



RESEARCH ARTICLE OPEN ACCESS

The Language of Greenwashing: SDG Omission and Opportunity-Oriented Environmental Tone as Alert Metrics in Green Bond Disclosures

Andrea Nicolodi^{1,2}  | Sandra Paterlini¹  | Monica Gentile³ | Vincenzo Foglia Manzillo³ | Gianluca Vittorioso³

¹University of Trento, Trento, Italy | ²Free University of Bolzano-Bozen, Bolzano, Italy | ³CONSOB, Rome, Italy

Correspondence: Andrea Nicolodi (andrea.nicolodi-1@unitn.it)

Received: 19 December 2025 | **Revised:** 17 May 2026 | **Accepted:** 25 May 2026

Keywords: AI | green bonds | greenwashing alert | NLP | SDG omission | sentiment analysis

ABSTRACT

Green bonds play a central role in sustainable finance, yet concerns about greenwashing raise questions about the credibility of issuers' sustainability disclosures. Using dictionary-based methods and domain-specific BERT transformer models, this paper proposes two greenwashing alert metrics and investigates their performance by analyzing sustainability reports of European corporate green bond issuers from 2019 to 2023. First, we develop the SDGs Omission Index (SDGOI), which measures the discrepancy between Sustainable Development Goals (SDGs) declared ex-ante in green bond frameworks and those subsequently reported in sustainability disclosures. We show that the use of SDG-specific language is associated with broader SDG coverage, whereas generic environmental language is not. Second, we introduce the Environmental Sentiment Metric (ESM), capturing opportunity-oriented environmental sentiment. We find that a more opportunity-oriented tone is positively related to SDG omission, ESG controversies, and greenwashing accusations. Together, SDGOI and ESM provide interpretable, disclosure-based indicators that can support issuing greenwashing alerts in the European green bond market.

JEL Classification: G34, G38, J33, K22, M52

1 | Introduction

Green bonds help finance the transition to a sustainable economy by channeling capital into environmentally beneficial projects in line with the Paris Agreement and the United Nations Sustainable Development Goals (SDGs) (Flammer 2021; Tang and Zhang 2020). To ensure transparency and impact, green bond proceeds should be allocated to eligible projects and accompanied by credible reporting on the use of proceeds and (where feasible) impact metrics (International Capital Market Association [ICMA] 2022). In 2024, green bonds emerged as the dominant category within the sustainable finance market, with issuances reaching USD 381 billion in the corporate sector (OECD 2025). Until the adoption of Regulation (EU) 2023/2631 on European

green bonds, the EU lacked a unified standard for certifying green financial instruments. Issuers relied on voluntary frameworks, such as the Green Bond Principles (ICMA) and the Climate Bonds Initiative (CBI). While these frameworks are widely recognized, concerns about greenwashing, a sustainability-related communication that creates an overly positive impression relative to underlying practices, have become central (Delmas and Burbano 2011; Lyon and Montgomery 2015). Common signs include vague claims, selective disclosure, and the omission of less favorable information, consistent with evidence on impression management in sustainability reporting (Marquis et al. 2016; Cho, Michelon, and Patten 2012). From a corporate social responsibility perspective, such practices undermine stakeholder trust and can distort capital allocation (Christensen et al. 2021,

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Corporate Social Responsibility and Environmental Management* published by ERP Environment and John Wiley & Sons Ltd.

Lagasio 2024, Liu et al. 2024), nevertheless greenwashing has transitioned from a theoretical concern into a systemic market risk. The European Securities and Markets Authority (ESMA, 2024) confirmed that greenwashing risks are growing across the sustainable investment value chain, identifying omissions, cherry-picking, and the use of “vague or aspirational language” as the primary drivers of misleading sustainability claims. Reflecting this trend, the European Banking Authority (EBA 2024) reports a 26.1% rise in alleged cases in the EU in 2023 versus 2022.

