# FROM CLAIMS TO REALITY: A DATA-DRIVEN APPROACH TO MEASURING GREENWASHING WITH LARGE LANGUAGE MODELS

by

Lida Ghasemi

A thesis

presented to the University of Waterloo

in fulfillment of the

thesis requirement for the degree of

Master of Applied Science

in

Management Science and Engineering

Waterloo, Ontario, Canada, 2025

© Lida Ghasemi 2025

## Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

# Abstract

The oil and gas industry, a major contributor to climate change, faces increasing scrutiny from regulators, investors, and the public to adopt sustainable practices. Many companies respond with ambitious environmental commitments yet concerns about greenwashing, raise doubts about the authenticity of these claims. Greenwashing allows firms to maintain a sustainable public image while continuing high-emission activities, creating a gap between corporate rhetoric and actual environmental action.

This study introduces a data-driven approach to measure greenwashing by analyzing corporate disclosures from 50 publicly traded oil and gas firms (2011–2020). Using Large Language Models (LLMs) and Natural Language Processing (NLP), we assess sustainability reports, annual filings, and 10-K reports, developing a quantitative greenwashing index to distinguish between symbolic and substantive environmental claims.

Findings reveal systematic discrepancies between voluntary sustainability reports and legally binding financial disclosures. Larger firms and those in highly regulated markets engage in more sophisticated greenwashing, and the practice has increased over time, particularly after the 2015 Paris Agreement. This research highlights the need for stricter regulations, standardized reporting, and independent audits to ensure corporate sustainability commitments lead to real environmental action.

# Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisors, Dr. Yang and Dr. Dimitrov. Dr. Yang, your unwavering support and encouragement have been instrumental in helping me explore innovative ideas and push the boundaries of my research. Your guidance has been invaluable, and I am truly grateful for your mentorship. Dr. Dimitrov, your meticulous attention to detail and insightful feedback have significantly shaped the quality of my work. Your constructive comments and valuable perspectives have strengthened my research, and I sincerely appreciate your dedication to my academic growth.

I would also like to thank Ahmadreza Jeddi, whose generosity and guidance played a crucial role in supporting my research. His help and encouragement were instrumental in overcoming technical challenges and refining my ideas.

To all who have supported me throughout this process, family, friends, and mentors, I extend my heartfelt thanks. This achievement would not have been possible without your kindness and encouragement.

# Dedication

This thesis is dedicated, above all, to the memory of my beloved parents, who always encouraged me to persevere in my studies. Their unwavering belief in the power of education continues to inspire me, and I hope this achievement makes them proud.

I also dedicate this work to my husband, Kazem Jeddi, whose constant support and encouragement have been the foundation of my journey. His belief in me, his patience, and his unwavering presence made this accomplishment possible.

Lastly, I dedicate this thesis to my daughter, Elara. Throughout my research, she was with me, both before her birth and after. Carrying her while carrying out my studies made this journey even more meaningful, and she remains my greatest source of motivation and joy.

Table of Contents

Author's Declaration ..... ii
Abstract ..... iii
Acknowledgements ..... iv
Dedication ..... v
List of Figures ..... viii
List of Tables ..... x
List of abbreviations ..... xi
1. Introduction ..... 1
2. Related Works ..... 4
2.1. Assessing Environmental Challenges and Corporate Responses in the Oil and Gas Sector ..... 4
2.2. Transparency in Oil and Gas Companies ..... 6
2.3. Greenwashing: Definitions, Measurement, and Limitations in Prior Research ..... 7
2.4. Investigating greenwashing using Text data analysis and language models ..... 8
3. Overview of the Data ..... 10
3.1. The Companies ..... 11
3.2. Market Data Features of the Oil and Gas Industry Companies (2011-2020) ..... 12
3.3. Overview of Datasets from Annual, 10-K, and Sustainability Reports ..... 13
3.3.1. Annual Reports ..... 14
3.3.2. 10-K Filings ..... 15
3.3.3. Sustainability Reports ..... 16
3.3.4. Summary ..... 17
4. Greenwashing ..... 18
4.1. Definition of Greenwashing ..... 18
4.2. Identification of Claims and Actions ..... 20
4.3. LLM-Based Classification and Model Justification ..... 20
4.4. Model Performance Evaluation for Llama 3-70B ..... 24

4.5. Justification and Methodological Considerations ..... 26
4.6. Experimental Setup Overview ..... 27
4.7. Ablation Studies ..... 29
5. Policy Analysis and Conclusion ..... 36
5.1. Experimental Setup ..... 36
5.1.1. Data Preparation ..... 37
5.1.2. Model Training and Validation ..... 37
5.1.3. Feature Importance Analysis ..... 38
5.2. Experimental Results ..... 39
5.2.1. Sustainability Dataset ..... 40
5.2.2 Annual Dataset ..... 42
5.2.3. 10K Dataset ..... 44
5.2.4. The Common Patterns ..... 46
5.3. Conclusion ..... 47
5.4. Implications of Findings and Policy Recommendations ..... 49
References ..... 51
Appendix A: Market Data Feature Descriptions ..... 56
Appendix B: Full Set of Partial Dependence Plots (PDPs) ..... 60

# List of Figures

Figure 4.1. Prompt structure provided to LLaMA 3-70B for classifying environmental claims and actions. The prompt includes explicit definitions, examples, and output formatting instructions aligned with manual annotation criteria. This structure enables the model to reason contextually and consistently across various corporate disclosures.…..……..22
Figure 4.2. Application of LLaMA 3-70B on a sample 10-K report page from ExxonMobil (2020). The model identifies and classifies textual segments as either symbolic claims or substantive actions, demonstrating its capacity to parse complex corporate language using prompt-based inference.…….……..23
Figure 4.3. Overview of the trends in four metrics, number of pages with environmental mentions, number of symbolic claims, number of substantive actions, and the greenwashing index, over the decade 2011–2020. The mean value of each metric is reported annually for each dataset……..32
Figure 4.4. Overview of the trend in the greenwashing index for oil and gas companies, categorized by the continent of their headquarters. The reported values represent the mean greenwashing index for all companies within each category…….….33
Figure 5.1. Training and validation loss curves (left) for the Sustainability dataset, showing the average MAE loss over five random seed partitions. Corresponding feature importances are shown on the right...…….40
Figure 5.2. SHAP values for the 13 independent features and the corresponding dependent feature, greenwashing, on the Sustainability dataset……..…….42
Figure 5.3. Training and Validation Loss Curves (left) for the Annual dataset, representing the average MAE loss value over the 5 random seed partitions., Corresponding Feature Importances (right).. ..... 42
Figure 5.4. SHAP values for the 13 independent features and the corresponding dependent feature, greenwashing, on the Annual dataset.…….…...44
Figure 5.5. . Training and Validation Loss Curves (left) for the 10K dataset, representing the average MAE loss value over the 5 random seed partitions. Corresponding Feature Importances (right)…….…….…….45
Figure 5.6. SHAP values for the 13 independent features and the corresponding dependent feature, greenwashing, on the 10K dataset…….……..46
Figure B.1. Partial Dependence Plots for the 13 independent features and the corresponding dependent feature, greenwashing, on the Sustainability dataset ……..……...60

Figure B.2. Partial Dependence Plots for the 13 independent features and the corresponding dependent feature, greenwashing, on the Annual dataset ……...61
Figure B.3. Partial Dependence Plots for the 13 independent features and the corresponding dependent feature, greenwashing, on the 10K dataset.…….……62

# List of Tables

Table 4.1. Llama 3-70B Model Performance Summary……..25
Table 4.2. Categories of Environmental Terms used for pre-screening content before feeding them to the LLM ..... 29
Table 4.3. Greenwashing statistics on the whole dataset. ..... 31
Table 4.4. Comparison between greenwashing in companies with sustainability reports and companies without sustainability report. ..... 31
Table B.1. Market data features, along with a description of what each feature represents…59

# List of abbreviations

| Abbreviation | Full Term |
|---|---|
| **LLM** | Large Language Model |
| **NLP** | Natural Language Processing |
| **ESG** | Environmental, Social, and Governance |
| **PDP** | Partial Dependence Plot |
| **SHAP** | SHapley Additive exPlanations |
| **SEC** | U.S. Securities and Exchange Commission |
| **MAE** | Mean Absolute Error |
| **GHG** | Greenhouse Gas |
| **CSR** | Corporate Social Responsibility |
| **MD&A** | Management Discussion and Analysis |
| **GI** | Greenwashing Index |
| **10-K** | U.S. Securities and Exchange Commission Annual Filing |
| **D/E Ratio** | Debt-to-Equity Ratio |

# Chapter 1

# Introduction

Environmental challenges such as climate change, biodiversity loss, and pollution are intensifying, threatening ecosystems and human well-being. Rising global temperatures, extreme weather events, and declining natural resources underscore the urgent need for corporate accountability in sustainability efforts (L. Zhang et al., 2022). The Paris Agreement (2015) aims to limit global warming to below 2°C, with efforts for 1.5°C, by reducing greenhouse gas emissions. It mandates rapid emission cuts, financial aid, and technology transfer from developed to developing nations. A transparency framework ensures accountability, with progress reviewed every five years (UNFCCC, 2015). The oil and gas industry, a dominant player in global energy production, has also been a major contributor to greenhouse gas emissions, accounting for over 35% of the energy sector's total emissions since 1965 (Taylor & Watts, 2019). This makes it critical to objectively assess their decarbonization efforts and sustainability claims. While many companies publicly commit to sustainability, their true environmental impact remains unknown, raising concerns about greenwashing, the strategic exaggeration or misrepresentation of environmental efforts.

Greenwashing allows corporations to maintain a sustainable public image while continuing practices that contribute to environmental degradation (Lyon & Maxwell, 2011). Prior studies indicate that sustainability reports, which are often voluntary, tend to present an overly optimistic view of corporate responsibility, while legally mandated financial disclosures

such as 10-K filings provide a more constrained and transparent assessment. Despite growing academic attention, existing methods for detecting greenwashing rely on manual content analysis, Environmental, Social, and Governance (ESG) scoring, or small-scale studies, limiting their scalability and objectivity.

This study introduces a large-scale, AI-driven approach to measure greenwashing systematically. Using LLMs and NLP, we analyze corporate sustainability claims across 10-K filings, annual reports, and sustainability reports, developing a quantitative greenwashing index that differentiates between symbolic and substantive environmental claims. Our findings indicate that larger companies and firms in highly regulated markets engage in more sophisticated greenwashing, shaping public perception without corresponding environmental actions. Greenwashing has also increased over time, particularly after the 2015 Paris Agreement, as firms adapt their sustainability narratives under regulatory and investor pressure.

This research makes several important contributions to the ongoing debate on corporate sustainability accountability by offering empirical evidence on the prevalence, structure, and drivers of greenwashing in the oil and gas sector. It introduces the largest greenwashing dataset compiled to date, covering 50 oil and gas firms over a ten-year period and across three types of corporate disclosures. As the first study to apply open-source large language models (LLMs) at this scale, it demonstrates the viability of using AI-powered analysis to systematically assess corporate sustainability communication. The study not only provides a scalable and replicable framework for identifying greenwashing but also offers a comprehensive policy analysis that links financial and geographic factors to firms' greenwashing tendencies. By establishing a benchmarking framework for future research, this work enhances the capacity of scholars,

investors, and policymakers to distinguish genuine sustainability efforts from strategic greenwashing. The findings underscore the urgent need for stronger regulatory oversight, standardized disclosure requirements, and independent audits to ensure that corporate sustainability commitments translate into measurable environmental impact.

# Chapter 2

# Related Works

This section critically reviews empirical research on the commitment of oil and gas companies to green energy. While the industry has acknowledged its role in global emissions (Boon, 2019), and made progress, significant challenges remain. The transition to sustainability requires assessing both stated commitments and actual practices.

The review is structured around key themes: Assessing Environmental Challenges and Corporate Responses in the Oil and Gas Sector, Transparency in oil and gas companies, and the prevalence of greenwashing, where companies highlight environmental efforts while continuing harmful practices. It also explores the role of text data analysis and language models in evaluating sustainability claims.

By addressing these themes, the review assesses the industry's true commitment to green initiatives and identifies gaps in literature, leading to the study’s novel methodological approach to corporate sustainability accountability.

## 2.1. Assessing Environmental Challenges and Corporate Responses in the Oil and Gas Sector

The oil and gas industry is a key economic driver but also a major contributor to greenhouse gas (GHG) emissions. Growing pressure from governments, investors, and the public has led companies to adopt climate commitments, yet their responses remain

inconsistent (Gielen et al., 2019). Understanding historical corporate behavior is essential for assessing future commitments. However, much of the existing research focuses on a limited number of Western oil majors, such as BP, Shell, ExxonMobil, and Chevron, overlooking broader industry trends. This narrow scope limits the generalizability of findings and leaves gaps in understanding how oil and gas firms worldwide approach sustainability.

Corporate responses to climate change vary widely. Some companies invest in renewable energy, while others delay climate-related regulations. Rowlands (2000) highlighted BP's early rebranding as "Beyond Petroleum" and its renewable energy investments. In contrast, Supran & Oreskes (2017, 2021) showed that ExxonMobil's internal research acknowledged climate risks while its public statements downplayed them, a tactic similar to misinformation used by the tobacco industry.

