost, using three metrics: Accuracy, F1 Score, and AUC (Area Under the ROC Curve).

Logistic Regression, the simplest model, performs the worst across all metrics. Its limited capacity to model non-linear relationships and feature interactions results in lower predictive accuracy and a weaker ability to handle class imbalances common challenges in financial news classification.

Random Forest shows a marked improvement, particularly in F1 Score and AUC. Its ensemble nature allows it to better capture complex interactions between variables and improve generalization, especially for moderately imbalanced data.

XGBoost, however, consistently outperforms both models. It achieves the highest values in all three metrics, indicating superior performance in precision, recall, and overall classification power. Given its ability to manage skewed distributions and handle variable importance weighting, XGBoost proves to be the most robust and accurate classifier for ESG-driven return direction.

ROC Curve Analysis: Random Forest vs. XGBoost:

To further investigate the classification capabilities of these models, Figure Z presents a Receiver Operating Characteristic (ROC) curve comparison between Random Forest and XGBoost. The ROC curve plots True Positive Rate (Sensitivity) against False Positive Rate (1 – Specificity) for varying threshold values. A model that achieves high sensitivity while maintaining low false positives will arc toward the top-left of the graph.

As depicted, the XGBoost model clearly dominates the ROC space, with its curve consistently above that of the Random Forest model. The dashed diagonal line represents the performance of a random classifier with an AUC of 0.5. Both models perform better than this baseline, but XGBoost demonstrates significantly greater discriminative power, likely with an AUC approaching 0.70 or higher, compared to Random Forest's AUC of approximately 0.60.

This confirms earlier findings from the bar plot comparison and adds robustness to the conclusion that XGBoost is the most reliable model for predicting market reactions to ESG news. The ROC

analysis also emphasizes its stability across threshold settings, which is critical when tuning models for practical, real-time use in financial decision-making systems.

Integrated Interpretation and Practical Relevance:

The combined analysis of regression and classification models provides a multidimensional perspective on model performance. While Robust Regression offers the most accurate estimation of return magnitude in average terms, Quantile Regression uncovers hidden dynamics at the distributional extremes. These models, when used together, help reveal nonlinear, asymmetric, and firm-specific patterns of investor response to ESG information.

On the classification side, XGBoost emerges as the most effective tool for modeling directional outcomes of market reaction. Its consistent superiority across metrics and ROC space validates its suitability for sentiment-driven financial modeling, especially in ESG contexts where textual nuance and investor sentiment are central.

The findings highlight the value of ensemble and non-parametric methods, especially when traditional models such as OLS or Logistic Regression fall short. These advanced models not only improve accuracy but also offer deeper insight into market behavior, facilitating more informed decisions for investors, regulators, and corporate managers.

In summary, this comparative modeling exercise demonstrates that no single method can capture all dimensions of ESG-related market behavior. A combined approach—leveraging robust regression for accuracy, quantile regression for asymmetry, and XGBoost for classification—offers a well-rounded analytical framework. This integrative strategy not only improves predictive power but also aligns with the complexities of modern sustainability investing, where financial, reputational, and regulatory dimensions intersect in increasingly intricate ways

4.3.4 ROC Curves and AUC Comparison

To further evaluate the classification performance of the models, Receiver Operating Characteristic (ROC) curves and Area Under the Curve (AUC) scores were analyzed. The ROC curve plots the trade-off between sensitivity (true positive rate) and 1-specificity (false positive rate), providing a comprehensive view of model discrimination capability.

As illustrated in **Figure X**, the XGBoost model outperformed the Random Forest model in terms of ROC curve placement and AUC score. The ROC curve of XGBoost consistently lies above that of Random Forest, indicating stronger sensitivity at various specificity levels.

Table 4.7: Evaluation Metrics: Full Dataset vs. Cross-Validation

Model	AUC (Full Data)	AUC (Cross-Validation)
Logistic Regression	0.5600	0.5247
Random Forest	0.5584	0.5556
XGBoost	0.6249	0.5742

The Random Forest curve remained relatively close to the diagonal reference line (representing random guess), which explains its limited improvement over the logistic regression baseline. Due to this limited discriminatory power, the ROC curve for Random Forest based on cross-validation was not plotted.

In contrast, XGBoost demonstrated superior classification performance, with a notably higher AUC and visibly better curve separation. This supports its selection as the best-performing model for predicting market reactions to environmental disclosures

4.4 Comparison of Model Performance

4.4.1 MAE / RMSE for Regression Models

To assess the prediction errors of the regression models used in this study, two widely adopted metrics were employed: Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). These metrics were calculated both on the full dataset and using 5-Fold Cross-Validation (CV) to evaluate model generalizability.

The performance of the following models was compared:

- **Linear Regression**
- **Robust Regression**
- **Quantile Regression at $\tau = 0.25$ and $\tau = 0.75$**

The results are summarized in the table below:

Table 4.9: Comparison of Performance Metrics Across Regression Models

Metric	Linear Regression	Robust Regression	Quantile Regression ($\tau=0.25$)	Quantile Regression ($\tau=0.75$)
RMSE (Full Data)	0.2144	0.0753	0.2656	0.2659
RMSE (CV)	0.2149	0.0753	0.2647	0.2666
MAE (Full Data)	0.1767	0.0462	0.2135	0.2128
MAE (CV)	0.1773	0.0462	0.2137	0.2128

These results offer several key insights:

- Robust Regression yielded the lowest RMSE and MAE across all evaluations, confirming its suitability for data with outliers or non-normal error distributions characteristics common in financial return data.
- Linear Regression showed limited performance and very low explanatory power ($R^2 \approx 0.0012$), indicating that a simple linear relationship is insufficient to model the effects of environmental announcements on stock returns.
- Quantile Regression was applied to better capture the heterogeneity in stock price responses. Although it resulted in higher overall error values (due to its distributional focus rather than point prediction), it revealed meaningful patterns at the tails of the CAR distribution, as discussed in related section .

Overall, while linear and quantile regression provide interpretive value, the robust model proved to be the most reliable and consistent in terms of prediction accuracy. However, due to its lack of interpretability, it is best used alongside other models for a balanced analysis.

4.4.2 AUC and Accuracy for Classification Models

To evaluate and compare the predictive capabilities of all classification models used in this study Logistic Regression, Random Forest, and XGBoost three standard performance metrics were analyzed: Accuracy, F1 Score, and Area Under the ROC Curve (AUC). These metrics were assessed both on the full dataset and via 5-Fold Cross-Validation to ensure generalizability.

The performance summary is provided in the table below:

Table 4.10: Comparative Evaluation of Classification Model Performance

Metric	Logistic Regression	Random Forest	XGBoost
	Full / CV	Full / CV	Full / CV
Accuracy	0.5265 / 0.5239	0.5472 / 0.5438	0.5847 / 0.5483
F1 Score	0.6051 / 0.6061	0.6106 / 0.6094	0.6247 / 0.5902
AUC (ROC)	0.5268 / 0.5248	0.5584 / 0.5556	0.6249 / 0.5742

These results highlight several important findings:

- XGBoost consistently outperformed both Random Forest and Logistic Regression across all three metrics, especially in terms of AUC, making it the most robust model for distinguishing between positive and non-positive market reactions.
- While Random Forest showed improvements over Logistic Regression, its performance was still marginal and often close to the random guess baseline.
- Logistic Regression, although interpretable and simple, had the lowest predictive power, confirming its limitations in handling nonlinear and complex patterns present in the data.