The academic literature has made important progress in conceptualizing greenwashing and ESG-washing, for example by comparing self-reported ESG disclosure to external ratings or environmental outcomes. However, empirical research on systematic *text-based* tools for issuing alert for potential greenwashing remains comparatively scarce. Two gaps are particularly salient. First, existing measures rarely focus on the omission of SDG-related information, despite the centrality of SDGs in the current sustainable finance architecture. Second, while several studies employ sentiment analysis, there is limited evidence on whether an opportunity-oriented environmental tone in sustainability reports can serve as a stand-alone alert indicator of greenwashing risk.

This study addresses these gaps by developing disclosure-based alert mechanisms for greenwashing in the context of European corporate green bond issuance. We propose an approach that combines Natural Language Processing (NLP) techniques with domain-specific transformer models based on BERT, including ClimateBERT (Bingler et al. 2022, 2024; Webersinke et al. 2022) and ESGBERT (Schimanski et al. 2024). In doing so, we rely exclusively on textual information extracted from sustainability reports, without imposing any ex-ante assumptions on actual environmental performance. A proprietary SDG-aligned dictionary is used to map targeted keywords to specific SDGs, leveraging the established alignment between the Green Bond Principles (GBPs) and the SDGs framework and building upon the prototype developed by CONSOB and the University of Trento (Paterlini et al. 2025).

Specifically, the study seeks to answer the following research questions. First, how can the omission of SDGs in sustainability disclosures be systematically analyzed and linked to factors such as type of sentence, sentiment, and other relevant textual characteristics in order to develop potential greenwashing alerts in green bond issuances? Second, can the tone and sentiment of sustainability-related disclosures, quantified through domain-specific BERT-based models, be leveraged to develop alert metrics for greenwashing risk independently of SDG-related disclosure content?

This paper makes three key contributions to the literature on greenwashing, corporate social responsibility, and sustainable finance. First, we shift the focus from the traditional ‘talk versus walk’ discrepancy toward a granular ‘talk-only’ analysis of corporate narratives, not relying on third-party scores to build alert.

Second, we introduce the SDGs Omission Index (SDGOI), a novel metric designed to help investors, market participants, and regulators assess SDG selective reporting practices.

Third, the study develops the Environmental Sentiment Metric (ESM) as a standardized alert indicator, showing a positive association between an excessive opportunity-oriented environmental tone and both ESG Controversies (ESGCs) and greenwashing accusations (GWAs). Both metrics are not aimed to prove or disprove the existence of greenwashing in general; rather, the purpose is to propose helpful tools to flag potentially problematic cases for further scrutiny.

The paper is organized as follows. Section 2 reviews the related literature. Section 3 presents the data. Section 4 introduces the SDGOI, develops the hypotheses, and reports empirical results on SDGs omission. Section 5 then introduces the ESM and investigates its association with ESG controversies and greenwashing accusations. Section 6 concludes.

A list of acronyms is provided in [Appendix](#).

2 | Literature Review

Corporate green bonds are an important focus for empirical research because they represent a rapidly expanding segment of sustainable finance and provide a direct mechanism for channeling capital toward environmentally beneficial projects. Flammer (2021) finds that investors react positively to the announcement of green bond issuances, with stronger reactions for first-time issuers and for bonds certified by external parties. After issuing green bonds, firms improve their environmental performance (such as achieving higher environmental ratings and reducing CO₂ emissions) and attract a greater share of long-term and environmentally oriented investors. Tang and Zhang (2020) provide a large-scale empirical investigations into whether corporate green bond issuance generates measurable benefits for shareholders. Using a comprehensive dataset covering firms across 28 countries from 2007 to 2017, the authors document that stock prices respond positively to green bond issuance announcements, indicating that markets generally perceive these instruments as value-enhancing corporate actions. Despite this favorable market reaction, the study finds no consistently significant yield premium on green bonds relative to conventional bonds. This implies that the positive shareholder response is not primarily driven by lower cost of debt, but, instead, reflects broader investor perceptions regarding the firm’s environmental positioning and credibility. Empirical evidence increasingly highlights that green bond issuance is often associated with improvements in firms’ environmental reporting quality, suggesting that these instruments can strengthen transparency and attract environmentally focused investors. For instance, Liu et al. (2025) document that corporate environmental disclosure quality rises substantially (by nearly 15%) after the issuance of green bonds, particularly when bonds are externally certified or issued for the first time, underscoring their potential role in reinforcing environmental governance mechanisms. The increasing prevalence of sustainability-oriented corporate communication has intensified concerns regarding the authenticity of environmental claims, a phenomenon widely known as greenwashing (Delmas and Burbano 2011). Greenwashing is a relevant phenomenon not only in equity markets but also in corporate debt, given its implications for financing costs and investor trust (Peng and Xie 2024; Roggi et al. 2024). Despite the growing relevance of