Broader classifications of firms have emerged. Pickl (2019) categorized companies into "renewable leaders" and "laggards," noting that firms with larger fossil fuel reserves transition more slowly due to infrastructure and capital constraints. Zhong & Bazilian (2018) examined renewable energy investments, finding that while some firms have engaged in clean energy, they face high costs, technological challenges, and policy uncertainties. A growing body of work has also highlighted the localized ecological consequences of extraction. For instance, the expansion of unconventional oil and gas development in the Appalachian region has led to significant land cover change and forest fragmentation, underscoring the tangible environmental footprint of fossil fuel operations (Grushecky et al., 2022).

Most prior studies focus on external influences like policy, investor pressures, and technological advancements, often neglecting emerging markets (M. Li et al., 2022a). To address these gaps, this study examines the top 50 publicly traded oil and gas firms across

multiple regions, offering a more comprehensive analysis of sustainability commitments and transparency. By expanding the scope beyond Western firms, this research provides new insights into the industry's transition toward sustainability.

## 2.2. Transparency in Oil and Gas Companies

The transparency of oil and gas companies in their sustainability commitments remains a subject of debate. Some researchers argue that these companies engage in greenwashing, selectively disclosing information in response to external pressures, while others contend that increasing scrutiny and regulations have led to greater transparency. This debate highlights the complexity of assessing corporate sustainability claims and the need for a more systematic evaluation of greenwashing behaviors.

Many studies suggest oil and gas firms lack full transparency, often using sustainability disclosures to manage their reputations rather than drive real change. Zhang (2022) found that green finance guidelines, intended to promote sustainability reporting, also incentivize selective disclosure, especially in polluting industries. Supran & Oreskes (2017, 2021) provided evidence of ExxonMobil’s contradictory approach, where internal research acknowledged climate risks, but public statements minimized them, demonstrating strategic manipulation of transparency.

Conversely, some scholars argue that concerns over greenwashing may be overstated. Ruiz-Blanco (2022) found that firms in environmentally sensitive industries face stringent regulations, making deceptive sustainability claims more difficult. Cohen et al. (2020) highlighted that despite low ESG scores, many oil and gas firms contribute to green

technologies like carbon capture, indicating a more complex transparency landscape than critics suggest.

Despite these differing perspectives, most research has focused on external factors like regulations and investor scrutiny, with limited attention to internal financial and geographical influences on greenwashing. This study addresses this gap by employing Gradient Boosting Regression (GBR) to analyze how market characteristics shape sustainability disclosures. By developing a greenwashing metric, this research provides a more comprehensive understanding of corporate transparency in the oil and gas sector and contributes to the ongoing debate on greenwashing.

## 2.3. Greenwashing: Definitions, Measurement, and Limitations in Prior Research

Greenwashing, introduced by Westerveld (1986), has been defined in various ways, with Lyon & Maxwell (2011) describing it as selectively presenting positive environmental performance while concealing negative aspects. However, the absence of a universally accepted definition (Lyon & Montgomery, 2015) has led to inconsistencies in measurement, making greenwashing difficult to assess systematically (Ruiz-Blanco et al., 2022).

Prior studies have used different methods to evaluate greenwashing, yet many face limitations in scope and applicability. Some rely on corporate disclosures and ESG scores, such as Ruiz-Blanco et al. (2022), who found that firms in highly regulated industries engage less in greenwashing, though their reliance on Bloomberg data limits generalizability. Kim & Lyon (2015) measured discrepancies between reported and actual emissions reductions in U.S. electric utilities while Marquis et al. (2016) used regression models to analyze environmental

disclosures, both of which faced data availability constraints. Cohen et al. (2020) examined green innovation through patent data but focused solely on the U.S. market. Zhang (2022) conducted a systematic review but excluded key terminology variations, limiting comprehensiveness.

A major gap in prior research is the lack of a standardized, measurable greenwashing metric applicable across industries. This study addresses this gap by using 10-K reports, annual reports, and sustainability reports from the top 50 global oil and gas companies, ensuring data accuracy through regulatory audits. Applying the decoupling definition of greenwashing, this research develops a scalable and reproducible greenwashing metric. Additionally, Gradient Boosting Regression (GBR) is used to analyze financial and geographical factors influencing greenwashing behaviors, contributing to a more systematic assessment of corporate transparency in the oil and gas sector.

## 2.4. Investigating greenwashing using Text data analysis and language models

Traditional methods of detecting greenwashing, such as ESG scores, manual content analysis, and regression-based approaches, face limitations in scalability, consistency, and objectivity (Kim & Lyon, 2015; Marquis et al., 2016). Recent advancements in natural language processing (NLP) and large language models (LLMs**)** offer new possibilities for systematically analyzing corporate sustainability disclosures at scale.

Early text-based approaches relied on sentiment analysis, topic modeling, and keyword extraction to assess discrepancies between corporate sustainability claims and external ESG

metrics (Ruiz-Blanco et al., 2022). However, such methods struggle to capture the complexity of strategic obfuscation in corporate language. The development of domain-specific language models, such as ClimateBERT, has improved the ability to detect vague, misleading, or exaggerated sustainability claims within corporate filings (Garrido-Merch'an et al., 2023). By leveraging transformer-based architecture, these models analyze corporate disclosures with higher semantic precision, identifying patterns indicative of selective disclosure and reputational risk management.

While prior studies have focused on small datasets or region-specific analyses (Cohen et al., 2020; Kim & Lyon, 2015), this study expands the scope by analyzing a global dataset of 50 oil and gas firms across multiple financial and regulatory environments. Unlike previous research, which often relies on a single document type, our approach integrates data from three key corporate disclosures: 10-K filings, annual reports, and sustainability reports, ensuring a more comprehensive representation of corporate greenwashing behavior.

Furthermore, to the best of our knowledge, this study is the first to apply open-source LLMs at this scale to analyze greenwashing trends in the oil and gas industry. By combining LLM-driven text analysis with financial and geographic factors, we systematically investigate how regulatory frameworks, market conditions, and firm-specific characteristics influence greenwashing behavior across different regions. This large-scale, data-driven approach provides a more rigorous, scalable, and replicable framework for assessing corporate environmental accountability.

# Chapter 3

# Overview of the Data

The role of oil and gas companies in climate change mitigation and corporate sustainability remains contested, given their long-standing involvement in obstructive climate actions and greenwashing (Lyon & Maxwell, 2011; Supran & Oreskes, 2017). While past studies have examined greenwashing in individual cases (Kim & Lyon, 2015; Marquis et al., 2016), a scalable, data-driven framework to quantify and monitor it over time is still absent. Most existing work relies on manual content analysis or financial heuristics, typically focused on a few major firms like ExxonMobil, BP, Shell, and Chevron (Supran & Oreskes, 2017; D. Zhang, 2022) lacking an industry-wide approach that links greenwashing to financial, geographic, or regulatory factors.

This study addresses that gap by introducing a novel dataset of 50 publicly traded oil and gas firms across ten countries and four continents, covering 2011–2020. It analyzes three key report types, 10-K filings, annual reports, and sustainability reports, as core communication channels for assessing corporate sustainability claims. The dataset also includes firm-level market data (revenue, profit, employee count, and location) to support policy analysis in Section 5 on the drivers of greenwashing.

Building on Li et al. (2024), who studied four firms' annual reports, this study scales the analysis to 50 firms and three document types. Although disclosures are increasingly

available, no standard method exists for evaluating greenwashing at this scale due to the constraints of manual analysis. With recent advances in Large Language Models (LLMs), which outperform humans on many text tasks (Brown et al., 2020), automated, scalable, and holistic analysis is now feasible.

## 3.1. The Companies

This section provides an overview of the 50 oil and gas companies analyzed in this study. As discussed earlier, our selection is based on the principle of corporate accountability, particularly in the context of stakeholder-oriented corporate reports. Given that publicly traded companies are legally obligated to disclose accurate and verifiable information in their financial and sustainability reports, their filings offer a reliable dataset for analyzing greenwashing trends and corporate sustainability claims.

To ensure a global and representative selection, we focused on publicly traded oil and gas companies listed in the United States during the period 2011–2020. These firms were required by regulatory frameworks such as the SEC's disclosure rules to publish annual reports, making them suitable candidates for systematic analysis.

Given the vast number of energy companies worldwide, we selected the top 50 firms based on Market Capitalization as of Fall 2023. Market Capitalization serves as a proxy for company size and global influence, ensuring that our dataset includes the most economically significant players in the industry. Larger oil and gas firms, such as BP and Shell, significantly contribute to global greenhouse gas emissions, ranking among the top 20 fossil fuel companies responsible for over one-third of such emissions since 1965 (Global Witness, 2023). These

firms face heightened investor scrutiny, with institutional investors challenging their LNG expansion strategies (Australasian Centre for Corporate Responsibility, 2025), while 60% of their public messaging includes 'green' claims despite misalignment with actual low-carbon investments (InfluenceMap, 2025). This intersection of environmental impact, investor oversight, and sustainability claims, positions these firms as key subjects for examining corporate greenwashing.

The dataset includes companies headquartered across a broad geographic range, reflecting diverse regional contexts. In North America, there are 32 firms based in the United States and 7 in Canada. The European subset includes 2 firms from the United Kingdom and 1 each from Italy, Norway, and France. South America is represented by 1 firm each from Colombia, Brazil, and Argentina. Additionally, the dataset includes 3 firms based in China, representing Asia.

This geographical diversity allows us to analyze regional variations in sustainability practices, regulatory influence, and differences in greenwashing behaviors across distinct market environments. In the next section, we detail the three types of corporate reports analyzed in this study and their role in assessing corporate sustainability claims.

## 3.2. Market Data Features of the Oil and Gas Industry Companies (2011-2020)

In this section, we examine the key financial and market features of the 50 largest publicly traded oil and gas companies from 2011 to 2020. These features offer a comprehensive view of the companies' financial health and market standing, which are critical for

understanding their operational dynamics and potential influence on environmental practices, including greenwashing tendencies. The data were meticulously collected from various credible sources, including corporate financial statements, industry reports, and databases such as Wharton Research Data Services (WRDS).

To keep the main body concise, the detailed descriptions of the numerical market features used in the analysis have been moved to Appendix A. These features include financial indicators such as revenue, income, and debt, which are later used to assess greenwashing behavior in Chapter 5.

These features are critical for understanding the financial and operational characteristics of each company in our study. As we proceed through the thesis, these features will serve as the foundation for analyzing potential correlations between company characteristics and greenwashing behaviors. By evaluating these features, we aim to uncover patterns that may reveal how certain financial metrics and market positions influence the likelihood and extent of greenwashing activities in the oil and gas industry.

## 3.3. Overview of Datasets from Annual, 10-K, and Sustainability Reports

In this section, we examine the three distinct types of datasets that form the backbone of our analysis on greenwashing practices within the oil and gas industry. These datasets, derived from Annual Reports, 10-K filings, and Sustainability Reports, offer varied perspectives on the financial, operational, and environmental performance of the companies under study. Understanding the nuances of each dataset is crucial, as they serve as the primary

sources of information for evaluating the companies' greenwashing activities over the 10-year period from 2011 to 2020.

### 3.3.1. Annual Reports

Annual Reports are comprehensive documents that publicly traded companies are required to file annually. These reports provide a detailed overview of the company's financial performance, including its income statement, balance sheet, cash flow statement, and accompanying notes that offer further insights into the company's financial health. The Annual Report also includes a Management Discussion and Analysis (MD&A), which offers a narrative explanation of the company's performance, significant events, and future outlook.

For this study, the Annual Reports serve as a reliable source of financial data, as all companies are mandated by law to produce them. Our analysis covers a decade, from 2011 to 2020, aiming to capture 500 Annual Reports, one for each company per year. However, due to several factors such as companies not being publicly traded for the entire duration, potential lapses in data collection, and some companies issuing 10-K reports as their annual reports, we were able to compile a total of 346 reports. These documents provide a baseline for comparing financial performance and sustainability efforts against the other types of reports.

Annual Reports are particularly valuable for this research as they offer a consistent and regulated source of information, making them ideal for longitudinal analysis. The financial data extracted from these reports are crucial for understanding the economic context in which potential greenwashing practices occur. Moreover, the narrative sections of these reports allow

us to evaluate the extent and nature of environmental claims made by the companies, providing a textual data source for assessing greenwashing tendencies.

### 3.3.2. 10-K Filings

The 10-K is a detailed annual report submitted by U.S. public companies to the Securities and Exchange Commission (SEC). It provides a more in-depth and legally mandated disclosure than the company's Annual Report. The 10-K includes extensive information on a company's financial condition, including a thorough examination of risks, legal proceedings, executive compensation, and market risk factors, alongside the standard financial statements.

The 10-K reports are particularly significant for our analysis because they are subject to stringent regulatory requirements, which means the data within them are highly detailed and structured. While the primary content overlaps with that of the Annual Reports, the 10-K provides additional disclosures that can be pivotal in identifying risks and strategic decisions that may influence or relate to greenwashing practices.

Environmental policies are often discussed in 10-K reports. While the 10-K is a legal document, companies often use separate sustainability reports, ESG reports, or corporate social responsibility (CSR) reports to highlight environmental policies in a more favorable way for investors and the public. These external reports may go beyond legal requirements, emphasizing green initiatives, carbon reduction goals, or sustainability efforts without the same level of regulatory scrutiny as the 10-K.

Out of the potential 500 reports, our dataset includes 496 10-K filings. The near-complete collection of these reports provides a robust dataset for examining the financial and

operational contexts in which these companies operate. The detailed risk factors and management discussions in the 10-K filings offer valuable insights into the potential motivations behind environmental claims, helping to identify whether these claims are backed by substantive actions or merely serve as a veneer to cover less sustainable practices.