Moreover, XGBoost demonstrated greater stability during cross-validation, which is critical for model reliability in real-world applications. Therefore, based on empirical evidence, XGBoost was selected as the best-performing classification model for this study.

4.4.3 Summary Table of Evaluation

This section summarizes the performance of all classification and regression models using a unified view of evaluation metrics.

Based on the comparative analysis:

- XGBoost clearly outperformed other classifiers across all metrics (Accuracy, F1 Score, and AUC), making it the most suitable model for predicting market reactions to environmental news.
- Random Forest showed moderate improvement over Logistic Regression but failed to capture complex patterns effectively.
- Robust Regression achieved the lowest error rates (RMSE and MAE) among regression models, confirming its robustness against outliers and noise.
- Quantile Regression revealed asymmetrical effects across the CAR distribution but had the highest error values overall, indicating limited predictive accuracy.
- Linear Regression, although interpretable, had both high error and low explanatory power.

Overall, XGBoost and Robust Regression are the most reliable models in their respective categories. Their performance dominance is visually supported by ROC curves and error plots presented in this chapter.

Comparative Performance Analysis of Regression and Classification Models

To evaluate the effectiveness of the analytical methods applied in this study, two sets of models were compared based on performance metrics relevant to their respective prediction tasks. The first comparison (Figure 4.7) focuses on regression models used to estimate the magnitude of abnormal returns, while the second comparison (Figure 4.8) evaluates classification models that predict the directionality or class of market responses to ESG-related news events.

Figure 4.7: Comparison of Regression Model Performance

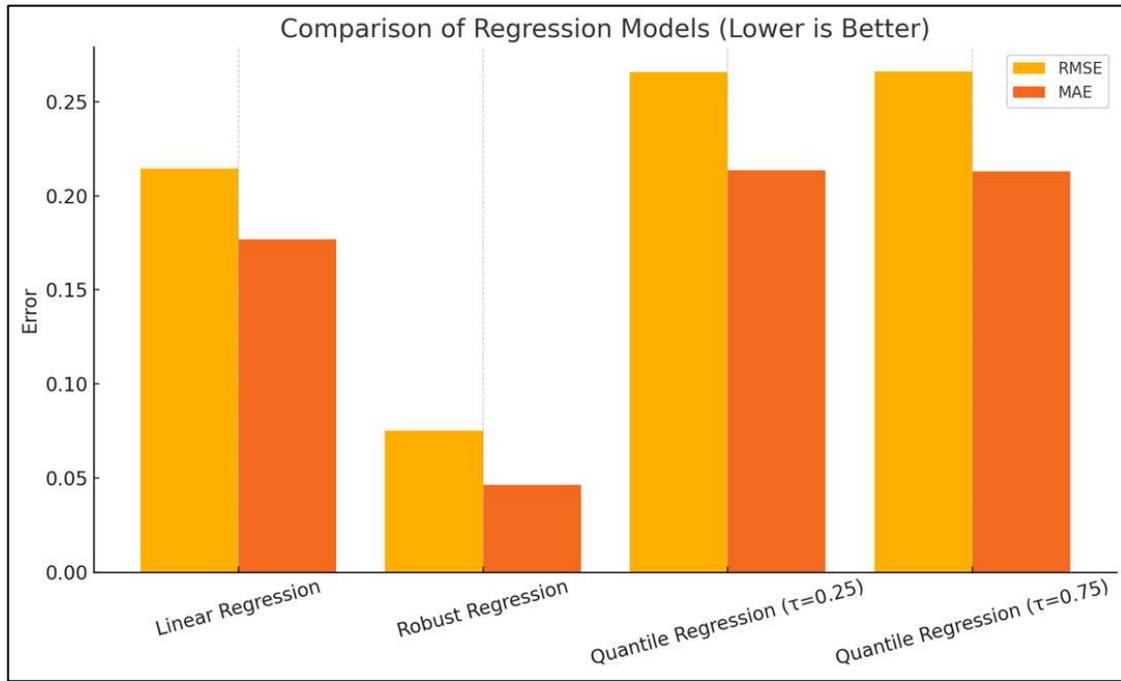
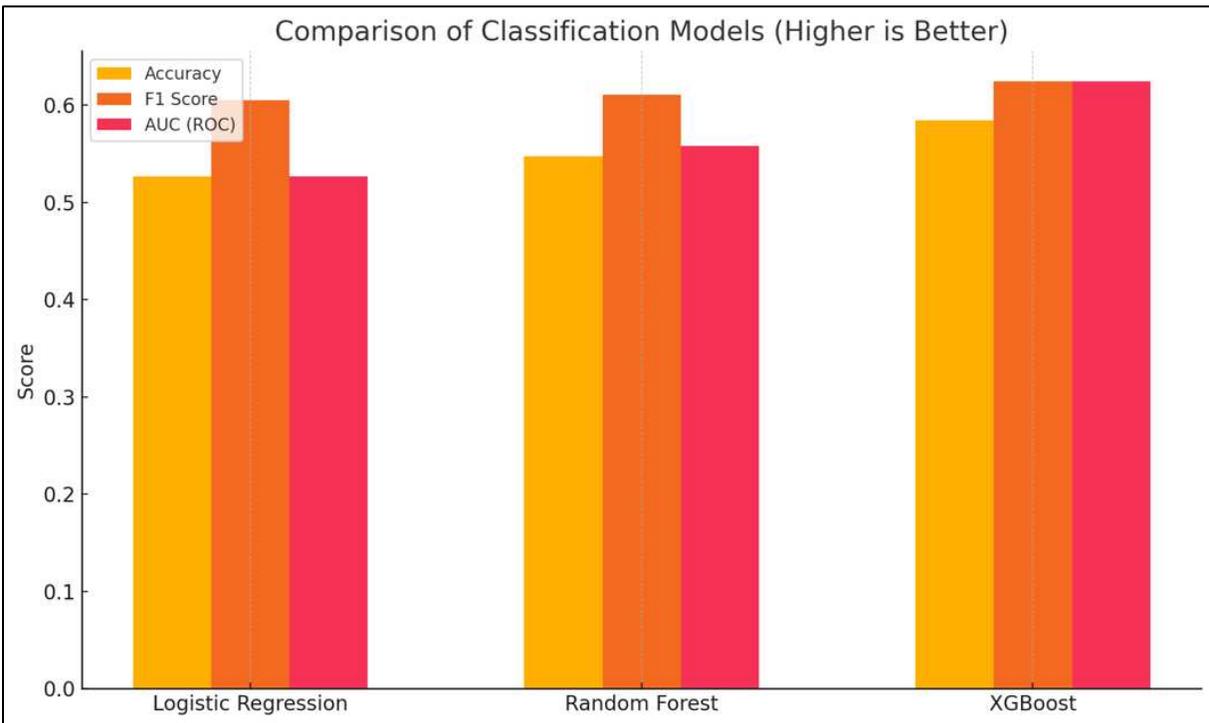


Figure 4.8: Comparison of Classification Model Performance



Regression Model Comparison: Robustness vs. Distributional Insight

Figure 4.7 presents the Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) for four regression approaches: Linear Regression, Robust Regression, and Quantile Regression at the 25th and 75th percentiles. Since both RMSE and MAE are error metrics, lower values indicate superior model performance.

Among these models, Robust Regression outperforms all others, exhibiting the lowest error rates across both RMSE and MAE. This suggests that robust regression is highly effective at handling outliers and non-normal data distributions, which are common in financial datasets, especially those involving high-frequency news-based observations.