sustainability reporting, the empirical identification and quantification of greenwashing remain methodologically challenging and conceptually fragmented. The absence of a universally accepted definition, coupled with its multifaceted nature, has led to a proliferation of measurement approaches. Existing methods range from ESG disclosure–performance comparisons (e.g., Cho, Guidry, et al. 2012) to perception-driven proxies (e.g., Testa et al. 2018) and sentiment-based evaluations (e.g., Arena et al. 2015), highlighting the complexity of assessing the gap between symbolic and substantive sustainability actions (Lublov et al. 2025). Within this debate, green bonds have attracted significant attention as emblematic instruments of sustainable finance. Ge et al. (2025) examine the impact of green bond issuance on the ESG performance of Chinese listed firms, finding improvements in environmental and social dimensions. The study also uncovers evidence suggestive of greenwashing: improvements do not differ significantly across firms with varying pollution intensity or political ties, raising concerns about the authenticity of ESG enhancements. These findings align with broader critiques that disclosures, ratings, and voluntary statements, being largely self-reported, are prone to strategic manipulation (e.g., Cho et al. 2015; Christensen et al. 2021; Dyck et al. 2019). A vast body of research emphasizes this disclosure–performance misalignment. Walker and Wan (2012) show that firms often engage in symbolic sustainability communication to protect legitimacy without substantive environmental change, ultimately undermining stakeholder trust and financial performance. Similarly, Yu et al. (2020) develop a peer-relative greenwashing score based on discrepancies between ESG disclosure and actual performance. Their findings reveal that large firms, despite extensive reporting, are particularly exposed to accusations of symbolic disclosure due to visibility and stakeholder pressure. Governance mechanisms, such as board independence and institutional ownership, can mitigate these risks, while external scrutiny, including foreign listings, reduces greenwashing incidence.

Another strand of literature highlights that omission patterns and opportunity-oriented environmental tone are consistent indicators of disclosure credibility risk. Poiriazzi et al. (2025) show through a bibliometric analysis that greenwashing weakens the credibility of ESG disclosures and destabilizes decision-making in sustainable finance, particularly in the context of green bonds. Complementing this, Feghali et al. (2025) provide a comprehensive systematic review identifying the strategic narrative tactics firms use to shape overly positive sustainability perceptions, emphasizing how such practices hinder regulatory and market efforts to curb misleading claims. Additional methodological evidence comes from Bernini et al. (2024), who highlight the difficulties of empirically measuring greenwashing and document the widespread reliance on symbolic and narrative-driven environmental communication. From a marketing and communication perspective, Persakis et al. (2025) illustrate how opportunity-oriented messaging shapes consumer perceptions and reinforces deceptive environmental positioning strategies, confirming the broader role of tone and sentiment in greenwashing dynamics.

Recent conceptual contributions expand the scope from “greenwashing” to “ESG-washing”. Todaro and Torelli (2024), for example, highlight how firms strategically appropriate ESG and

circular economy terminology to construct reputational legitimacy in the absence of substantive alignment. This reinforces the argument that greenwashing is multidimensional, spanning environmental, social, and governance dimensions. At the same time, systematic reviews confirm both conceptual fragmentation and methodological diversity in this field. Ziolo et al. (2024) stress that reliance on theoretical frameworks such as stakeholder theory, legitimacy theory, and institutional theory is frequent but inconsistent, which complicates generalization and comparability.