### 3.3.3. Sustainability Reports

Sustainability Reports are voluntary disclosures that many companies produce to communicate their environmental, social, and governance (ESG) efforts to stakeholders. These reports typically outline a company's initiatives and performance in areas such as carbon footprint reduction, energy efficiency, waste management, community engagement, and governance practices.

Unlike the legally mandated Annual Reports and 10-K filings, Sustainability Reports are often crafted with public relations in mind, and therefore, the level of transparency and the metrics used vary widely from one company to another. The narrative nature of these reports, combined with selective disclosures, makes them a critical resource for assessing greenwashing. By analyzing the content of these reports, we can evaluate whether the companies' stated environmental goals align with their actions, or if these reports are used to create a misleading impression of environmental stewardship.

For this study, we collected 182 Sustainability Reports, which, while less than the desired 500, provide a significant sample for examining trends and identifying discrepancies between environmental claims and actual practices. The voluntary nature of these reports also introduces an additional layer of complexity; companies that produce these reports may be

more inclined to engage in greenwashing, as they control the narrative, and the metrics disclosed.

### 3.3.4. Summary

The combination of these three datasets, Annual Reports, 10-K filings, and Sustainability Reports, provides a comprehensive view of the oil and gas companies' financial and environmental practices over the last decade. While each type of report offers unique insights, together they allow for a multi-dimensional analysis of how these companies present themselves publicly versus their actual performance and impact. The disparities and consistencies found across these datasets will be crucial in our subsequent analysis of greenwashing levels and the identification of any patterns or correlations between company characteristics and their environmental practices.

As we move forward, these datasets will be used to calculate greenwashing scores for each company, which will then be analyzed in conjunction with the market data features discussed in Section 3.2. This analysis will help us determine whether certain financial or operational characteristics are predictive of higher levels of greenwashing, thereby contributing to the broader understanding of how corporate behavior in the oil and gas industry intersects with environmental accountability.

# Chapter 4

# Greenwashing

This section aims to quantify greenwashing in oil and gas companies using the three datasets introduced earlier: 10-K filings, annual reports, and sustainability reports. We begin by defining greenwashing based on existing literature and use this foundation to construct a scalable metric. This metric is then applied to evaluate patterns across report types and firm behaviors, identifying which disclosures are more prone to greenwashing. The resulting index will later support policy-level analysis in Section 5, examining its relationship with firm-level market characteristics.

## 4.1. Definition of Greenwashing

Greenwashing, broadly defined as misleading or exaggerated claims about an organization's environmental efforts to appear more responsible than it is, remains loosely defined in literature. Despite its widespread recognition, the lack of a standardized mathematical measure poses challenges for consistent quantification and cross-sector comparisons. Researchers have proposed various definitions and types of greenwashing, the most prevalent being selective disclosure, decoupling, attention deflection, and deceptive manipulation.

Selective disclosure, as described by Yang et al. (2020), involves companies selectively highlighting favorable information while omitting less favorable content, making it difficult to measure greenwashing due to incomplete or skewed data. Decoupling, in contrast, focuses on the disconnect between symbolic statements and substantive actions. According to de Freitas Netto et al. (2020), greenwashing occurs when companies make environmental claims without implementing meaningful actions.

This study introduces a Greenwashing Index (GI) grounded in the decoupling framework, conceptualizing greenwashing as the divergence between what firms say and what they do. The GI quantifies this by comparing the number of substantive actions to the total number of environmental claims, expressed as:

$$Greenwashing = 1 - \left( \frac{Number\ of\ Actions}{Total\ Number\ of\ Claims} \right) \quad (4.1)$$

A higher GI reflects greater reliance on symbolic rhetoric, while a lower score indicates stronger alignment between claims and actions. The index serves as a scalable proxy for evaluating environmental communication across firms, time periods, and report types, and forms the analytical foundation of this study.

By applying this framework, the research provides a systematic and scalable method for detecting greenwashing, contributing to the broader discourse on corporate sustainability accountability.

## 4.2. Identification of Claims and Actions

To classify environmental language into symbolic claims or substantive actions, both manual and automated approaches were used. A human-annotated dataset of 100 corporate report pages served as a benchmark. Sentences were classified as follows:

- **Actions**: Statements that included specific, operational commitments. These typically mentioned budget allocations, implementation timelines, or measurable outcomes. For example, "We have allocated $10 million to improve our renewable energy infrastructure." was categorized as an action.

- **Claims**: Statements that were general, aspirational, or lacked measurable detail. Examples include "We are committed to sustainability" or "We aim to reduce our carbon footprint."

The annotation process followed structured guidelines to ensure consistency. While subjectivity cannot be entirely removed, the use of clear classification criteria reduced ambiguity and provided a robust foundation for model validation.

## 4.3. LLM-Based Classification and Model Justification

With this definition in hand, the challenge of measuring greenwashing becomes one of processing company documents to count the number of claims and actions accurately. Previous studies have typically relied on environmental lexicon sets, where an algorithm scans the text and flags words that match a predefined set of terms. Based on the frequency of these terms and their categorization, a greenwashing measure is computed. However, this method is

limited by its inability to capture the broader context of the reports, and lexicon sets often need to be tailored to each specific report, making the approach unscalable and prone to errors.

To address these limitations, for large-scale analysis, the classification process was automated using LLaMA 3-70B via Groq Cloud. Specifically, we employ the 70-billion-parameter version of Llama 3, developed and open-sourced by Meta Inc. Llama 3-70B has demonstrated strong performance on key reasoning and comprehension benchmarks, achieving 94.8% on the Abstraction and Reasoning Corpus (ARC) Challenge and 86% on the Massive Multitask Language Understanding (MMLU) dataset, making it one of the most powerful open-source language models suitable for prompt-based instruction tasks available at the time of writing. With appropriate instructions and prompts, the performance of these models is nearly on par with that of human experts. The model was prompted to classify sentences according to the same claim-action definitions used in manual annotation. LLaMA's ability to reason contextually makes it particularly suited for analyzing formal business language. Figure 4.1 presents the prompt structure.

**Prompt Given to the Language Model:**

*You are an intelligent assistant specializing in environmental sustainability. Your task is to help me measure the scale of greenwashing by oil and gas companies by reviewing their public reports. You will receive one page of their report and, relying solely on the context within this page, you should identify and count the frequency of promises, pledges, and claims, as well as concrete actions related to environmental sustainability.*

*Greenwashing will be assessed based on the frequency of environmental claims and promises versus the number of actual concrete environmental actions.*

**Definitions:**

- **Environmental Claim:** An environmental claim, pledge, or promise (collectively referred to as claims) occurs when a company states they care about the environment or plan to work on environmental goals without providing concrete numbers or a budget to support these plans. Claims must be explicitly about environmental topics, such as reducing emissions, increasing renewable energy use, protecting ecosystems, or similar activities.
  *Examples of Environmental Claims:*
  - "We are committed to reducing our carbon footprint."
  - "We aim to increase renewable energy use in the future."
- **Environmental Action:** An environmental action is when a company specifies a clear budget or number for money or resources allocated to their environmental plan. Actions must include a specific number regarding the budget or resources and be explicitly about environmental topics, such as reducing emissions, increasing renewable energy use, protecting ecosystems, or similar activities.
  *Examples of Environmental Actions:*
  - "We have allocated $10 million to improve our renewable energy infrastructure."
  - "We will reduce emissions by 20% by 2025."

**Instructions:**

- *Count the frequency of environmental claims and environmental actions.*
- *Ensure the output is in the following format: OUTPUT={"claims": NUMBER_OF_CLAIMS, "actions": NUMBER_OF_ACTIONS}, where NUMBER_OF_CLAIMS and NUMBER_OF_ACTIONS indicate the number of environmental claims and actions, respectively.*
- *Identify distinct environmental claims and actions, meaning that if the exact same claim or action is repeated within the same page, count it only once.*
- *Report the parts of the text that are environmental claims and actions back to me.*

**Input:**

- *The data you need to process is any text following the <REPORT_PAGE_DATA> marker.*
- *Once again, the output MUST have the following format: OUTPUT={"claims": NUMBER_OF_CLAIMS, "actions": NUMBER_OF_ACTIONS}*

Figure 4.1. Prompt structure provided to LLaMA 3-70B for classifying environmental claims and actions. The prompt includes explicit definitions, examples, and output formatting instructions aligned with manual annotation criteria. This structure enables the model to reason contextually and consistently across various corporate disclosures.

Figure 4.2 also illustrates the model's performance on a page from a 10-K report, where it accurately identifies claims and actions, highlighting the model's ability to comprehend complex language. Given the instructions as a prompt, the LLM consumes the given text and identifies claims (symbolic claims) vs. actions; the model can also identify why a given phrase in the initial report is flagged as a claim and not an action. This report is taken from page 41 of ExxonMobil's 10K report to United States Securities and Exchange Commission in the year 2020. This accuracy is a result of extensive prompt tuning and the inherent expressiveness of LLMs.

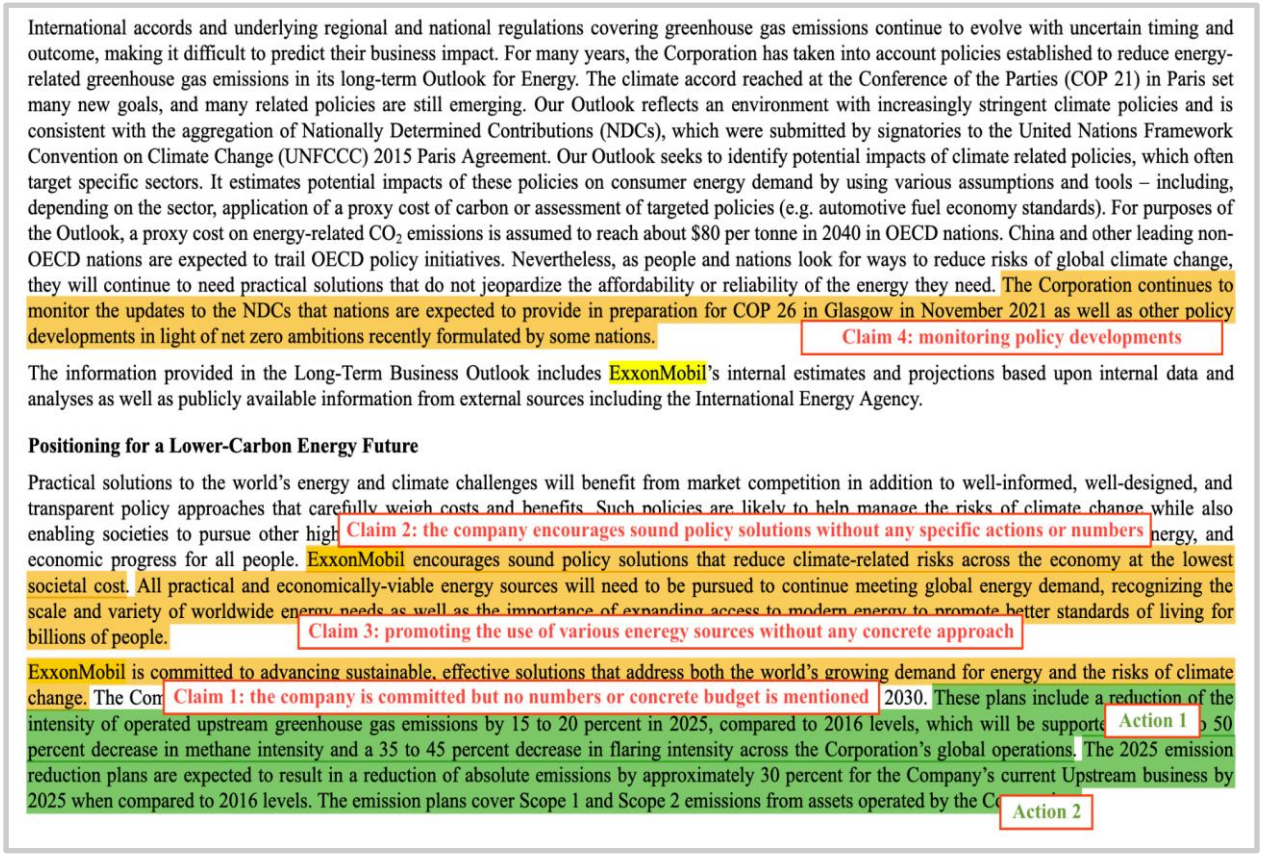



Figure 4.2. Application of LLaMA 3-70B on a sample 10-K report page from ExxonMobil (2020). The model identifies and classifies textual segments as either symbolic claims or substantive actions, demonstrating its capacity to parse complex corporate language using prompt-based inference.

## 4.4. Model Performance Evaluation for Llama 3-70B

To evaluate the effectiveness of our Llama 3-70B model in detecting greenwashing-related text, we applied the model to 100 pages of documents collected from three different datasets. The model was tasked with identifying two key textual components: 1. Number of Claims, 2. Number of Actions. Sentences were labeled as actions if they exhibited concrete operational commitments, such as specific implementation plans, timelines, resource allocations, or performance targets. In contrast, sentences were labeled as claims if they were general or aspirational in nature, lacking measurable indicators or defined plans.