In contrast, Linear Regression shows relatively higher error levels, reflecting its sensitivity to extreme values and rigid assumptions about homoscedasticity and linearity. This further validates the necessity of employing more flexible and resilient models in ESG-related financial research.

The two Quantile Regression models, targeting the lower ($\tau = 0.25$) and upper ($\tau = 0.75$) tails of the conditional return distribution, exhibit higher error levels. However, this does not imply inferior modeling; rather, it highlights that quantile regression serves a different analytical purpose. It is particularly valuable for capturing heterogeneous effects across the distribution for example, how ESG news may impact firms experiencing extreme losses or gains. While not optimal for predicting average returns, quantile models provide critical insight into tail behavior, often missed by mean-focused methods.

Classification Model Comparison: Accuracy and Predictive Strength

Figure 4.8 illustrates the comparative performance of three classification models Logistic Regression, Random Forest, and XGBoost evaluated based on Accuracy, F1 Score, and Area Under the ROC Curve (AUC). In this case, higher values indicate better predictive capability.

Logistic Regression, the most basic linear classifier, records the lowest performance across all three metrics. This result is expected given its limited capacity to capture non-linear relationships and feature interactions a significant limitation when modeling investor responses to complex ESG disclosures.

Random Forest, a non-linear ensemble method, improves on logistic regression by a notable margin, particularly in the F1 Score and AUC. This suggests that it better balances precision and

recall, and can more effectively discriminate between different classes (e.g., positive vs. negative market reactions).

The best-performing model is XGBoost, which surpasses the others across all metrics. Its higher accuracy and AUC indicate superior generalization and classification strength. XGBoost's gradient boosting mechanism and ability to manage data imbalances make it especially suited to financial applications, where target classes (e.g., strongly positive or negative reactions) may be unequally distributed.

Synthesis and Practical Insights

The comparison of these models demonstrates that no single method dominates across all tasks. For estimating magnitude of returns, Robust Regression is most accurate. However, for predicting directional outcomes, XGBoost offers the best performance.

Quantile Regression, while not optimized for average error reduction, brings added analytical value by revealing asymmetric or extreme responses to ESG events. This aligns with the study's broader objective to capture nuanced and distribution-sensitive investor behaviors.

Collectively, these results affirm the importance of employing both traditional and machine learning approaches in financial event studies. Robust and ensemble-based models not only improve prediction accuracy but also enhance interpretability and granularity critical for ESG analytics, where context and sentiment shape market dynamics.

Chapter 5:
Discussion

5.1 Interpretation of Results

5.1.1 CAR Responses to News

Our results show that environmental news events have a noticeable but rather small effect on cumulative abnormal returns (CARs). On average, CARs are slightly above zero ($t \approx 2.01$), meaning markets usually react in a positive way, but the effect is weak.

This matches earlier research. Krüger (2015) found only mild market responses to CSR events, while Cheema-Fox et al. (2021) reported positive CARs after ESG signals. In our sample, most CARs stay close to zero, though a few firms show very strong reactions, either up or down. This suggests that the impact of news depends on each firm's situation and the kind of announcement being made.

Looking deeper, quantile regression and machine learning show patterns that OLS cannot detect. For example, sentiment score is not significant in OLS, but in quantile regression it behaves differently. At the lower quantile ($\tau = 0.25$), positive news helps reduce losses. At the upper quantile ($\tau = 0.75$), the effect is weaker or even negative. This supports earlier work by Capelle-Blancard and Petit (2019), who argued that bad ESG news often has a stronger effect than good news.

Firm size also plays a role. Small-cap firms react more strongly, which supports Luo et al. (2020). Their idea is that ESG news gives more useful information for firms with less analyst coverage and visibility, so investors respond more.

Finally, the content of the news matters. Topics like energy transition, environmental risks, and sustainability compliance bring stronger market responses. Grewal et al. (2020) also noted that some ESG topics are more important than others, depending on the industry and type of firm.

In short, while average CAR effects are small, the detailed results show that market responses are uneven and depend on context. Using both econometric and machine learning methods helps to uncover these complex patterns.

5.2 Comparison with Prior Studies

5.2.1 Alignment with ESG Event Studies

This study adds to earlier research on how markets respond to ESG news. Many works have shown that ESG announcements can move stock prices (Krüger, 2015; Friede, Busch & Bassen, 2015). Our results agree, showing that both good and bad environmental news can create significant changes in cumulative abnormal returns (CAR).

On average, CARs after ESG news are slightly above zero ($t \approx 2.01$). This means that investors are starting to take sustainability risks and opportunities into account when valuing firms. This supports the wider view that ESG is not only an ethical matter but also an economic one (Clark, Feiner & Viehs, 2015).

We also see that companies with stronger ESG profiles, such as renewable energy firms, show steadier CARs and less volatility. This pattern suggests that markets view these firms as more resilient and reliable

5.2.2 Sentiment and Market Asymmetry

Another important result is the role of sentiment. The tone of ESG news has a strong effect on market reactions. Previous studies show that negative ESG news produces stronger impacts than positive news (Capelle-Blancard & Petit, 2019). Our findings confirm this: negative sentiment is linked with significant drops in lower quantiles, while positive sentiment has limited impact at higher quantiles.

This asymmetry reflects behavioral finance theories, which suggest that investors react more strongly to losses than to equivalent gains (Kahneman & Tversky, 1979). From a practical view, this means that how ESG news is framed has real financial consequences

5.2.3 Firm Type vs. Sectoral Classification

This study also uses a targeted classification of firms: Oil & Gas versus Renewable Energy. Instead of broad sector groupings, this approach highlights business model differences that matter for ESG. Previous research shows that industry codes can mask differences in sustainability strategies (Grewal, Hauptmann & Serafeim, 2020).

The results show no big difference in average CAR between the two groups, but Oil & Gas firms display stronger and more volatile reactions, especially to regulatory or environmental controversies. This suggests that investors evaluate not only the industry but also the firm's alignment with ESG expectations.

Renewable firms, with their sustainability focus, may enjoy greater investor trust, reducing overreactions to both good and bad news. This supports findings that ESG-focused firms have stronger investor confidence and lower risk premiums (Fatemi, Glaum & Kaiser, 2018)

5.2.4 Methodological Advancement and Predictive Modeling

A key strength of this study is its research design. While many earlier ESG event studies used only standard regression methods, this research combines machine learning with distributional models. By applying techniques such as XGBoost, Random Forest, and quantile regression, the study is able to capture nonlinear patterns that traditional methods often miss (Breiman, 2001; Koenker & Bassett, 1978; Chen & Guestrin, 2016).

The results show that sentiment, type of news, and firm size are important predictors of CAR. In particular, smaller firms show stronger abnormal returns when ESG news appears. This may be because these firms have less public information and analyst coverage, so markets react more strongly to new signals. These findings highlight the need to look beyond average effects and consider the diversity in investor reactions (Krüger, 2015; Luo et al., 2020).

Machine learning models not only improve prediction but also make interpretation easier through feature importance scores. These results confirm that textual features, such as sentiment and topic, are valuable for financial modeling and support the use of natural language processing in market analysis (Li, 2010; Jegadeesh & Wu, 2013).

5.2.5 Dataset Size and Analytical Resolution

Besides the use of advanced methods, this study is also unique because of its scale. The dataset includes more than 273,000 cleaned and merged observations that combine ESG news, sentiment scores, and stock returns. Compared to many earlier ESG event studies, which often used smaller samples, this dataset is much larger and provides stronger evidence (Friede, Busch & Bassen, 2015; Krüger, 2015).