Another important strand of literature critiques the reliability of ESG ratings. Since ratings rely heavily on self-reported disclosures, they are exposed to strategic manipulation. Moreover, methodological discrepancies among agencies further complicate their interpretation. Drempetic et al. (2020) show that firm size inflates ESG scores due to better disclosure capabilities rather than superior sustainability. This concern is reinforced by Kathan et al. (2025), who disentangle apparent from real environmental performance in European listed firms. While ESG scores correlate with narrative disclosure (apparent performance), they negatively associate with objective ecological outcomes, particularly among large firms. Treepongkaruna et al. (2024) reach similar conclusions in the U.S. context, showing that high ESG scores are not associated with lower carbon emissions, suggesting that ratings often function as reputational tools rather than performance indicators. These documented limitations of ESG scores have motivated the search for alternative measures of sustainability performance and misrepresentation.

To address these limitations, scholars increasingly turn to natural language processing (NLP) and sentiment-based approaches. These methods allow systematic analysis of textual data such as sustainability reports, press releases, earnings calls, and risk factor disclosures. By leveraging sentiment analysis, topic modeling, and semantic similarity, researchers can identify discrepancies between corporate narratives and operational outcomes.

Building on Loughran and McDonald (2016), recent sustainability literature uses sentiment dictionaries and machine learning tools to measure the tone, polarity, and readability of ESG communication. Studies such as Luccioni et al. (2020), Bingler et al. (2022), Bingler et al. (2024), and Moodaley and Telukdarie (2023) apply NLP-based metrics and transformer models to detect divergences between corporate claims and actual performance, demonstrating their value as early warning signals. Conceptual frameworks such as that of Dorfleitner and Utz (2024) explicitly operationalize greenwashing as the gap between “green talk” and “green walk”, while Lagasio (2024) propose the ESG-Washing Severity Index, which combines sentiment metrics with third-party ESG ratings to detect symbolic disclosure. Empirical evidence suggests that larger firms exhibit higher ESGSI scores, consistent with reputational incentives. Methodologically, this strand of work underscores the potential of advanced techniques for scrutinizing sustainability reports and uncovering patterns of strategic communication.

Recent applications further highlight the strategic use of textual tone. Gorovaia and Makrominas (2024) show that U.S. firms subject to environmental penalties expand the length and positivity of CSR reports while reducing transparency, suggesting

deliberate obfuscation. Similarly, bibliometric studies (e.g., Sneideriene and Legenzova 2025) confirm that AI-based approaches are becoming a dominant frontier in detecting reputational manipulation. Importantly, these tools align with broader academic efforts to move from perception-based to verifiable, impact-oriented metrics. Yet, the scarcity of confirmed green-washing cases remains a key challenge in validating proposed indicators.

3 | Data

3.1 | Green Bond Data

Our empirical analysis is conducted in the European green bond market. We begin by clarifying the operational definition of a “green” bond. To ensure a robust classification, we rely on two independent data providers: LSEG and FactSet. LSEG applies the standards established by the International Capital Market Association (ICMA), whereas FactSet identifies green bonds by examining disclosures in the use-of-proceeds section of the bond prospectus or in the final terms and conditions of the issuance.

On the basis of this classification, issuers were ranked by the number of green bonds issued between 2013 and 2023. This starting point was chosen because prior to 2013 the issuance of green bonds was extremely limited, which renders it negligible for research purposes. Once issuers were sorted according to their issuance activity, their sustainability reports were retrieved through LSEG and supplemented with targeted manual web searches, restricted to the period 2019–2023. This temporal choice aligns with the evolution of the EU’s regulatory framework following the 2018 ‘Action Plan: Financing Sustainable Growth.’¹ This initiative marked the beginning of a structured regulatory framework, making the subsequent period the most relevant for assessing disclosure practices and alignment with sustainability standards.