The model was evaluated based on how well its classifications aligned with those generated through manual annotation. A sentence was considered correctly classified if the model assigned it to the same category, either a claim or an action, as determined in the human-labeled dataset; otherwise, it was recorded as a misclassification.

To assess the model's performance, we computed three standard evaluation metrics: precision, recall, and F1-score. Precision measures the proportion of correctly identified claims or actions among those flagged by the model, while recall evaluates the proportion of all actual claims or actions that were successfully detected. The F1-score, as the harmonic mean of precision and recall, provides an overall measure of classification performance.

As shown in Table 4.1, the model demonstrated strong alignment with human judgment across both claim and action detection tasks. For claims, the model achieved a relatively high recall, indicating its effectiveness in capturing the majority of true claims within the dataset. Similarly, the results for action detection reveal the model's capability to identify substantive

content with high accuracy. While minor trade-offs between precision and recall were observed, the overall F1-scores support the model's robustness and suitability for large-scale greenwashing classification across diverse types of corporate disclosures.

These results underscore LLaMA 3-70B's effectiveness as a scalable and reliable tool for analyzing complex corporate sustainability reports, where manual classification would be resource-intensive and susceptible to subjectivity.

Table 4.1. Llama 3-70B Model Performance Summary

| | Standard Evaluation metrics | Model performance on Claims | Model performance on Actions |
|---|---|---|---|
| 1 | Precision | 0.78 | 0.75 |
| 2 | Recall | 0.85 | 0.88 |
| 3 | F1-score | 0.813 | 0.809 |

Before we present our experimental results, it is important to note that, to the best of our knowledge, this work is the first to apply decoupling analytics for greenwashing using LLMs. We believe that the continued advancement of natural language processing and the improving performance of language models will significantly benefit environmental analytics, enabling the extraction of valuable insights that would otherwise require extensive manual labor, which is often cumbersome, unproductive, unscalable, and prone to errors due to biased judgments and unclear definitions. It is also worth mentioning that recent efforts, such as the

development of ClimateBERT, have begun to introduce such technologies to environmental studies. ClimateBERT is a general-purpose language model fine-tuned on 2 million climate-related texts to enhance its environmental understanding. Although investigating the potential of ClimateBERT and its derivatives to further improve greenwashing analytics is beyond the scope of this study, it remains an intriguing avenue for future research.

## 4.5. Justification and Methodological Considerations

A critical consideration in this study is the justification for using a general-purpose large language model (LLM) rather than a fine-tuned sustainability-specific model. While fine-tuning could potentially improve performance by adapting the model to the specific language and nuances of ESG reporting, it also requires a large volume of labeled training data, which is not currently available for this task. Moreover, fine-tuning carries the risk of overfitting, which could reduce generalizability across firms, sectors, or report formats. Instead, a zero-shot approach using prompt engineering was adopted, allowing the model to apply learned reasoning patterns to unseen sustainability content while maintaining flexibility and scalability.

The prompt used in the classification process clearly defined both symbolic claims and substantive actions using the same criteria applied during manual annotation. The model read each page and evaluated sentences based on these definitions, enabling a scalable and consistent classification approach.

A key limitation of the method is that it does not verify whether forward-looking statements were eventually fulfilled. Some statements classified as actions may refer to plans

that are yet to be implemented, such as “We will reduce carbon emissions by 20% by 2030.” While follow-up tracking would offer additional insight, it is not feasible given the scope of this study, particularly for recent reports where such timelines have not yet elapsed. To ensure consistency across years, all documents were evaluated based solely on the content within the reporting year.

Despite this limitation, the method captures essential features of greenwashing by focusing on the structure and specificity of the environmental language. It represents a novel application of advanced LLMs to measure the rhetorical composition of corporate sustainability communication at scale. Unlike prior methods based on keyword frequencies or ESG scores, this approach enables sentence-level analysis grounded in semantic understanding and contextual reasoning. The result is a replicable and interpretable Greenwashing Index that can inform both academic research and policy-oriented investigations.

## 4.6. Experimental Setup Overview

To accurately calculate greenwashing using large language models (LLMs), we hypothesize that a broader context window is advantageous for the model to effectively capture the nuances and context within the text. Accordingly, we choose to process large corporate reports on a page-by-page basis, ensuring that each page's content is fully utilized by the model. For this study, we utilize the 70-billion-parameter version of the Llama 3 model, which supports a context size of up to 8192 tokens, providing a substantial capacity to handle the detailed and dense nature of these reports.

The experimental workflow begins by parsing each document sequentially, where each page is individually extracted and then fed into the LLM for inference, guided by the prompt outlined in Figure 4.2. To optimize computational resources and processing time, we implement a pre-screening phase prior to LLM processing. This pre-screening involves scanning each page for the presence of environmentally relevant terminology, which serves as an indicator of whether the page likely contains discussions pertinent to environmental claims or actions. Following the methodology of Li et al. (2022b), we employ a predefined list of terms, as shown in Table 4.2, to conduct this initial filtering. The list is deliberately simple and categorized loosely based on former lexicon set analysis studies, encompassing terms related to emissions and climate, energy sources, and agreements and policies. The filtering process is intentionally lenient, requiring only the occurrence of these terms for a page to proceed to the next stage, thus minimizing the risk of excluding potentially relevant content.

After prescreening, only the pages that meet these criteria are processed by the LLM, which identifies and counts the environmental claims and actions as defined by our methodology. This approach ensures that non-informative pages, such as blank or irrelevant content, are excluded from analysis, streamlining the process while maintaining focus on environmentally significant data. The outcome of this strategy is twofold: first, we obtain the number of pages from each document that were processed, and second, we calculate the total number of environmental claims and actions identified by the model. These metrics are critical for the subsequent stages of analysis, where they will be used to quantify and compare the extent of greenwashing across different reports.

Table 4.2. Categories of Environmental Terms used for pre-screening content before feeding them to the LLM

| **Category** | **Terms** |
|---|---|
| Emissions & Climate | environment, emission, footprint, carbon, greenhouse, ghg, climate, net zero, net-zero, efficiency |
| Energy Sources | clean, clean fuel, biofuel, renewable, clean energy, green energy, alternative energy, ccus |
| Agreements & Policies | paris agreement, unfccc, ipcc |

## 4.7. Ablation Studies

In this section, we perform a thorough analysis of the experimental results obtained from applying our greenwashing scoring algorithm to the three datasets of the oil and gas industry collected in Section 3. Each report within these datasets is associated with three key metrics: the number of pages, the number of environmental claims, and the number of environmental actions. Using these metrics, we calculate the greenwashing score for each report according to Equation 4.1. This score is a continuous value within the [0, 1] range, where a lower score indicates a lower level of greenwashing for the corresponding company. It is important to note that while this score serves as our primary measure of greenwashing,

alternative metrics based on the relative frequency of claims and actions could also be considered; however, the exploration of such metrics falls outside the scope of this study.

The objective of this section is to explore the nuances of greenwashing across the datasets by posing and answering specific research questions grounded in our experimental findings. These questions will help us uncover patterns, outliers, and any systematic differences in greenwashing behaviors across the different types of reports. For instance, we will investigate whether certain companies consistently show lower or higher greenwashing scores across all report types, or if particular report types (e.g., 10-K vs. Sustainability reports) tend to exhibit more pronounced greenwashing.

To support this analysis, we first share the broader data from our experiments. Table 4.3 summarizes the overall statistics for each dataset, i.e., every metric is calculated over the entire records in that dataset over 10 years. As can be seen, these statistics provide only a very rough baseline for comparing the three datasets, and the very high variance for each record indicates the significant differences among the reports across companies and over the years. Figures 4.3 and 4.4 offer a finer-grained view into these three datasets; however, there are some interesting findings and comparisons that we can draw from Table 4.3. As expected, the Sustainability dataset discusses environmental claims and actions far more than the other two datasets, with the Annual Reports scoring slightly higher than the 10-K reports in this regard. Another interesting observation is that for all three datasets, the mean and median greenwashing scores are remarkably close to the 0.5 mark, which indicates that the number of claims and actions are relatively balanced. However, it is worth mentioning that the Sustainability reports have a lower variance in greenwashing levels than the other two datasets,

potentially due to the smaller number of companies and the fact that all these companies are based in the same geographical area (North America).

Table 4.3 shows greenwashing statistics on the whole dataset. The 10K dataset has 496 records over the 50 companies in the period between 2011-2020, the Annual dataset has 346 records over 38 companies, and the Sustainability dataset has 182 records over 23 companies.

Table 4.3. Greenwashing statistics on the whole dataset.

| **Dataset** | **Num Pages** | | | **Num Claims** | | | **Num Actions** | | | **Greenwashing** | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | **Mean** | **Med** | **Std** | **Mean** | **Med** | **Std** | **Mean** | **Med** | **Std** | **Mean** | **Med** | **Std** |
| **10K** | 60.75 | 43 | 63.29 | 50.82 | 22 | 93.14 | 48.91 | 24.58 | 86.83 | 0.51 | 0.49 | 0.15 |
| **Annual** | 49.42 | 37 | 40.14 | 69.47 | 32.5 | 114.06 | 69.11 | 34.5 | 112.9 | 0.51 | 0.51 | 0.12 |
| **Sustainability** | 56.06 | 46.5 | 38.43 | 243.4 | 217.5 | 158.6 | 235.2 | 211 | 149.1 | 0.51 | 0.51 | 0.08 |

Figure 4.3 provides more insight into these statistics by tracking them yearly. As can be seen, each metric has its own chart, and the Y-axis in each one represents the average value for that metric in a given year. Some of the key findings from Figure 4.3 are that the attention to environmental topics has been trending upward, as evidenced by the increasing number of pages, claims, and actions reported. However, at the same time, it appears that greenwashing has been gradually increasing across the industry, which raises concerns as the number of symbolic claims is accelerating past the actual actions.

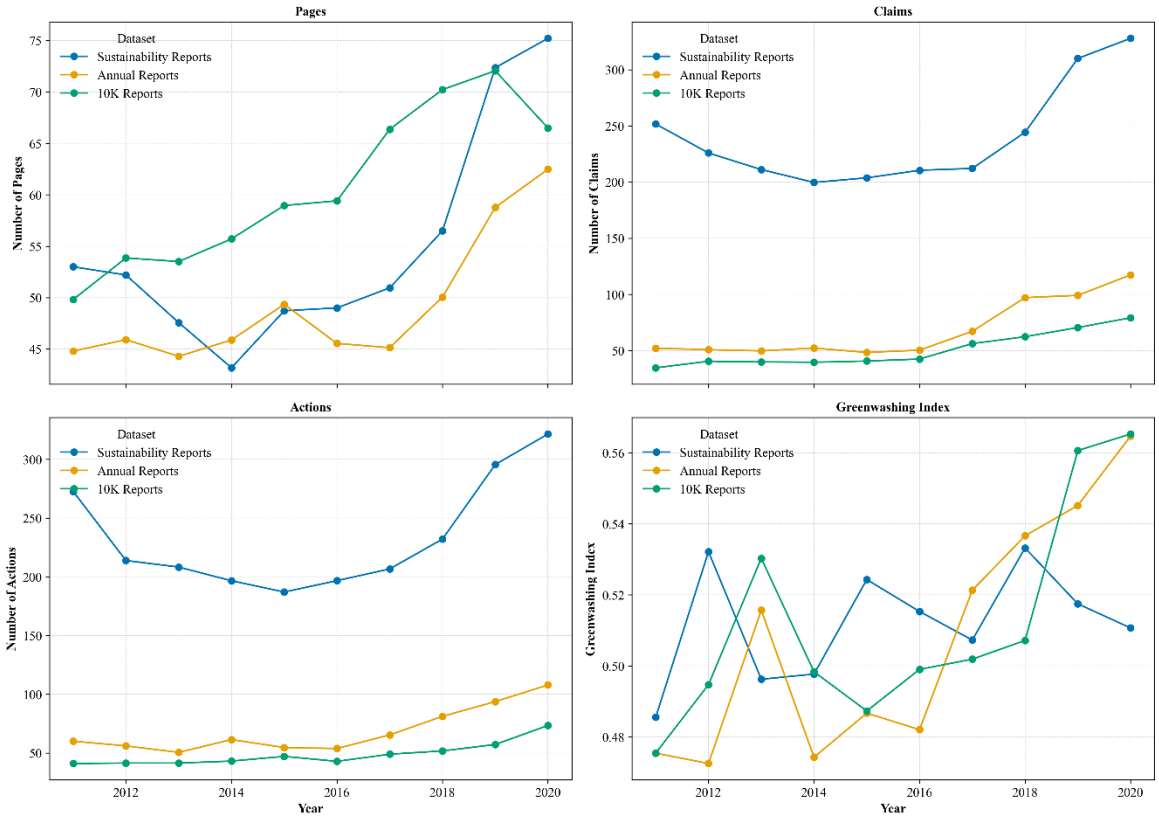



Figure 4.3. Overview of the trends in four metrics, number of pages with environmental mentions, number of symbolic claims, number of substantive actions, and the greenwashing index, over the decade 2011–2020. The mean value of each metric is reported annually for each dataset

The findings reported in Figure 4.3 are alarming. To gain further insights into how this trend evolved, Figure 4.4 analyzes the greenwashing index based on geographical separations, examining the changes in the index per continent. It is worth mentioning that, as reported in Section 3, the majority of the 50 companies studied in this work are from North America, and this is even more pronounced in the cases of the Sustainability and Annual datasets, where fewer companies from outside North America provide their reports. As can be seen, the trend of greenwashing is rising across almost all datasets and continents, with North American and European companies exhibiting the highest levels of greenwashing.