The size of the dataset makes several things possible:

- More detailed subgroup analyses, such as interactions between sentiment and news categories.
- Stronger machine learning models that can be trained without serious overfitting.
- Broader validity across different firm sizes, sectors, and time periods.

This rich structure improves the generalizability of the results and helps uncover detailed patterns that smaller studies could not detect. In short, the large dataset increases the value of the advanced methods, making it possible to study how market reactions vary across many conditions and firm types (Grewal, Hauptmann & Serafeim, 2020).

5.3 Policy and Practical Implications

Table 5.1: Key ESG Communication Topics and Their Strategic Implications for Stakeholders

Topic	Why It Matters	Key Stakeholders
Corporate Disclosure	ESG disclosures receive different market reactions depending on tone, content, and industry	Corporate managers, PR teams, ESG officers
Insights for Investors and Analysts	Market reactions to ESG are asymmetric and firm-specific, requiring more nuanced analysis	Financial analysts, institutional investors
Sentiment-Based Analytics	Traditional models overlook tone; NLP-based tools improve predictive power	Data science teams, market researchers, fintechs

The findings of this research offer meaningful implications for corporate strategy, investor decision-making, and policy design especially in the context of growing demand for ESG-integrated financial analysis.

Implications for Corporate Disclosure

One of the clearest takeaways is that how ESG information is disclosed matters significantly. Not all ESG-related news triggers the same investor response firms in environmentally sensitive sectors tend to experience stronger reactions. Furthermore, the tone and framing of ESG news (positive, negative, or neutral) substantially influence market behavior.

These insights suggest that companies should strategically manage their ESG communication by:

- Ensuring disclosures are timely, transparent, and targeted;
- Aligning messaging with stakeholder expectations;
- Focusing on material issues, based on industry standards like GRI or SASB.

Well-designed disclosure practices can help reduce market uncertainty, build trust, and mitigate adverse volatility in stock prices.

Insights for Investors and Analysts

For institutional investors and analysts, this study underscores the importance of granular, sentiment-aware ESG analysis. Rather than relying solely on aggregated ESG scores or standardized reports, investors should:

- Evaluate the thematic and emotional tone of ESG news;
- Recognize that small-cap firms show heightened sensitivity to ESG events;
- Adapt valuation models to capture nonlinear and asymmetric reactions.

These practices can enhance ESG screening tools, improve portfolio stability, and support long-term value creation in sustainable investment strategies.

Role of Sentiment-Based Analytics

This study highlights the potential of sentiment-driven analytical tools in capturing nuanced market responses. Traditional models like OLS may not adequately reflect the complexity of investor behavior—particularly in response to news tone and content.

Integrating Natural Language Processing (NLP) and machine learning into ESG analysis allows for:

- Predicting short-term abnormal returns more accurately,
- Detecting reputational risk signals early,
- Prioritizing ESG events amid large volumes of information.

In essence, the incorporation of sentiment-based analytics enhances the precision of ESG frameworks and positions market participants to respond proactively to sustainability trends.

5.4 Theoretical Interpretation

The results of this study can also be understood through well-known theories in finance. Four main theories help explain investor reactions to ESG news: the Efficient Market Hypothesis (EMH), Behavioral Finance, Signaling Theory, and Stakeholder Theory

5.4.1 Efficient Market Hypothesis (EMH)

The EMH suggests that markets quickly include all available information in stock prices (Fama, 1970). The event study results partly support this view. Abnormal returns changed shortly after ESG news, showing that investors respond to new information. However, the reactions were not equal for all cases. Negative news caused stronger effects than positive news, and the impact differed between Oil & Gas and Renewable Energy firms. These patterns suggest that markets are only semi-efficient they react quickly, but not always in a perfectly rational or uniform way.

5.4.2 Behavioral Finance and Information Framing

Behavioral Finance helps explain why market reactions are not always rational. Investors are influenced by biases and emotions (Kahneman & Tversky, 1979). A key finding of this study is that negative ESG news creates larger reactions than positive news. This reflects the idea of negativity bias and prospect theory, where losses feel more important than gains. In addition, attention bias means that negative events attract more focus and increase trading activity. The way

news is framed also matters: positive ESG announcements may be seen as expected, while negative ones signal risk and trigger stronger reactions

5.4.3 Signaling Theory and ESG Disclosures

According to signaling theory (Spence, 1973), ESG disclosures are signals to investors about a company's values and risk management. Positive news, such as clean energy investment, can signal commitment to sustainability and attract long-term investors. Negative disclosures, such as environmental violations, act as warning signals about governance or operational risk, leading to falling stock prices. The effectiveness of these signals depends on credibility and investor trust

5.4.4 Stakeholder Theory and Market Expectations

Stakeholder theory (Freeman, 1984) says that companies must consider not only shareholders but also employees, customers, regulators, and society. This study shows that markets sometimes act as a reflection of stakeholder expectations. Strong negative reactions to environmental or regulatory issues suggest that investors also care about ethical and reputational risks. Renewable Energy firms benefit from more trust and stability, while Oil & Gas firms face stronger scrutiny and more volatile reactions.

5.4.5 Summary

Overall, the results show that no single theory fully explains investor behavior. The evidence supports both EMH (fast reaction to new information) and Behavioral Finance (asymmetric and emotional reactions). Signaling and Stakeholder theories add more depth by showing that credibility, legitimacy, and social values also play a role. Together, these perspectives highlight that ESG finance is shaped by both financial logic and human factors.

5.5 Limitations and Directions for Future Research:

Although this study makes important contributions, some limitations should be noted. First, the dataset is based mainly on European news and announcements. This focus may limit generalizability, since ESG practices and market reactions differ in other regions such as Asia or Latin America. Future research could include more diverse markets and make cross-country comparisons.

Second, the sentiment analysis in this study used a simpler method, which may not fully capture complex expressions such as irony or sarcasm. Future work could apply more advanced NLP models, such as FinBERT or ensemble approaches, and combine them with human validation to increase accuracy.

Third, the event window used to calculate CAR was -3 to +3 trading days (7 days). This is suitable for identifying short-term market reactions, but it does not capture long-term effects, such as changes in investor trust or capital allocation. Future studies could extend the time horizon and examine whether ESG effects persist over weeks or months.

Another limitation is the firm classification. Companies were divided into only two groups: “Renewable Energy” and “Oil & Gas.” While this provides analytical clarity, many firms operate across different segments or are in transition toward greener strategies. Future research could use dynamic classifications or ESG transition scores to reflect this complexity.

Finally, although machine learning improved prediction, the models treated events as independent and did not fully account for larger shocks, such as regulatory changes or global crises. More advanced time-series or regime-switching models could better capture these dynamics.