We exhaustively include all firms with seven or more green bond issuances. The remaining population was stratified into two groups (4–6 bonds and ≤ 3 bonds) and randomly sampled in proportion to their prevalence within the lower-issuance population. These proportions were applied exclusively during document retrieval; issuers without an accessible or English-language sustainability report were not replaced within their stratum. This language restriction is driven by the algorithm used for the analysis (see Section 3.2), which is pre-trained for English. While this selection may over-represent internationally exposed issuers with higher disclosure capabilities, it implies that our results are conservative: if disclosure anomalies are detectable in a population of structurally advantaged, high-quality English reports, they may be even more pronounced in less standardized reporting contexts. Although this method departs from the original issuance distribution, due to the complete inclusion of high-frequency issuers, it ensures broad market coverage, capturing approximately 70% of the total issued green-bond volume in the period. This approach prioritizes representativeness in terms of issuance magnitude rather than issuer count and follows standard stratified sampling principles for skewed distributions (Tillé and Wilhelm 2017). The final dataset consists of 898 reports from 195 unique issuers, comprising 396 Sustainability Reports, 310 Other

Non-Financial Reports, 77 Annual & Sustainability Reports, and 115 Annual Reports which also include sustainability information. For simplicity, the term “sustainability reports” will be used henceforth to refer to the documents analyzed. The distribution of issuers’ domicile and sector by year is reported in Tables 1 and 2.

3.2 | NLP Methodology

To ensure the reliability of our textual analysis, the NLP pipeline builds on domain-adapted transformer models developed for

TABLE 1 | Sector distribution per year.

Sector	2019	2020	2021	2022	2023
Agency	1	1	1	1	1
Banks	48	52	56	57	54
Electric power	18	21	23	22	21
Energy company	4	4	4	4	4
Gas distribution	2	2	2	2	2
Manufacturing	36	41	42	42	39
Other financial	26	34	34	35	30
Service company	15	18	19	19	19
Communication	4	4	5	5	5
Transportation	2	2	4	5	5
Total	156	179	190	192	180

Note: Green bonds, which belong to the ‘Agency’ sector on the basis of LSEG, are classified as corporate bonds by FactSet, and for this reason they are included.

TABLE 2 | Domicile distribution per year.

Domicile	2019	2020	2021	2022	2023
AT	5	6	7	7	7
BE	3	4	5	5	4
CZ	0	0	0	1	1
DE	29	29	31	30	29
DK	8	8	8	8	8
ES	11	12	12	13	12
FI	7	7	7	7	7
FR	19	24	28	27	27
GR	5	5	5	5	3
HU	1	1	1	1	1
IE	2	5	5	5	4
IT	18	20	20	20	18
LU	3	4	4	4	4
NL	12	17	16	18	17
PL	3	4	4	4	3
PT	1	1	1	1	0
SE	29	31	35	35	34
SI	0	1	1	1	1
Total	156	179	190	192	180

ESG and climate-related disclosure analysis. The procedure involves multiple stages.

First, we use ESGBERT, specifically the ESGBERT/EnvironmentalBERT-environmental classifier, to identify environmentally related sentences in sustainability reports (label: “*environmental*” or “*none*”). The model is based on ESG-related corporate text and is fine-tuned on expert-labelled ESG datasets, achieving an accuracy of 95.0% for the environmental disclosure classifier (Schimanski et al. 2024).

Second, we rely on the ClimateBERT environmental-claims classifier to identify environmental claims in the reports (label: “*environmental claim*” or “*none*”).

An environmental claim is defined as a statement that conveys, explicitly or implicitly, an environmental benefit, commitment, or impact. It therefore differs from a general environmental sentence, which may merely mention an environmental topic without making a substantive claim.

This model is fine-tuned on an expert-annotated dataset of environmental claims drawn from corporate sustainability reports, annual reports, and earnings call transcripts. In the corresponding validation exercise, the ClimateBERT model achieves an F1-score of 0.838 and an accuracy of 0.909 on the test set for environmental claim detection (Stammach et al. (2023)).