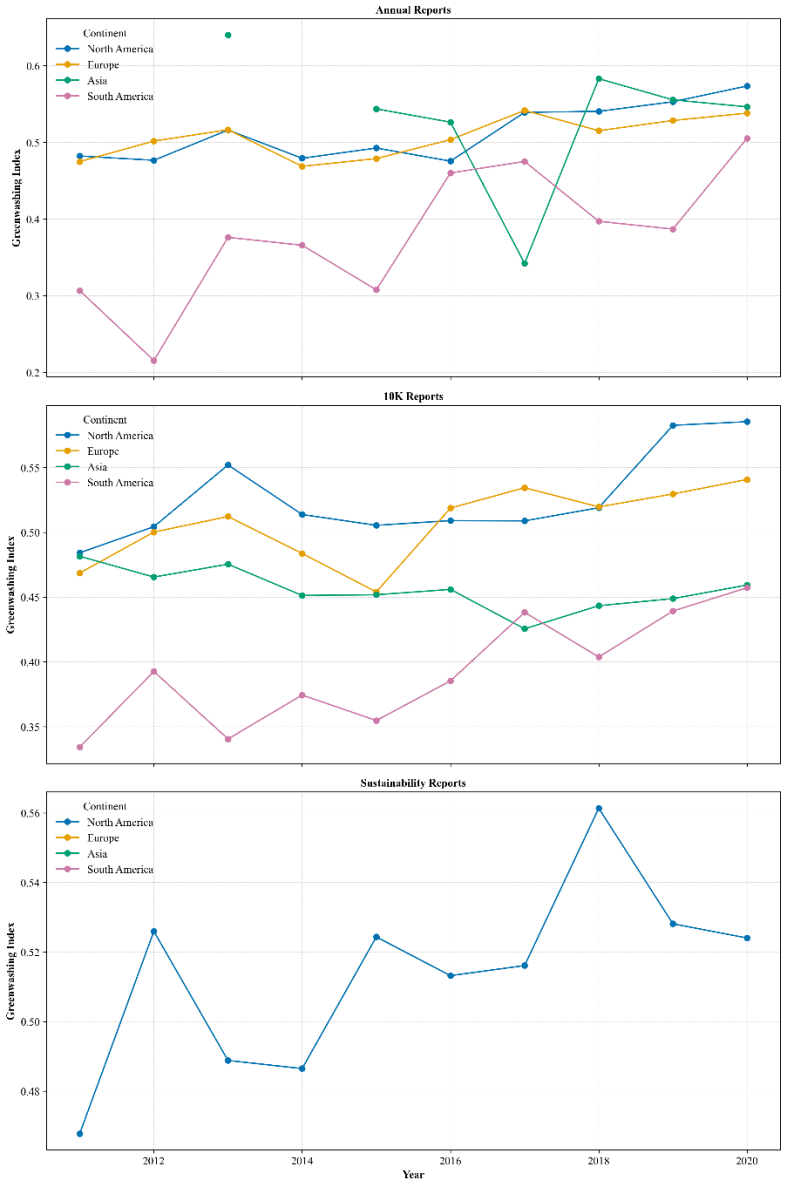



Figure 4.4. Overview of the trend in the greenwashing index for oil and gas companies, categorized by the continent of their headquarters. The reported values represent the mean greenwashing index for all companies within each category.

Now, given the above insights as well as the data from our experiments, we turn our attention to further ablation of these experiments by posing and answering the following question.

**Do companies that publish sustainability reports exhibit lower greenwashing in the 10-K and annual datasets?**

As it can be seen in table 4.4, The analysis indicates that companies publishing sustainability reports exhibit slightly lower greenwashing levels in both 10-K and annual reports. Firms with sustainability reports have a mean greenwashing score of 0.50, compared to 0.53 (10-K) and 0.52 (annual) for those without. This suggests that sustainability reporting may enhance transparency and reduce greenwashing, likely due to greater corporate accountability.

However, the small variance in scores implies that sustainability reports alone do not eliminate greenwashing, as firms may still engage in selective environmental disclosure. Additionally, the standard deviations across datasets indicate variability in corporate behavior, implying that the extent to which sustainability reporting mitigates greenwashing is influenced by firm-specific factors.

Table 4.4. Comparison between greenwashing in companies with sustainability reports and companies without sustainability report

| Dataset | Num Pages | | | Num Claims | | | Num Actions | | | Greenwashing | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Mean | Med | Std | Mean | Med | Std | Mean | Med | Std | Mean | Med | Std |
| **10K with Sustainability report** | 57.23 | 45 | 40.92 | 52.13 | 31.50 | 65.99 | 53.60 | 33 | 57.96 | 0.5 | 0.49 | 0.12 |
| **10K without Sustainability report** | 62.98 | 41 | 77.01 | 49.44 | 18.50 | 110.72 | 44.78 | 20 | 104.74 | 0.53 | 0.5 | 0.16 |
| **Annual with Sustainability report** | 50.09 | 42 | 34.53 | 67.41 | 43 | 79.06 | 65.62 | 44 | 63.73 | 0.5 | 0.5 | 0.1 |
| **Annual without Sustainability report** | 48.71 | 34 | 45.60 | 71.67 | 26.00 | 142.86 | 72.87 | 25 | 148.98 | 0.52 | 0.52 | 0.13 |

# Chapter 5

# Policy Analysis and Conclusion

This section analyzes the relationship between greenwashing and firm-level market characteristics in the oil and gas sector using the data and greenwashing scores developed earlier. We apply a regression model, GradientBoostingRegressor, to assess how factors like geographic location, market capitalization, workforce size, revenue, and year correlate with greenwashing levels (measured on a 0–1 scale). The model captures complex interactions, allowing us to identify which market features most strongly influence greenwashing behaviors across firms.

## 5.1. Experimental Setup

In this section, we outline the experimental setup used to analyze the correlation between market characteristics and greenwashing behaviors in the oil and gas industry. This analysis is conducted using three datasets: the Sustainability dataset (182 records), the Annual dataset (346 records), and the 10-K dataset (482 records). The objective is to model the relationship between various market features and the level of greenwashing using the GradientBoostingRegressor algorithm, a robust gradient boosting method known for its effectiveness in handling structured data.

### 5.1.1. Data Preparation

The datasets encompass a range of financial and operational metrics for companies in the oil and gas industry, spanning the years 2011-2020. For brevity and clarity, these definitions are now relocated to Appendix A. These features cover standard financial and operational metrics (e.g., assets, revenue, debt, ROE, and D/E ratio), alongside country-level categorical data and the year of observation. These variables are used to model the relationship between firm characteristics and greenwashing levels.

The target variable for the regression model is the greenwashing score, which is calculated as the ratio of the number of environmental claims made by the company to the sum of claims and actual environmental actions. This score is a continuous variable ranging between 0 and 1, representing the extent of greenwashing behavior.

To prepare the data, the categorical feature country was encoded into numerical values using label encoding, which assigns a unique integer to each country. This transformation allows the categorical data to be effectively utilized in the regression model. Additionally, all numerical features were normalized using MinMax scaling, ensuring that they are on a comparable scale, which is crucial for the performance of the gradient boosting model.

### 5.1.2. Model Training and Validation

The GradientBoostingRegressor from the sklearn.ensemble library was selected for this regression task due to its effectiveness in handling structured data and capturing complex interactions between features. The dataset was split into training and validation sets, with 80% of the data allocated for training and 20% for validation. This split was repeated 5 times, each

with a different random seed, to ensure that the results are robust and not dependent on a single data partition.

The model was configured with specific parameters to optimize performance and manage complexity. The number of estimators was set to 100, representing the total number of boosting rounds. The learning rate was fixed at 0.1, which determines the contribution of each tree to the final model. To control model complexity, the maximum tree depth was limited to 5. Finally, the model was optimized using Mean Absolute Error (MAE) as the loss function, selected for its robustness to noise and its suitability for continuous target variables.

During the training phase, the model's performance was evaluated at each boosting iteration (i.e., each number of estimators), and the mean absolute error (MAE) was recorded for both the training and validation sets. These losses were averaged across the 5 runs and plotted against the number of boosting iterations to visualize the model’s learning process. This approach enables us to determine the optimal number of boosting iterations by identifying the point where the validation loss stabilizes or begins to increase, indicating potential overfitting.

### 5.1.3. Feature Importance Analysis

Following the model training, we conducted a comprehensive feature importance analysis to identify the market characteristics that most significantly influence the greenwashing score. The GradientBoostingRegressor inherently provides feature importance measures, which quantify the contribution of each feature to the reduction in loss during the boosting process. These importance scores were averaged across the five experimental runs to

ensure robustness and reliability, offering a clear ranking of market features based on their influence on greenwashing behavior. This ranking provides valuable insights into the factors driving greenwashing practices in the oil and gas industry.

To deepen our understanding of the model's behavior and the specific role of each feature, we employed two advanced interpretability techniques: SHapley Additive exPlanations (SHAP) and Partial Dependence Plots (PDP). SHAP values, rooted in cooperative game theory, offer a precise, local explanation of how each feature contributes to individual predictions by measuring the marginal contribution of each feature across all possible feature combinations. This allows us to visualize both global feature importance and localized effects on model predictions. Meanwhile, Partial Dependence Plots provide a global view by illustrating the marginal effect of a particular feature on the predicted outcome, averaged over all other features in the dataset. This dual approach not only enables us to confirm the model's reliance on key features but also to understand the nature and direction of these features' effects, thereby offering a nuanced interpretation of the factors that influence greenwashing in the sector. These methods enhance our model's transparency, ensuring that the findings are interpretable and actionable in the context of sustainability practices in the oil and gas industry.

## 5.2. Experimental Results

This section presents the experimental results obtained from analyzing the correlation between greenwashing behaviors and various market characteristics for the three datasets: Sustainability, Annual, and 10-K. For each dataset, we discuss the training and validation loss curves, feature importances, and how greenwashing scores vary with respect to each feature.

Due to their repetitive nature across the Sustainability, Annual, and 10-K datasets, the full set of Partial Dependence Plots has been relocated to Appendix B. In the main text, we retain only key insights drawn from these plots to streamline interpretation while maintaining technical depth.

### 5.2.1. Sustainability Dataset

- **Training and Validation Loss Curves and Feature Importances**

Illustrated in Figure 5.1, the training and validation loss curves for the Sustainability dataset show that the regression model is successful in learning and converging, despite a visible degree of overfitting, likely due to the small dataset size. For this dataset, the average MAE on the validation data converges to 0.05. Moreover, as shown in Figure 5.1, the feature importance analysis reveals that Market Capitalization is the feature with the most significant impact on model training, suggesting that Market Capitalization is a key determinant of the level of greenwashing. For the Sustainability dataset, the three features of Debt, Number of Employees, and Revenue closely follow Market Capitalization in importance. Next, we will analyze how each of these features and their direction impact the level of greenwashing.

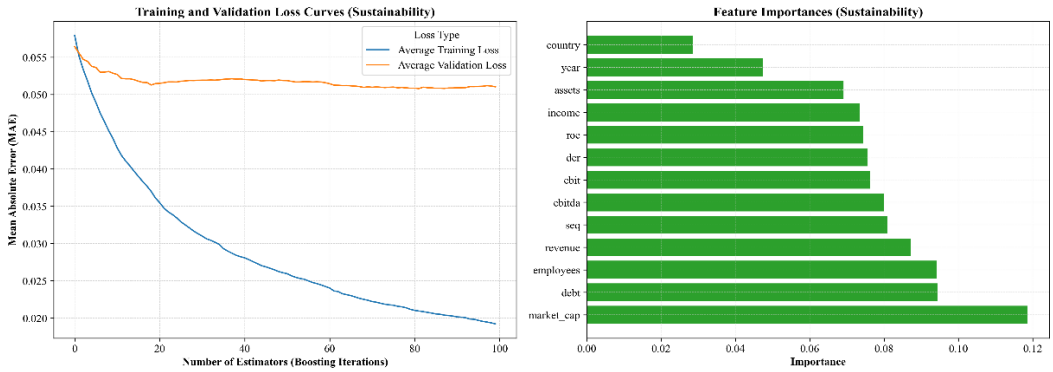



Figure 5.1. Training and validation loss curves (left) for the Sustainability dataset, showing the average MAE loss over five random seed partitions. Corresponding feature importances are shown on the right.

- **Greenwashing Response to Features**

PDP figure of sustainability dataset (See Appendix B for full PDP visualizations) clearly illustrates how changes (increase or decrease) in the value of each feature affect greenwashing, reflecting the model's ability to learn this relationship. For example, as shown, as the value for the feature "Year" increases, the level of greenwashing slightly increases, indicating that the level of greenwashing, when compared over time, has been rising. Another interesting observation is that as a company's "Debt" increases, this also positively correlates with the greenwashing score, meaning that a company with higher "Debt" is likely to exhibit a higher level of greenwashing.

To further elucidate and support the findings above, we also present the chart of the SHAP values for each dataset. Figure 5.2 shows these values for the Sustainability dataset. These values also confirm the observations in PDP plots; for example, as previously mentioned, Market Capitalization is the most important and descriptive indicator for greenwashing in the Sustainability dataset. As Figure 5.2 illustrates, the larger a company's Market Capitalization, the more susceptible it is to engage in greenwashing activities.