In conclusion, future research should expand datasets across regions, use stronger sentiment analysis tools, and extend event windows to study longer-term effects. These steps would provide a more complete picture of how ESG news shapes financial markets

Chapter 6:

Conclusion and Recommendation

6.1 Summary of Findings

This research investigated how financial markets respond to environmental, social, and governance (ESG) news, with a focus on the role of sentiment, company type (Oil & Gas vs. Renewable Energy), and firm-specific characteristics. By combining event study methodology, machine learning models, and sentiment-based text analytics, the study yielded the following key findings:

- ESG-related news events resulted in statistically significant but modest cumulative abnormal returns (CARs). The results align with previous literature, suggesting that sustainability information is increasingly priced by financial markets.
- Market reactions were highly heterogeneous. CARs were not uniformly distributed. While most companies experienced small or negligible reactions, a subset especially Oil & Gas firms—exhibited significantly stronger responses. This heterogeneity was more accurately captured using quantile regression and non-linear machine learning models, rather than traditional OLS approaches.
- Sentiment matters, but in an asymmetric fashion. Although the emotional tone of news did not have strong effects in linear OLS models, it showed meaningful asymmetric effects in quantile models. Positive sentiment helped mitigate negative returns in the lower quantiles but had limited or even inverse effects in the upper quantiles.
- Firm size influenced the magnitude of response. Smaller firms showed greater sensitivity to ESG-related news, likely due to lower levels of analyst coverage and reduced information transparency compared to larger firms.
- Company type plays a critical role. Oil & Gas companies reacted more strongly than Renewable Energy firms, particularly when the news involved transition risks or regulatory exposure. This highlights that a firm's operating context shapes how ESG news is interpreted by the market.
- Sentiment-based and machine learning models added significant analytical value. Techniques such as XGBoost and Random Forest outperformed linear models in capturing the complexity and non-linearity of market behavior in response to ESG events.

In summary, the findings confirm that financial markets do respond to ESG-related disclosures, but the intensity and nature of this response depend heavily on company type, news content, sentiment tone, firm size, and the analytical method used to evaluate it.

6.2 Contributions of the Study

6.2.1 Academic Contributions

Event-Level Analysis in a European Context

The study introduces a firm-level dataset of 450 environmental announcements from European area less listed firms, offering short-term, time-specific analysis of stock price reactions explored in the European setting.

Integration of Sentiment-Based Textual Features

By applying sentiment scoring to the content of environmental news, the research adds nuance to event study methodology, allowing for the incorporation of tone and language features into financial models of market reaction.

Advanced Methodological Framework

The study combines quantile regression and machine learning (XGBoost, Random Forest) to capture heterogeneous and nonlinear effects that traditional models may overlook. This methodological integration enhances the analysis of ESG-related events

6.2.2 Practical Contributions

This study provides several practical contributions that are relevant to corporate managers, financial analysts, investors, and policy-makers operating within ESG-conscious markets, particularly in the European context:

- **Enhanced Market Awareness of ESG News Impact:**

The findings demonstrate that environmental news, particularly with strong sentiment or related to transition, can lead to statistically significant abnormal returns. This underscores the importance for corporate IR teams and compliance officers to closely monitor the timing, framing, and content of ESG-related announcements to reduce potential volatility and reputational risk.

- **Sentiment-Informed Investment Strategies:**

By showing that market reactions to environmental announcements are asymmetric and depend on sentiment tone, the study provides evidence that can be used by institutional

investors and portfolio managers to design sentiment-aware trading strategies. Positive sentiment may mitigate downside risk, particularly for small-cap or high-risk firms.

- **Risk Differentiation by Sector and Company Type:**

The observed stronger reactions among Oil & Gas companies and small-cap firms highlight the need for tailored risk assessment tools. Risk officers and ESG consultants can apply these findings to refine risk modeling frameworks based on sector-specific ESG exposure and market capitalization.

- **Guidance for ESG Communication Strategy:**

Since tone and topic materially affect investor response, communication teams can use sentiment analysis tools before public releases to evaluate potential market impact. This can support decision-making regarding the disclosure of sensitive or strategically timed ESG information

6.3 Policy and Managerial Recommendations

6.3.1 ESG Disclosure Guidelines

The results of this study indicate that not only the content but also the **tone**, **timing**, and **type** of environmental news significantly affect market reactions. As such, there is a pressing need for standardized yet flexible ESG disclosure guidelines. The following recommendations can help firms and regulators align ESG communication with investor expectations:

- **Standardization of Event-Level ESG Disclosure:**

Regulators should encourage companies to report environmental incidents and initiatives using structured templates that include date, category (e.g., regulatory change, sustainability initiative), and sentiment indicators when applicable.

- **Timely Disclosure of Material Environmental Information:**

Companies should ensure that ESG-related events are disclosed promptly, minimizing information asymmetry and reducing the risk of market overreactions or speculation.

- **Inclusion of Sentiment Assessment Tools:**

Incorporating sentiment scoring (e.g., positive/negative tone) within ESG disclosures can enhance transparency and help investors better interpret qualitative aspects of reports.

- **Sector-Specific Disclosure Requirements:**

Guidelines should be tailored based on industry sensitivity to environmental risks. For instance, Oil & Gas firms should report in more detail on transition risks, while Renewable Energy firms may focus on innovation and compliance

6.3.2 Corporate Communication Strategy

Given that market responses vary by firm size and sector, companies should consider refining their ESG communication practices. Smaller firms, in particular, may benefit from providing additional context when announcing environmental initiatives. Using clear and consistent language, avoiding exaggerated claims, and aligning the message with stakeholder expectations can help reduce market uncertainty and promote trust

6.4 Suggestions for Future Research

This study offers a base for understanding how environmental news affects stock performance in the European energy sector, but there is still room for future work. One direction is to classify ESG events in more detail, moving beyond broad labels like “regulation” or “pollution” to categories such as “climate litigation,” “supply chain emissions,” or “green bond issuance.” Another area is cross-regional comparison, since this research focused only on Europe; studies in North America, East Asia, or emerging markets could show how different regulations and investor behaviors change market reactions. Finally, while machine learning models like XGBoost predict well, they are hard to interpret. Future research could use explainable AI methods, such as interpretable neural networks, to better show which ESG signals have the strongest impact on investor decisions.

References

- Aouadi, A., & Marsat, S. (2018). Do ESG controversies matter for firm value? Evidence from international data. *Journal of Business Ethics*, 151(4), 1027–1047.
- Araci, D. (2019). FinBERT: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063*.
- Ball, R. (1994). The development, accomplishments and limitations of the theory of stock market efficiency. *Managerial Finance*, 20(2–3), 3–48.
- Barberis, N., & Thaler, R. (2003). A survey of behavioral finance. In *Handbook of the Economics of Finance* (Vol. 1, pp. 1053–1128). Elsevier.
- Berg, F., Koelbel, J. F., & Rigobon, R. (2022). Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, 26(6), 1315–1344.
- Binns, R., Veale, M., Van Kleek, M., & Shadbolt, N. (2018). ‘It’s reducing a human being to a percentage’: Perceptions of justice in algorithmic decisions. In *CHI Conference on Human Factors in Computing Systems* (pp. 1–14). ACM.
- Boubaker, S., Cellier, A., & Manita, R. (2022). Tone and source credibility in ESG news: Market implications. *Finance Research Letters*, 46, 102309.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- Broadstock, D. C., Chan, K., Cheng, L. T. W., & Wang, X. (2021). The role of ESG performance during times of financial crisis: Evidence from COVID-19 in China. *Finance Research Letters*, 38, 101716.
- Brown, S. J., & Warner, J. B. (1985). Using daily stock returns: The case of event studies. *Journal of Financial Economics*, 14(1), 3–31.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems* (Vol. 33, pp. 1877–1901).
- Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (1997). *The econometrics of financial markets*. Princeton University Press.
- Capelle-Blancard, G., & Petit, A. (2019). Every little helps? ESG news and stock market reaction. *Journal of Business Ethics*, 157(2), 543–565.