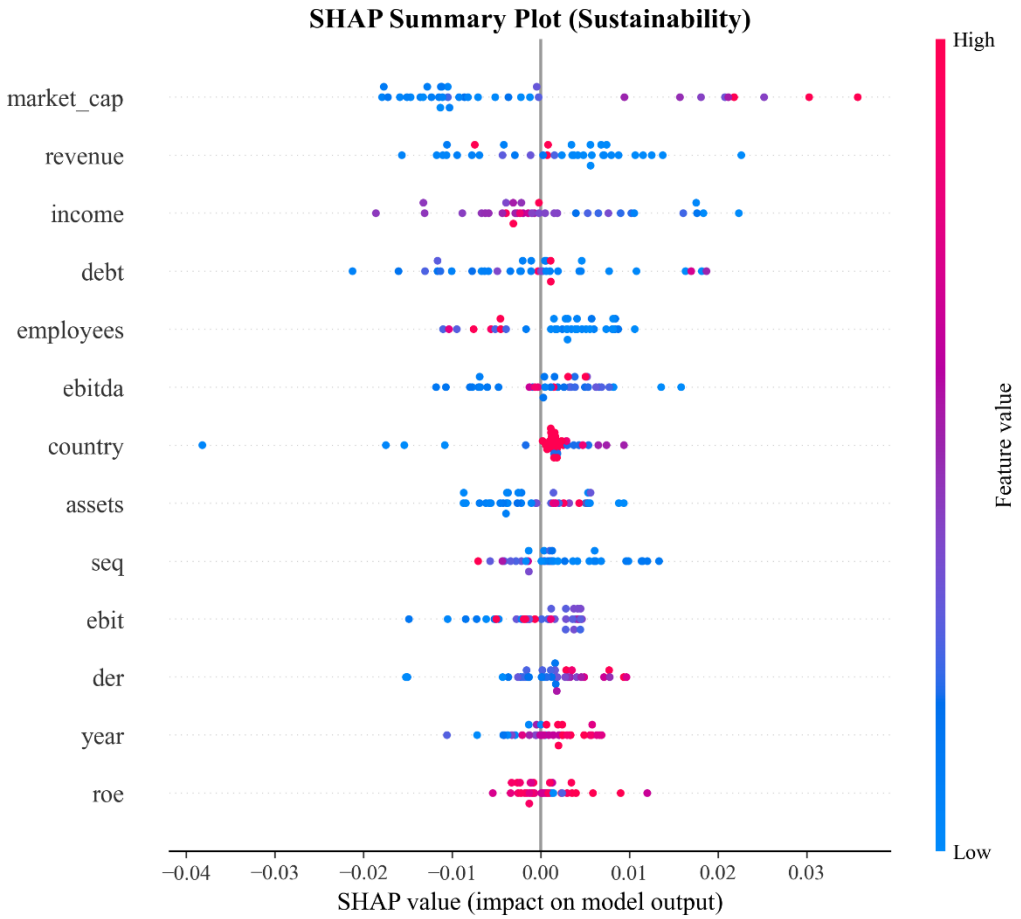



Figure 5.2. SHAP values for the 13 independent features and the corresponding dependent feature, greenwashing, on the Sustainability dataset.

### 5.2.2 Annual Dataset

Given the explainability and transparent nature of our analysis and results, for the Annual and 10-K datasets, we will focus on reporting the results charts and highlighting the most notable findings for each. However, we invite the reader to further analyze our charts and identify the patterns at play. After presenting all the charts, we will summarize the most common themes observed across all three datasets.

- **Training and Validation Loss Curves and Feature Importances**

Figure 5.3 illustrates the loss values for the Annual dataset, followed by the corresponding Feature Importances. As can be seen, the MAE value converges to 0.08, indicating that the model's predictions would, on average, be this distance from the actual value. This value is higher than that for the Sustainability dataset, which further confirms our findings in Section 4 regarding the variances of the greenwashing score in each dataset. Additionally, we observe that The Ratio of Debt to Equity is the most important feature for this dataset, followed by Market Capitalization.

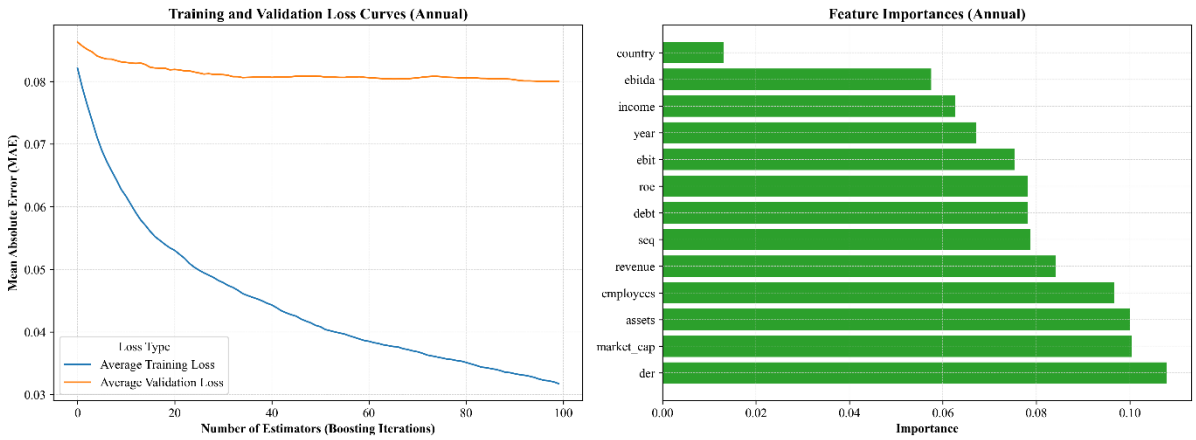



Figure 5.3. Training and Validation Loss Curves (left) for the Annual dataset, representing the average MAE loss value over the 5 random seed partitions., Corresponding Feature Importances (right).

- **Greenwashing Response to Features**

PDP figure of annual dataset (See Appendix B for full PDP visualizations), and SHAP values shown in figure 5.4 show some visible patterns for this dataset including an increase in greenwashing over time, a higher revenue leading to a smaller greenwashing score, and Debt

showing a small negative correlation with the greenwashing score, opposite to the positive correlation observed in the sustainability dataset.

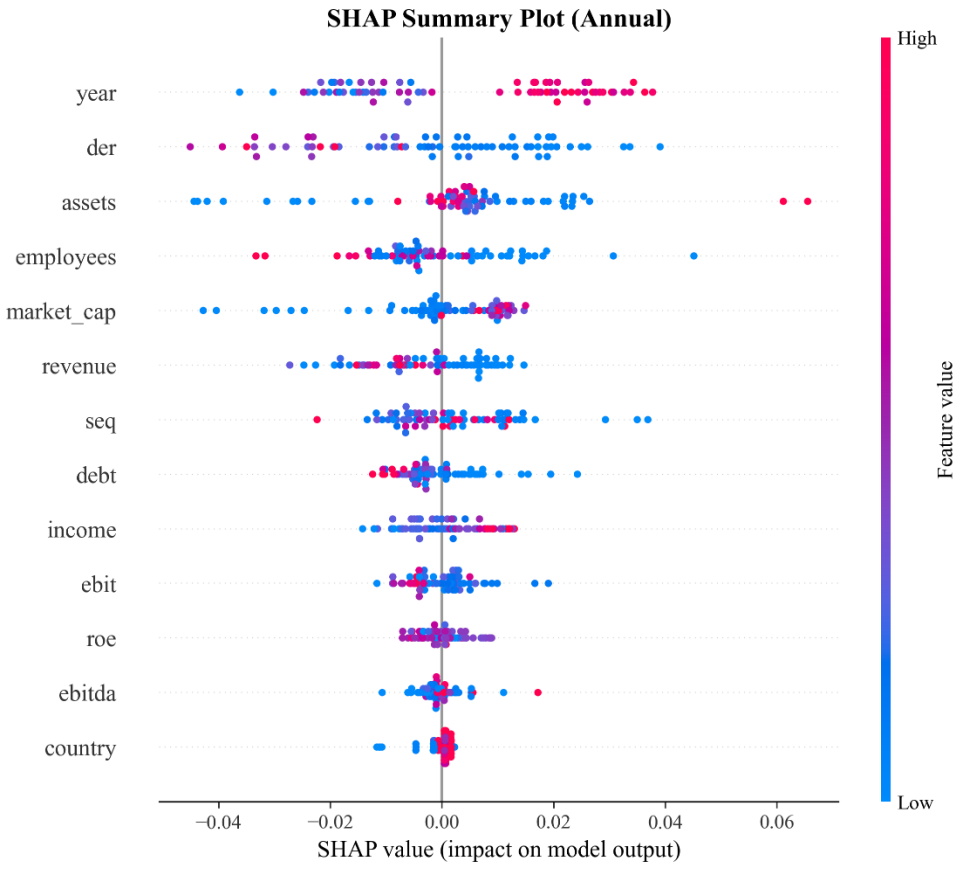



Figure 5.4. SHAP values for the 13 independent features and the corresponding dependent feature, greenwashing, on the Annual dataset.

### 5.2.3. 10K Dataset

- **Training and Validation Loss Curves and Feature Importances**

Figure 5.5 reports the loss and importance values for the 10-K dataset. The MAE loss converges to 0.09, which is consistent with the noise and variance levels reported in Section 4.

"Number of Employees," "Total Assets," and "Debt to Equity Ratio" are the top three most important features for this dataset.

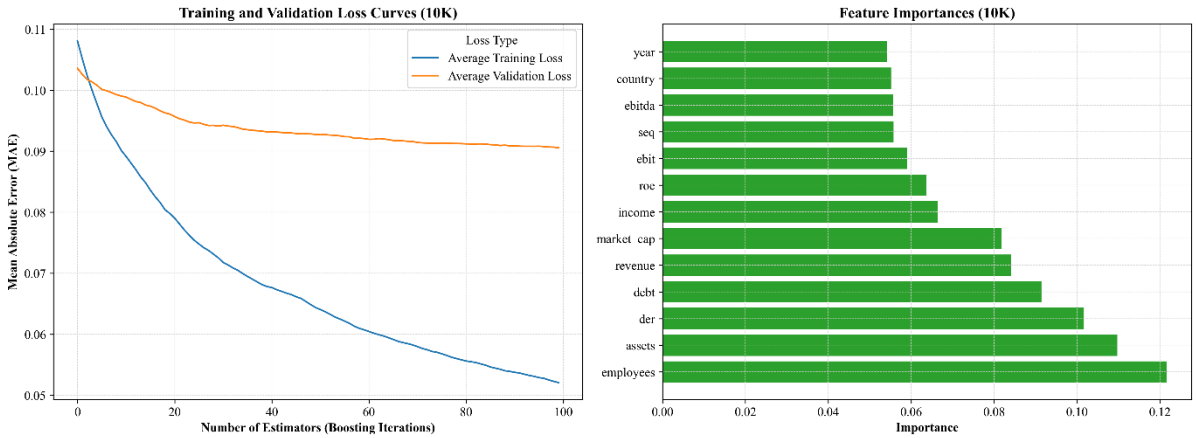



Figure 5.5. Training and Validation Loss Curves (left) for the 10K dataset, representing the average MAE loss value over the 5 random seed partitions. Corresponding Feature Importances (right).

- **Greenwashing Response to Features**

PDP figure of 10K dataset (See Appendix B for full PDP visualizations), and figure 5.6 show how greenwashing scores respond to changes in various features for the 10-K dataset. It is evident that the larger a company is (higher Number of Employees and Market Capitalization), the more likely it is to engage in greenwashing activities. Additionally, the greenwashing score increases over time. However, if a company is performing well financially (Better Revenue or Debt to Equity Ratio), it is more likely to exhibit fewer greenwashing behaviors.

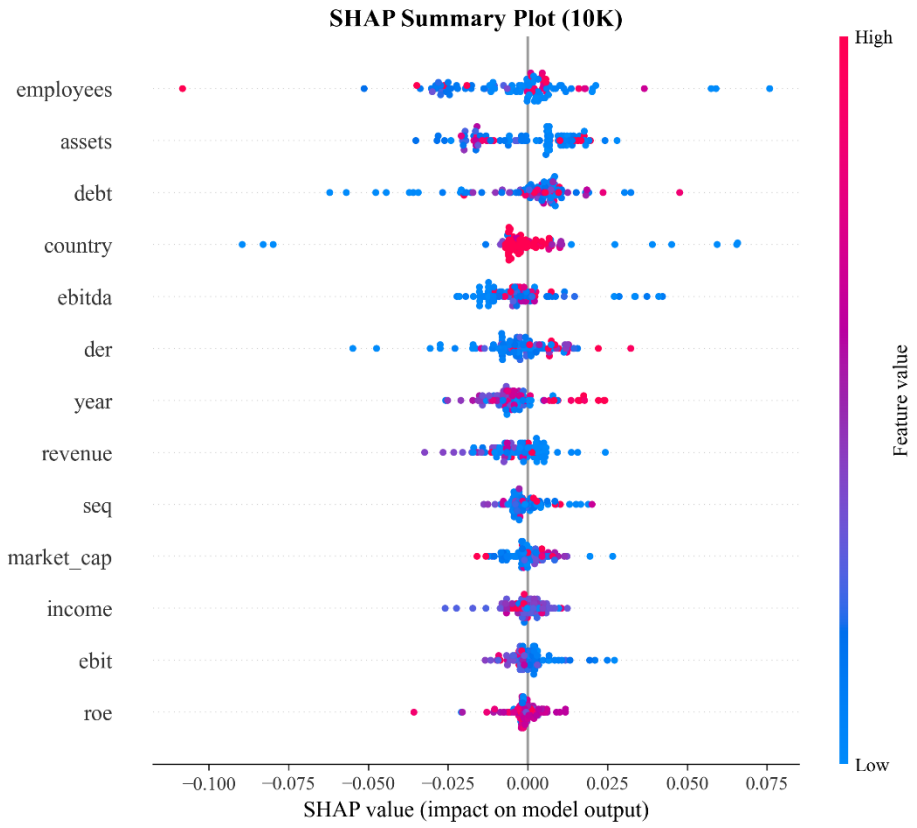



Figure 5.6. SHAP values for the 13 independent features and the corresponding dependent feature, greenwashing, on the 10K dataset.