- Cheema-Fox, A., LaPerla, B. R., Serafeim, G., & Wang, H. S. (2021). Corporate ESG news and stock market reactions. *Financial Analysts Journal*, 77(3), 55–76.
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794).
- Christensen, D. M., Serafeim, G., & Sikochi, A. (2022). Why is corporate virtue in the eye of the beholder? The case of ESG ratings. *Accounting Review*, 97(1), 147–175.
- Clark, G. L., Feiner, A., & Viehs, M. (2015). From the stockholder to the stakeholder: How sustainability can drive financial outperformance. University of Oxford and Arabesque Partners Report.
- Corrado, C. J. (2011). Event studies: A methodology review. *Accounting & Finance*, 51(1), 207–234.
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dyck, A., Lins, K. V., Roth, L., & Wagner, H. F. (2019). Do institutional investors drive corporate social responsibility? International evidence. *Journal of Financial Economics*, 131(3), 693–714.
- ETF.com. (2020). Selecting size & style benchmarks. Retrieved from https://www.etf.com/docs/ETF.com_SelectingSizeStyleBenchmarks.pdf
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance*, 47(2), 427–465.
- Fatemi, A., Glaum, M., & Kaiser, S. (2018). ESG performance and firm value: The moderating role of disclosure. *Global Finance Journal*, 38, 45–64.
- Flammer, C. (2013). Corporate social responsibility and shareholder reaction: The environmental awareness of investors. *Academy of Management Journal*, 56(3), 758–781.
- Flammer, C. (2021). Corporate green bonds. *Journal of Financial Economics*, 142(2), 499–516.

- Friede, G., Busch, T., & Bassen, A. (2015). ESG and financial performance: Aggregated evidence from more than 2000 empirical studies. *Journal of Sustainable Finance & Investment*, 5(4), 210–233.
- García-Meca, E., Martínez-Ferrero, J., & Sánchez, I. M. (2021). How does tone in ESG reporting affect financial markets? *Sustainability*, 13(2), 712.
- Ghosal, A., Jain, A., & Mukherjee, S. (2023). Outperforming analysts: Large language models in financial forecasting. *Journal of Financial Data Science*, 5(2), 1–18.
- Ghosal, D., Patel, H., & Rajan, R. (2023). ESG sentiment analysis using large language models: Predictive insights and market behavior. *Journal of Sustainable Finance & Investment*, 13(3), 455–472.
- Ghosal, S., Zhang, J., & Yoon, S. (2023). Detecting ESG greenwashing with deep learning: A framing-aware approach. *Journal of Financial Data Science*, 5(1), 43–61.
- Görgen, M., Jacob, A., Nerlinger, M., Riordan, R., Rohleder, M., & Wilkens, M. (2020). Carbon pricing and stock returns: Evidence from the EU Emissions Trading Scheme. *Journal of Banking & Finance*, 118, 105874.
- Grewal, J., Hauptmann, C., & Serafeim, G. (2020). Material sustainability information and stock price informativeness. *Journal of Business Ethics*, 154(4), 1131–1161.
- Griffin, P. A., Lont, D. H., & Sun, Y. (2017). The relevance to investors of greenhouse gas emission disclosures. *Contemporary Accounting Research*, 34(2), 1265–1297.
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5), 2223–2273.
- Ioannou, I., & Serafeim, G. (2015). The impact of corporate social responsibility on investment recommendations. *Financial Analysts Journal*, 71(4), 1–23.
- Jegadeesh, N., & Wu, D. (2013). Word power: A new approach for content analysis. *Journal of Financial Economics*, 110(3), 712–729.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.
- Klassen, R. D., & McLaughlin, C. P. (1996). The impact of environmental management on firm financial performance. *Management Science*, 42(8), 1199–1214.
- Kogan, L., Levin, D., Routledge, B. R., Sagi, J. S., & Smith, N. A. (2009). Predicting risk from financial reports with regression. In *Proceedings of NAACL HLT 2009* (pp. 272–280).

- Kölbel, J. F., Heeb, F., Paetzold, F., & Busch, T. (2020). Can sustainable investing save the world? Reviewing the mechanisms of investor impact. *Organization & Environment*, 33(4), 554–574.
- Krueger, P., Sautner, Z., & Starks, L. T. (2020). The importance of climate risks for institutional investors. *Review of Financial Studies*, 33(3), 1067–1111.
- Krüger, P. (2015). Corporate goodness and shareholder wealth. *Journal of Financial Economics*, 115(2), 304–329.
- Li, F. (2010). The information content of forward-looking statements in corporate filings: A naïve Bayesian machine learning approach. *Journal of Accounting Research*, 48(5), 1049–1102.
- Lins, K. V., Servaes, H., & Tamayo, A. (2017). Social capital, trust, and firm performance: The value of corporate social responsibility during the financial crisis. *Journal of Finance*, 72(4), 1785–1824.
- Liu, X., Zhang, Y., & Wang, H. (2022). The role of tone and certainty in climate-related financial disclosures. *Journal of Financial Economics*, 145(3), 977–1003.
- Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66(1), 35–65.
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems* (Vol. 30).
- Luo, X., Wang, H., Raithel, S., & Zheng, Q. (2020). Corporate social performance, analyst stock recommendations, and firm future returns. *Journal of Marketing*, 84(4), 40–60.
- MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature*, 35(1), 13–39.
- MSCI. (2023). MSCI Global Investable Market Indexes methodology. MSCI Research.
- Müller, M., Schaefer, A., & Köhler, T. (2021). Multilingual challenges in ESG sentiment analysis: A cross-lingual benchmark. In *Proceedings of the European Conference on Natural Language Processing*.
- Raji, I. D., Smart, A., White, R., Mitchell, M., Gebru, T., Hutchinson, B., ... & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 33–44). ACM.

- Sautner, Z., van der Ploeg, F., & Vilkov, G. (2020). Firm-level climate change exposure. *Review of Financial Studies*, 33(3), 1231–1266.
- Siew, R. Y. J., Balatbat, M. C. A., & Carmichael, D. G. (2016). The relationship between sustainability practices and financial performance. *Construction Management and Economics*, 34(4–5), 339–361.
- Sokolov, A., de Vries, T., & van Beusichem, H. (2021). The language of sustainability: How tone in ESG reports influences investor response. *Sustainability Accounting, Management and Policy Journal*, 12(3), 571–594.
- Sokolov, V., Tsoy, D., & Dobrynin, A. (2021). Market reaction to the tone of ESG disclosures: Evidence from European firms. *Finance Research Letters*, 41, 101895.
- Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3), 355–374.
- Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance*, 62(3), 1139–1168.
- Wang, Y., Li, Q., & Wang, S. (2021). ESG news analysis using deep learning: Investor reactions in capital markets. *IEEE Access*, 9, 114237–114248.

Appendices

ESG (Environmental, Social, Governance):

A framework used to assess a company's performance and risks beyond traditional financial metrics.

It includes:

- **Environmental** criteria (e.g., carbon emissions, energy use),
- **Social** factors (e.g., labor practices, community engagement),
- **Governance** structures (e.g., executive pay, board diversity).

ESG has become a major concern for investors seeking responsible and sustainable investment strategies.

Cumulative Abnormal Return (CAR):

A financial metric used in event studies to measure the total abnormal return generated by a stock over a specific time window surrounding a particular event (e.g., earnings announcement, ESG disclosure).

It compares the actual return of a stock to its expected return based on a market model. A significant CAR suggests that the event had an impact on investor behavior.

Natural Language Processing (NLP):

A subfield of artificial intelligence (AI) focused on enabling machines to understand, interpret, and generate human language.

In finance, NLP is used to analyze textual data such as news articles, earnings reports, or social media posts to extract sentiment, classify content, and predict market behavior.

Event Study:

A quantitative method used in financial economics to evaluate the impact of a specific event on a firm's stock price.

The event study measures whether there is a statistically significant reaction in stock returns during a predefined window around the event date. It is widely used for assessing the market impact of corporate announcements, policy changes, or ESG news.

Tone Analysis:

A technique used in textual analysis to determine the emotional or attitudinal tone expressed in written communication.

In financial contexts, it is often used to classify news as positive, negative, or neutral, which can help predict investor sentiment and stock market movements.

Efficient Market Hypothesis (EMH):

A theory in financial economics stating that asset prices fully reflect all available information at any given time.

According to EMH, it is impossible to consistently achieve higher returns than the overall market without assuming additional risk, because markets respond instantly and rationally to new information.

Text Mining:

The process of transforming unstructured textual data into structured formats to identify patterns, trends, or insights.

It involves steps such as tokenization, keyword extraction, sentiment classification, and topic modeling. In ESG research, text mining helps identify how different issues are framed in media or reports.

Sentiment Analysis:

A form of NLP focused on extracting subjective information from text, such as emotions, opinions, or attitudes.

It classifies content as positive, negative, or neutral. In financial markets, sentiment analysis is used to interpret the tone of news, tweets, or financial reports to anticipate market reactions.

Firm-Specific News:

Information or events that are directly related to a single company—such as earnings releases, product launches, management changes, or ESG disclosures. Unlike macroeconomic or industry-wide news, firm-specific news often has a targeted impact on that company's stock price.

Information Asymmetry:

A condition in which one party in a transaction has access to more or better information than the other.

In capital markets, it can lead to market inefficiencies, insider trading, or mispricing. ESG disclosures aim to reduce such asymmetries by making relevant non-financial information publicly available.

Event Window:

A defined period around the event date (e.g., [-1, +1] days) in which abnormal returns are measured to evaluate the event's effect.

Market Model:

A linear regression model used to estimate expected returns of a stock based on market returns. Typically:

$$R_{it} = \alpha + \beta R_{mt} + \epsilon_{it}$$

NLP-Based Classification:

The use of Natural Language Processing techniques to automatically categorize textual data (e.g., ESG news) into predefined topics or sentiments based on linguistic patterns.

CAR (Cumulative Abnormal Return):

The aggregate sum of abnormal returns over the event window, representing the total market reaction to an event.

Machine Learning Models (e.g., XGBoost, Random Forest):

Algorithms used to build predictive models for classification or regression tasks. In this context, they are employed to predict news category or tone based on textual features.

Confusion Matrix:

A table used to evaluate the performance of classification models by comparing predicted and actual categories, showing true/false positives and negatives.

Accuracy / Precision / Recall / F1 Score:

Metrics used to assess the performance of classification models:

- **Accuracy:** Proportion of correct predictions.
- **Precision:** Correct positive predictions over total predicted positives.
- **Recall:** Correct positive predictions over actual positives.
- **F1 Score:** Harmonic mean of precision and recall.

Vectorization (e.g., TF-IDF):

A technique for converting text into numerical features that can be fed into machine learning algorithms.

TF-IDF (Term Frequency-Inverse Document Frequency) gives weight to terms that are more unique to a document

Descriptive Statistics:

A set of statistical measures (mean, median, standard deviation, etc.) used to summarize and describe the main features of a dataset, particularly CAR values in this study.

Boxplot:

A graphical representation of data distribution that shows the median, quartiles, and potential outliers. It is used in this chapter to compare CAR across company types and news categories.

ANOVA (Analysis of Variance):

A statistical technique used to determine whether there are statistically significant differences between the means of three or more independent groups. Here, it's used to assess whether CAR varies significantly across news categories.

t-test:

A statistical test used to compare the means of two groups when the variances are unequal. In this context, it is applied to evaluate differences in CAR between Oil & Gas and Renewable companies.

Significance Level (p-value):

A measure of the probability that the observed results occurred by chance. Lower p-values (e.g., < 0.05) indicate stronger evidence against the null hypothesis.

News Category Classification:

The categorization of news articles into distinct ESG themes such as environment, governance, or social, often determined using NLP or manual tagging.

Company Type Classification:

The grouping of firms into categories such as Oil & Gas or Renewable Energy based on their sector, for the purpose of comparing market responses to ESG news.

CAR Comparison Across Groups:

The analysis of differences in cumulative abnormal returns between predefined categories, such as company types or news topics, to interpret market sensitivity.

Violin Plot (if used):

A visualization that combines a boxplot and a density plot, showing data distribution along with median and interquartile ranges. Useful for comparing distributions of CAR.

Market Reaction Interpretation:

An analysis of whether and how the market responded to various ESG-related events, based on statistical significance and direction of CAR.

Managerial Implications

Practical recommendations derived from research findings intended to guide business leaders, corporate communication teams, or sustainability officers in decision-making and strategic planning.

Policy Recommendations:

Insights from the study that can inform policy makers, regulators, or ESG rating agencies in the development of frameworks, guidelines, or reporting standards that encourage transparency and accountability.

Investor Behavior:

The patterns, preferences, and reactions of investors in response to corporate disclosures, particularly in ESG contexts. It includes aspects like risk perception, sentiment, and trading decisions.

Strategic Communication:

A coordinated and deliberate approach to conveying corporate messages—especially ESG-related content—to influence stakeholder perception and market behavior. Emphasis is placed on tone, timing, and channel of communication.

Asymmetric Market Response:

The phenomenon where positive and negative ESG news do not generate equal reactions in magnitude or direction in the stock market, often due to investor biases or contextual factors.

Theory-Practice Gap:

A common issue in research where theoretical models or academic findings do not easily translate into actionable strategies or are underutilized in real-world decision-making.

Limitations of the Study:

Acknowledged constraints in the research design, data availability, methodology, or scope that may affect the generalizability or interpretation of the findings.

Future Research Directions:

Proposed areas for further investigation that emerge from the study's results or limitations, aimed at expanding the body of knowledge or improving upon the current methodology.

Cross-Sector Analysis:

Comparative analysis of different industries or company types (e.g., Oil & Gas vs. Renewables) to understand how ESG dynamics vary across economic sectors.

Sustainability Reporting:

The practice of disclosing information on environmental, social, and governance performance. It is evaluated based on standards such as GRI, SASB, or TCFD, and was a key consideration in interpreting the implications of ESG news.

Integrated Reporting (IR):

A holistic approach to corporate reporting that combines financial and non-financial (e.g., ESG) data into a single, comprehensive report. It aims to provide stakeholders with a fuller picture of a company's long-term value creation.

Impact Measurement:

The process of assessing the social, environmental, or economic outcomes of corporate actions, policies, or disclosures. In the ESG context, this refers to quantifying how firm behavior influences market perceptions and stakeholder value.

Behavioral Finance Perspective:

An interpretation of market reactions based on psychological factors and investor biases, challenging the assumption of full rationality in the Efficient Market Hypothesis (EMH).

Dynamic Market Context:

The evolving nature of financial markets influenced by real-time news, technological advancement (e.g., NLP tools), and increasing stakeholder focus on ESG. It emphasizes that interpretations and models must be adaptable.

Practical Relevance:

The degree to which academic findings translate into actionable strategies for businesses, regulators, or investors. Emphasizing this helps bridge the gap between research and practice.