### 5.2.4. The Common Patterns

Comparing the results across the three datasets, we can observe several patterns, though some may be more prevalent than others. Multiple implicit factors could influence this behavior, such as differing motives behind each report, the fact that some companies do not report certain data, and the scale of the data along with inherent interpretation noise. However, we can broadly identify these three patterns:

- The better a company performs financially, the less likely it is to engage in greenwashing activities.

- Over time, the level of greenwashing has been increasing, with the years 2015-2016 emerging as a potential turning point, after which greenwashing levels generally rise.
- Larger companies, with higher Market Capitalization and Number of Employees, are more likely to engage in greenwashing activities.

This analysis gains urgency in light of projected future landscape disruptions. For example, forecasts suggest that continued oil and gas development in the Permian Basin could lead to over 6,000 km² of new landscape alteration by 2050 under high-impact scenarios, posing significant risks to ecologically sensitive regions like the Chihuahuan Desert (Johnston et al., 2022)

## 5.3. Conclusion

The findings of this thesis reveal the prevalence and sophistication of greenwashing in the oil and gas industry, emphasizing the gap between corporate sustainability claims and actual environmental actions. By leveraging Large Language Models (LLMs) and Natural Language Processing (NLP), this study systematically analyzed corporate disclosures from 50 publicly traded oil and gas firms (2011–2020), identifying patterns of selective disclosure and symbolic environmental commitments. The results show that larger companies and those in highly regulated markets tend to engage in more sophisticated greenwashing, strategically shaping their sustainability narratives to maintain a positive public image. Additionally, greenwashing has increased over time, particularly after the 2015 Paris Agreement, as companies respond to heightened scrutiny with enhanced, yet often misleading, environmental messaging.

This research contributes to the ongoing discourse on corporate accountability and environmental transparency by introducing a scalable, data-driven framework for detecting greenwashing. The proposed quantitative greenwashing index provides an objective measure to differentiate substantive corporate actions from symbolic claims, offering insights for policymakers, investors, and regulatory bodies. The findings underscore the need for stronger regulatory frameworks, standardized reporting requirements, and independent audits to prevent strategic greenwashing and ensure corporate sustainability commitments translate into meaningful climate action. Related studies on offshore platform decommissioning also emphasize that even partial structural removal can lead to measurable changes in benthic and fish communities, suggesting that corporate sustainability reports may understate such ecological consequences (Pierre et al., 2020).

While this study advances the methodology for detecting greenwashing, some limitations remain. The analysis is limited to textual corporate disclosures and does not account for non-textual actions such as actual investment in green technology or operational carbon footprints. Future research could expand by integrating financial performance data, industry-wide emission records, and real-time media monitoring to develop a more comprehensive assessment of corporate sustainability practices. Additionally, fine-tuning LLMs on sustainability-related datasets could enhance the accuracy of greenwashing detection.

As the world moves toward a low-carbon economy, distinguishing genuine sustainability efforts from misleading corporate narratives is critical. This study provides a foundation for systematic corporate greenwashing detection, equipping stakeholders with the tools to hold companies accountable and drive meaningful environmental change.

## 5.4. Implications of Findings and Policy Recommendations

The findings of this study provide not only a scalable framework for detecting corporate greenwashing but also yield meaningful implications for understanding firm behavior and shaping policy. First, the observation that larger firms and those operating in highly regulated markets tend to exhibit higher greenwashing levels suggests that increased scrutiny may lead companies to adopt more sophisticated messaging strategies without necessarily improving their environmental performance. This points to a dynamic where firms respond to pressure with enhanced symbolic disclosures, not always matched by substantive action.

Second, the increase in greenwashing behavior over time, especially after the 2015 Paris Agreement, indicates that global climate accords, while pivotal, may inadvertently incentivize reputational management over genuine operational change unless accompanied by strict enforcement and disclosure standards.

From a policy standpoint, these insights support the need for stronger regulatory mechanisms, including standardized disclosure frameworks that limit selective reporting. Regulatory bodies should consider requiring companies to report quantitative environmental targets and follow-up metrics over time. Furthermore, integrating independent third-party audits into sustainability reporting can improve the credibility of corporate disclosures and deter greenwashing.

Lastly, the regression findings linking greenwashing levels to specific financial indicators (such as high debt or large market capitalization) offer actionable insights for investors and watchdog organizations. These features can be used as risk signals to prioritize

which firms require closer ESG scrutiny, guiding both investment decisions and regulatory audits.

## References


Australasian Centre for Corporate Responsibility. (2025, January 7). *Shell's LNG strategy under scrutiny as institutional investors file shareholder resolution*. ACCR. https://www.accr.org.au/news/shell's-lng-strategy-under-scrutiny-as-institutional-investors-file-shareholder-resolution/

Boon, M. (2019). A climate of change? The oil industry and decarbonization in historical perspective. *Business History Review*, *93*(1), 101–125. Scopus. https://doi.org/10.1017/S0007680519000321

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., … Amodei, D. (2020). Language Models are Few-Shot Learners. *CoRR, abs/2005.14165*. https://arxiv.org/abs/2005.14165

Cohen, L., Gurun, U. G., & Nguyen, Q. H. (2020). *The ESG-Innovation Disconnect: Evidence from Green Patenting* (Working Paper 27990). National Bureau of Economic Research. https://doi.org/10.3386/w27990

de Freitas Netto, S. V., Sobral, M. F. F., Ribeiro, A. R. B., & Soares, G. R. L. (2020). Concepts and forms of greenwashing: A systematic review. *Environmental Sciences Europe*, *32*(1). Scopus. https://doi.org/10.1186/s12302-020-0300-3

Garrido-Merch'an, E. C., Gonz'alez-Barthe, C., & Vaca, M. C. (2023). Fine-tuning ClimateBert transformer with ClimaText for the disclosure analysis of climate-related financial risks. *ArXiv, abs/2303.13373*. https://api.semanticscholar.org/CorpusID:257687513

Gielen, D., Boshell, F., Saygin, D., Bazilian, M. D., Wagner, N., & Gorini, R. (2019). The role of renewable energy in the global energy transformation. *Energy Strategy Reviews*, *24*, 38–50. https://doi.org/10.1016/j.esr.2019.01.006

Global Witness. (2023, February 15). *Fossil fuel greenwash*. Global Witness. https://globalwitness.org/en/campaigns/fossil-fuels/fossil-fuel-greenwash/

Grushecky, S. T., Harris, K. J., Strager, M. P., Wang, J., & Mesa, A. N. (2022). Land Cover Change Associated with Unconventional Oil and Gas Development in the Appalachian Region. *Environmental Management*, *70*(5), 869–880. Scopus. https://doi.org/10.1007/s00267-022-01702-y

InfluenceMap. (2025, September 1). *Shell shareholders file resolution asking oil giant for transparency on LNG strategy, climate goals*. ESG Dive. https://www.esgdive.com/news/shell-shareholders-file-resolution-asking-for-transparency-on-lng-strategy-climate-goals/736919/

Johnston, M. A., Nuttall, M. F., Hickerson, E. L., O'Connell, K., Blakeway, R. D., Embesi, J. A., MacMillan, J., Peter, D., & Schmahl, G. P. (2022). Characterizing the Biological Community before and after Partial Removal of an Offshore Gas Platform in the Northwestern Gulf of Mexico. *Environmental Management*, *70*(6), 1078–1092. Scopus. https://doi.org/10.1007/s00267-022-01714-8

Kim, E.-H., & Lyon, T. P. (2015). Greenwash vs. Brownwash: Exaggeration and undue modesty in corporate sustainability disclosure. *Organization Science*, *26*(3), 705–723. Scopus. https://doi.org/10.1287/orsc.2014.0949

Li, M., Trencher, G., & Asuka, J. (2022a). The clean energy claims of BP, Chevron, ExxonMobil and Shell: A mismatch between discourse, actions and investments. *PLoS ONE*, *17*(2 February). Scopus. https://doi.org/10.1371/journal.pone.0263596

Li, M., Trencher, G., & Asuka, J. (2022b). The clean energy claims of BP, Chevron, ExxonMobil and Shell: A mismatch between discourse, actions and investments. *PLoS ONE*, *17*(2 February). Scopus. https://doi.org/10.1371/journal.pone.0263596

Li, W., Xia, Y., Li, N., Chang, J., Liu, J., Wang, P., & He, X. (2024). Corrigendum to "Temporal assembly patterns of microbial communities in three parallel bioreactors treating low-concentration coking wastewater with differing carbon source concentrations" [Journal of Environmental Sciences, Volume 137, March 2024, Pages 455–468] (Journal of Environmental Sciences (2024) 137 (455–468), (S1001074223001146), (10.1016/j.jes.2023.03.005)). *Journal of Environmental Sciences (China)*, *137*, 701–702. Scopus. https://doi.org/10.1016/j.jes.2023.03.034

Lyon, T. P., & Maxwell, J. W. (2011). Greenwash: Corporate environmental disclosure under threat of audit. *Journal of Economics and Management Strategy*, *20*(1), 3–41. Scopus. https://doi.org/10.1111/j.1530-9134.2010.00282.x

Lyon, T. P., & Montgomery, A. W. (2015). The Means and End of Greenwash. *Organization and Environment*, *28*(2), 223–249. Scopus. https://doi.org/10.1177/1086026615575332

Marquis, C., Toffel, M. W., & Zhou, Y. (2016). Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, *27*(2), 483–504. Scopus. https://doi.org/10.1287/orsc.2015.1039

Pickl, M. J. (2019). The renewable energy strategies of oil majors – From oil to energy? *Energy Strategy Reviews*, *26*. Scopus. https://doi.org/10.1016/j.esr.2019.100370

Pierre, J. P., Andrews, J. R., Young, M. H., Sun, A. Y., & Wolaver, B. D. (2020). Projected Landscape Impacts from Oil and Gas Development Scenarios in the Permian Basin, USA. *Environmental Management*, *66*(3), 348–363. Scopus. https://doi.org/10.1007/s00267-020-01308-2

Rowlands, I. H. (2000). Beauty and the beast? BP's and Exxon's positions on global climate change. *Environment and Planning C: Government and Policy*, *18*(3), 339–354. Scopus. https://doi.org/10.1068/c9752

Ruiz-Blanco, S., Romero, S., & Fernandez-Feijoo, B. (2022). Green, blue or black, but washing–What company characteristics determine greenwashing? *Environment, Development and Sustainability*, *24*(3), 4024–4045. Scopus. https://doi.org/10.1007/s10668-021-01602-x

Supran, G., & Oreskes, N. (2017). Assessing ExxonMobil's climate change communications (1977-2014). *Environmental Research Letters*, *12*(8). Scopus. https://doi.org/10.1088/1748-9326/aa815f

Supran, G., & Oreskes, N. (2021). Rhetoric and frame analysis of ExxonMobil's climate change communications. *One Earth*, *4*(5), 696–719. Scopus. https://doi.org/10.1016/j.oneear.2021.04.014

Taylor, M., & Watts, J. (2019, October 9). Revealed: The 20 firms behind a third of all carbon emissions. *The Guardian*. https://www.theguardian.com/environment/2019/oct/09/revealed-20-firms-third-carbon-emissions

UNFCCC. (2015). *Paris Agreement. United Nations Framework Convention on Climate Change*. https://unfccc.int/sites/default/files/english_paris_agreement.pdf

Yang, Z., Nguyen, T. T. H., Nguyen, H. N., Nguyen, T. T. N., & Cao, T. T. (2020). Greenwashing behaviours: Causes, taxonomy and consequences based on a systematic literature review. *Journal of Business Economics and Management*, *21*(5), 1486–1507. Scopus. https://doi.org/10.3846/jbem.2020.13225

Zhang, D. (2022). Green financial system regulation shock and greenwashing behaviors: Evidence from Chinese firms. *Energy Economics*, *111*. Scopus. https://doi.org/10.1016/j.eneco.2022.106064

Zhang, L., Xu, M., Chen, H., Li, Y., & Chen, S. (2022). Globalization, Green Economy and Environmental Challenges: State of the Art Review for Practical Implications. *Frontiers in Environmental Science*, *10*. Scopus. https://doi.org/10.3389/fenvs.2022.870271

Zhong, M., & Bazilian, M. D. (2018). Contours of the energy transition: Investment by international oil and gas companies in renewable energy. *Electricity Journal*, *31*(1), 82–91. Scopus. https://doi.org/10.1016/j.tej.2018.01.001

# Appendix A: Market Data Feature Descriptions

Table B.1. Market data features, along with a description of what each feature represents

| **Feature** | **Description** | **Feature Type** |
|---|---|---|
| Country | ISO standard country code where the company's headquarters are located | Categorical |
| Employees | The total number of employees working for the company, reported annually, and measured in thousands. | Numerical |
| Assets | The total value of the company's assets as reported on the balance sheet, measured in millions. | Numerical |

| | | |
|---|---|---|
| Revenue | Represents the company's total income from sales and operating revenues, excluding costs, measured in millions. | Numerical |
| EBITDA | A measure of the company’s profitability from core operations, excluding the effects of capital structure, tax rates, and depreciation policies, reported in millions | Numerical |
| EBIT | Another profitability measure that includes depreciation and amortization but excludes interest and taxes, reported in millions. | Numerical |
| Income | The total income or loss after all expenses, taxes, reported in millions. | Numerical |

| | | |
|---|---|---|
| Total Debt | The total long-term debt of the company, including current portions, provides a view of the company's leverage and financial obligations. | Numerical |
| Total Parent Stockholders' Equity | Represents the equity held by common and preferred shareholders, reflecting the net worth of the company from shareholders' perspectives. | Numerical |
| Return on Equity (ROE) | A ratio that measures profitability relative to shareholders' equity, indicating how effectively the company is using its equity base to generate profits. | Numerical |
| Debt to Equity Ratio (D/E) | A leverage ratio that compares the company's total debt to its total shareholders' equity, indicating the balance between debt and equity financing. | Numerical |

| | | |
|---|---|---|
| Market Capitalization | The total market value of the company's outstanding shares provides a quick measure of its size and market value. | Numerical |

# Appendix B: Full Set of Partial Dependence Plots (PDPs)

This appendix presents the full versions of Figures B.1–B.3, showing the marginal effects of 13 features on greenwashing scores across the Sustainability, Annual, and 10-K datasets. These plots support the feature importance analysis in Chapter 5.