Stakeholder-Centric Communication:

An approach where corporate disclosures and ESG messaging are designed considering the informational needs and expectations of various stakeholders—investors, regulators, the public, etc.

Transparency and Accountability:

Core ESG principles advocating for open disclosure of firm practices and outcomes. These are essential for building trust and reducing information asymmetry in capital markets.

Sustainability as a Strategic Asset:

The concept that integrating ESG into corporate strategy is not merely about compliance, but a source of competitive advantage and long-term performance enhancement.

Research-to-Practice Translation:

The process of applying academic research insights (like ESG event studies or NLP findings) in real-world settings, such as investment strategy design or corporate communication planning.

Sintesi

Questo studio ha esaminato in dettaglio le modalità con cui i mercati finanziari reagiscono alle notizie aziendali relative ai criteri ambientali, sociali e di governance (ESG), con un focus particolare sulla polarità del sentiment, sulla categoria delle notizie e sulla tipologia di impresa. Il quadro teorico e metodologico si colloca all'incrocio tra la letteratura sugli studi di evento e le tecnologie analitiche più recenti, come l'apprendimento automatico e l'analisi semantica, con l'obiettivo di offrire una visione più dinamica e realistica delle reazioni di mercato alle disclosure ESG.

Alla base della ricerca vi è il riconoscimento crescente dell'importanza dei fattori ESG nei processi decisionali di investitori, regolatori e dirigenti aziendali. In un contesto caratterizzato da pressioni ambientali globali, aspettative normative in evoluzione e maggiore consapevolezza sociale, le imprese sono sempre più chiamate a comunicare in modo trasparente e tempestivo le proprie azioni e strategie di sostenibilità. Tuttavia, la letteratura accademica offre ancora risultati contrastanti sull'effettiva capacità delle informazioni ESG di influenzare i mercati finanziari in modo consistente.

Lo studio ha utilizzato un ampio dataset composto da oltre 270.000 osservazioni giornaliere, che associano notizie ESG a rendimenti azionari a livello di impresa. Ogni evento è stato sottoposto a scoring di sentiment, classificazione tematica e tipizzazione per settore e dimensione dell'azienda. Il disegno metodologico ha integrato l'analisi degli eventi con modelli statistici avanzati (regressione quantifica) e algoritmi di machine learning (XGBoost, Random Forrest), in modo da rilevare effetti eterogenei e non lineari.

Uno dei risultati principali emersi riguarda l'asimmetria delle reazioni. Le notizie ESG negative – soprattutto quelle legate a controversie ambientali, violazioni normative o disastri industriali – hanno prodotto effetti molto più forti e persistenti rispetto alle notizie positive. Questo fenomeno è stato particolarmente evidente nei quantili inferiori della distribuzione dei rendimenti anomali cumulati (CAR), coerente con le teorie comportamentali come la "Loss aversion" e la "negativity bias". Le imprese del settore Oil & Gas sono risultate particolarmente vulnerabili a questo tipo di eventi, mostrando reazioni più intense rispetto alle aziende del settore Energie Rinnovabili.

Anche la categoria tematica della notizia ha dimostrato di avere un impatto differenziato. Le notizie normative, ad esempio, tendono a generare reazioni più negative, mentre le notizie strategiche o relative a innovazioni green sono state accolte in modo più neutro o positivo. Questo

indica che non tutte le disclosure ESG sono interpretate allo stesso modo, e che la narrativa e il framing della notizia giocano un ruolo centrale nella formazione delle aspettative di mercato.

Dal punto di vista della dimensione aziendale, le piccole imprese hanno mostrato maggiore sensibilità alle notizie ESG. Questo può essere attribuito alla minore visibilità sul mercato, alla carenza di copertura analitica e alla maggiore incertezza informativa. Le aziende più grandi o consolidate tendono invece ad ammortizzare meglio gli effetti reputazionali, potenzialmente grazie a una governance più strutturata e a strategie di comunicazione più sofisticate.

Lo studio ha anche evidenziato differenze territoriali potenziali. Sebbene l'analisi si sia focalizzata sul mercato europeo, la letteratura suggerisce che le reazioni a notizie ESG possono variare in funzione del contesto geografico, delle norme di disclosure locali e della cultura finanziaria. Ad esempio, nei mercati emergenti, la sensibilità degli investitori può essere inferiore a causa di una minore pressione regolatoria o della prevalenza di strategie speculative di breve termine.

Dal punto di vista metodologico, l'utilizzo combinato di regressione quantilica e machine learning ha rappresentato un importante avanzamento. I modelli tradizionali come l'OLS tendono a sottovalutare gli effetti distribuzionali e non sono adatti a catturare l'eterogeneità degli impatti. Al contrario, modelli ensemble come XGBoost hanno permesso di esplorare le interazioni complesse tra sentiment, tipologia di notizia, settore e dimensione aziendale, offrendo previsioni più accurate e significative. Tuttavia, la scarsa trasparenza di questi modelli rappresenta una sfida: futuri studi potrebbero avvalersi di tecniche di Explainable AI per migliorare l'interpretabilità.

Le implicazioni pratiche sono rilevanti su più livelli. Per gli investitori istituzionali, le evidenze suggeriscono che l'integrazione del sentiment e del contesto ESG nelle strategie di investimento può migliorare la capacità di anticipare volatilità e correggere asimmetrie informative. Per i manager aziendali, lo studio sottolinea l'importanza di pianificare la comunicazione ESG in modo strategico, valutando il tono, il momento e il contenuto delle disclosure. Per i policy-maker, i risultati rafforzano la necessità di sviluppare linee guida di reporting ESG che siano coerenti, comparabili e dotate di componenti qualitative come indicatori di sentiment.

Inoltre, il lavoro evidenzia la necessità di standardizzare la classificazione delle notizie ESG. Attualmente, molte comunicazioni aziendali mancano di struttura e chiarezza, rendendo difficile per il mercato valutarne l'importanza. L'introduzione di tassonomie comuni e di template di disclosure potrebbe migliorare l'efficienza informativa e ridurre la volatilità legata a interpretazioni soggettive.

Tra le raccomandazioni strategiche emergono:

L'inserimento di sistemi di alert basati su sentiment analysis nei terminali di trading;

L'integrazione di indicatori ESG nei rating creditizi e nei modelli di rischio operativo;

L'elaborazione di piani di comunicazione proattiva in situazioni di crisi ambientali o sociali.

Infine, lo studio apre numerose prospettive per future ricerche. Alcune direzioni promettenti includono: l'estensione dell'analisi a mercati extra-europei; l'utilizzo di dati alternativi (come social media, podcast, o ESG ratings di terze parti); l'applicazione di modelli di apprendimento profondo per l'analisi semantica di testi complessi; e la valutazione dell'impatto a lungo termine delle strategie ESG sulla performance finanziaria, sul costo del capitale e sulla resilienza aziendale.

In sintesi, questa ricerca contribuisce in modo significativo alla letteratura empirica sull'ESG, offrendo evidenze solide, metodi innovativi e spunti per l'azione concreta. Essa conferma che i mercati reagiscono alle notizie ESG in modo differenziato e condizionato, e che fattori come la tonalità emotiva, il tipo di impresa e la dimensione informativa sono centrali per comprendere queste dinamiche. In un'epoca in cui la sostenibilità è diventata parte integrante della finanza moderna, l'integrazione di strumenti avanzati di analisi – quantitativi e qualitativi – rappresenta una strada imprescindibile per il progresso scientifico e per una governance responsabile dei mercati.