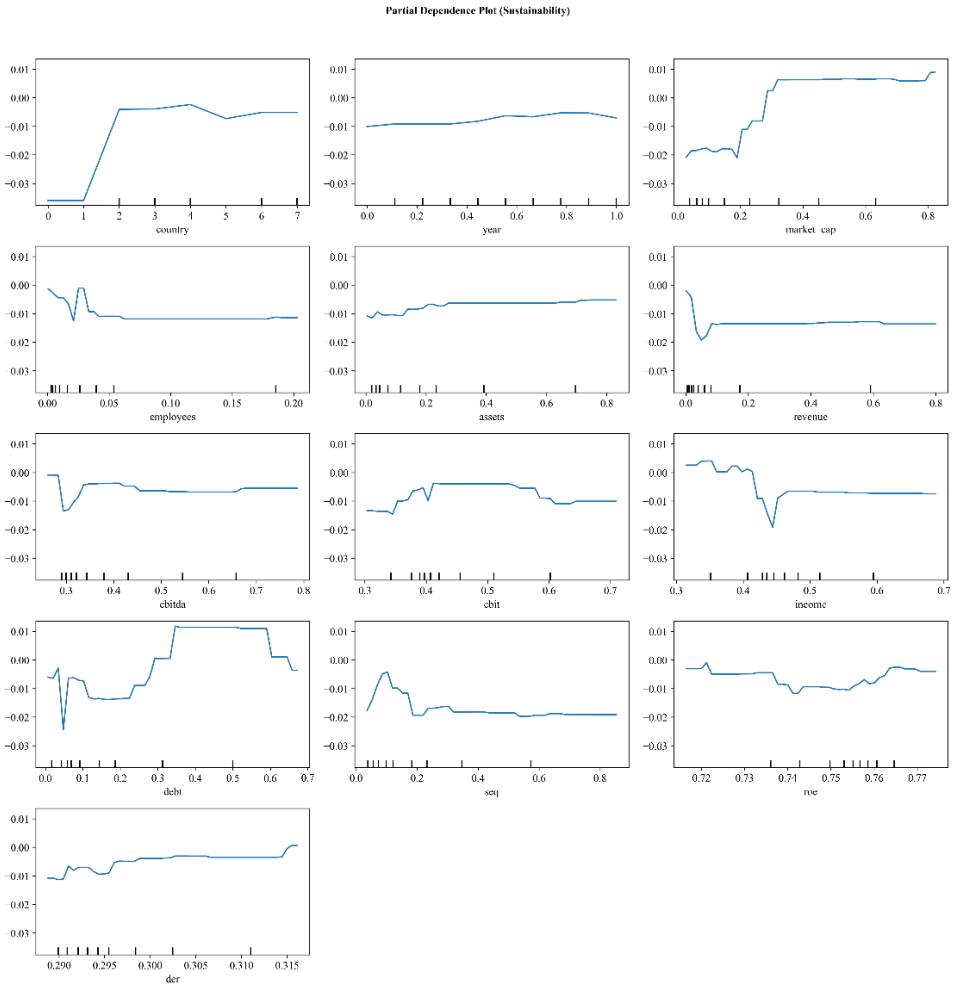



Figure B.1. Partial Dependence Plots for the 13 independent features and the corresponding dependent feature, greenwashing, on the Sustainability dataset

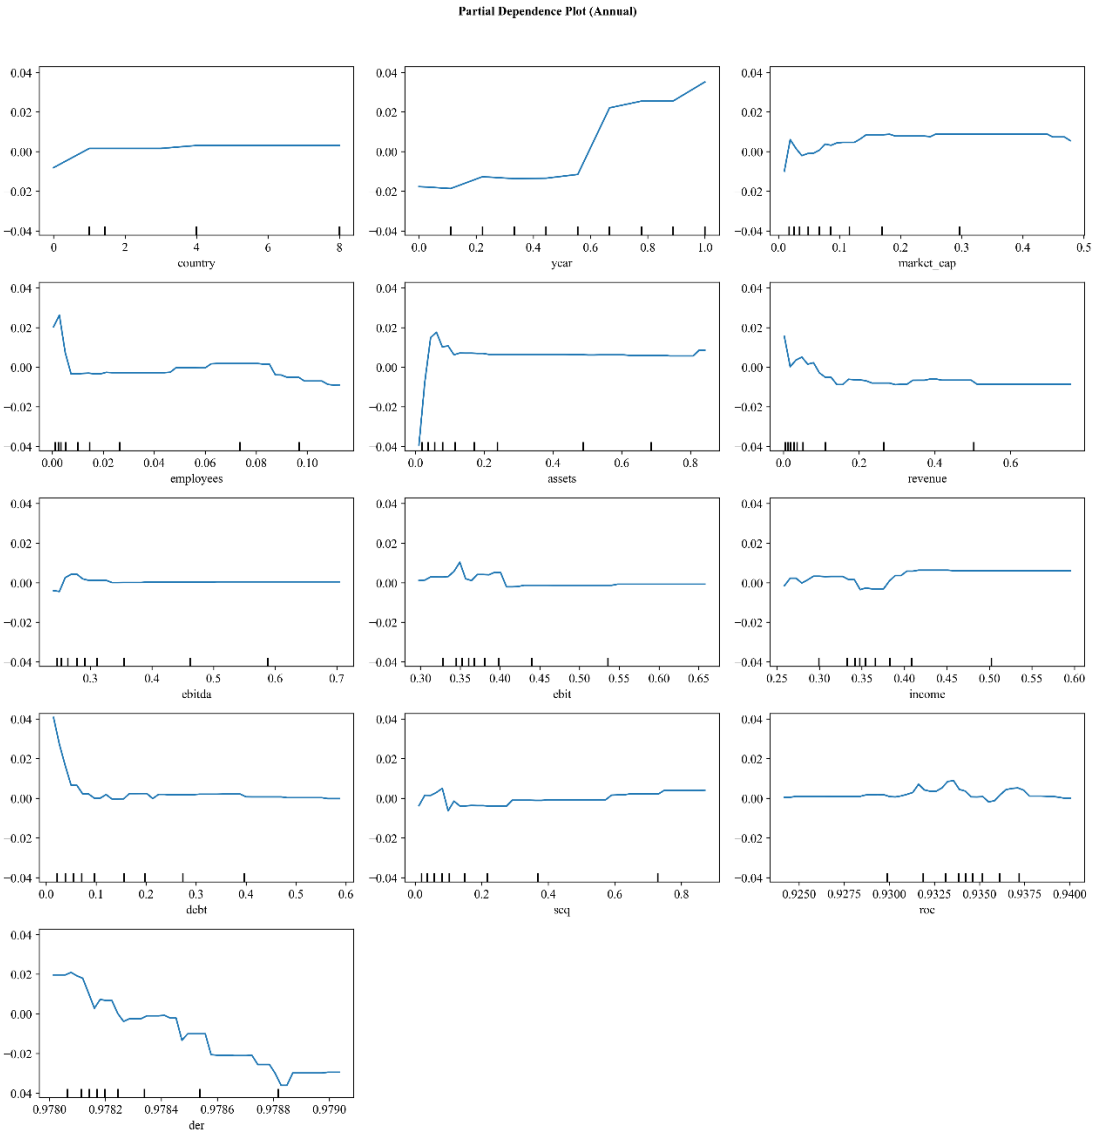



Figure B.2. Partial Dependence Plots for the 13 independent features and the corresponding dependent feature, greenwashing, on the Annual dataset.

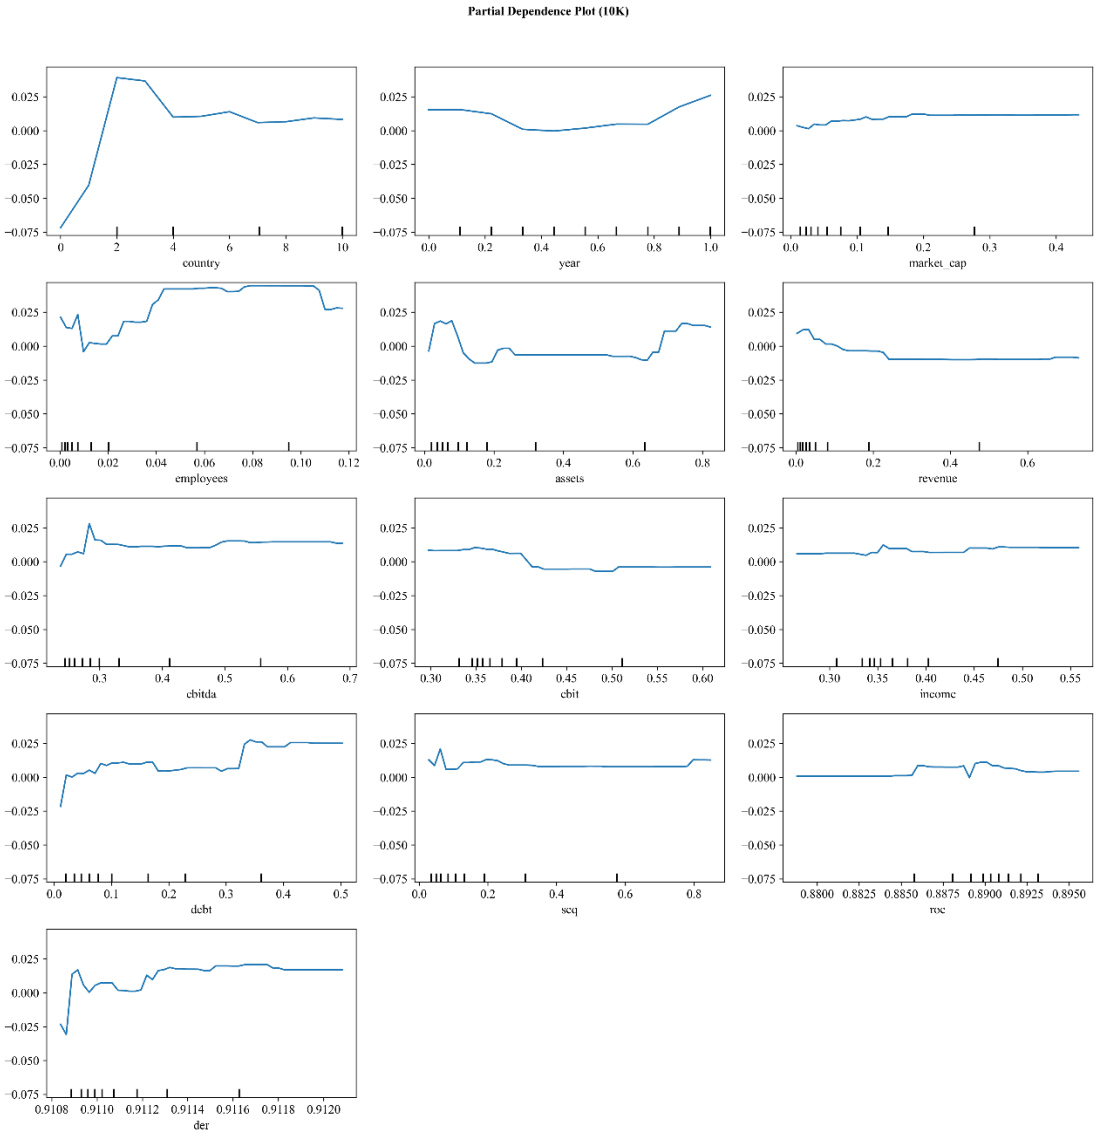



Figure B.3. Partial Dependence Plots for the 13 independent features and the corresponding dependent feature, greenwashing, on the 10K dataset.