Third, climate-related sentiment labels are assigned using ClimateBERT, which is built on the same ClimateBERT backbone but fine-tuned on a different labelled dataset for climate-sentiment classification enabling to classify text as emphasizing risks (label: *risk*), opportunities (label: *opportunity*), or neutral climate-related content (label: *neutral*).

Since our corpus consists of English-language sustainability reports that is, the same type of corporate disclosure domain on which these models were trained and validated, we rely on the external expert-annotation validation reported in the original studies rather than conducting an additional manual validation exercise.

Finally, SDG-related phrases are identified through a dictionary-based search in the sustainability report (i.e., *SDGs Found*) and compared with the a priori declared SDGs (i.e., *SDGs Declared*) of active green bonds in the reference year.

Our SDG dictionary follows the methodology described in Paterlini et al. 2025, where SDG-related keywords are constructed by mapping the Green Bond Principles to SDGs and integrating this mapping with the LSEG Use-of-Proceeds taxonomy. Once the dictionary of keywords associated with each SDG was defined, we searched the sustainability documents to identify which SDGs from the set declared *in advance* (i.e., in the green bond documentation) could be found in the reports. The NLP pipeline is deterministic and rule-based for SDG extraction, and transformer-based for environmental content, claims, and sentiment (see Paterlini et al. 2025 for details).

3.3 | Variables

Based on the NLP-based methodology, we construct a set of text-based variables to quantify both the presence and the tone of content related to environmental themes and the SDGs. Specifically, we compute the proportion of environmentally labeled sentences relative to the total number of sentences in each document (EnvContent_Of_Total), as well as the proportion of environmentally labeled claims (EnvClaims_Of_Total). In parallel, we calculate analogous metrics for SDG-related content, including the proportion of SDG-themed sentences (SDGContent_Of_Total) and the proportion of SDG-themed claims (SDGClaims_Of_Total).

To assess the tone of the content, we further compute the share of environmentally labeled sentences that are associated with opportunity-oriented or risky sentiment (Sent_Opp_Of_EnvContent and Sent_Risk_Of_EnvContent, respectively). Similarly, we derive the corresponding proportions for SDG-related sentences (Sent_Opp_Of_SDGContent and Sent_Risk_Of_SDGContent).

In addition to these text-based indicators, we incorporate a structural variable: the number of green bonds actively outstanding as of December 31 of the reporting year (Active_Bonds). While the relationship between the volume of outstanding bonds and the quality of sustainability disclosure has received limited attention in prior research, it is reasonable to expect that firms with a larger presence in the green bond market are subject to greater scrutiny by both investors and regulators. This increased visibility may, in turn, create stronger incentives for issuers to enhance the comprehensiveness, transparency, and reliability of their sustainability reporting. We refer to Table 3 for variables definition and Table 4 for the correlation matrix among variables.

4 | SDGs Omission: Index, Hypotheses and Empirical Results

This section introduces the SDGs Omission Index (SDGOI), develops hypotheses on its relation to content and sentiment, and presents empirical evidence from a fixed-effects panel model.

4.1 | SDGs Omission Index and Hypotheses

We focus on two aspects of sustainability communication: SDG omission (i.e., misalignment between ex-ante declared SDG commitments and subsequent reporting) and the tone of sustainability narratives. We examine how textual tone and content features co-vary with SDG omission to identify communication patterns associated with larger SDG-related disclosure gaps, emphasizing how firms talk about sustainability rather than what they claim to achieve. Firms can strategically shape sustainability reports to steer perceptions and gain competitive advantages. In the European green bond context, market standards (e.g., GBPs) and regulations (e.g., the EU Taxonomy) encourage issuers to state which SDGs bond proceeds support. A lack of clear references to those stated SDGs in shareholder reporting can therefore be interpreted as a greenwashing alert, raising concerns about

TABLE 3 | Variables definitions for the SDGOI fixed-effects panel model.

Variable	Definition
SDGOI	Sustainable Development Goals Omission Index, defined as $1 - \frac{SDGsFound}{SDGsDeclared}$, where SDGsFound are the distinct number of the 12 relevant SDGs identified by the NLP approach within the sustainability report and SDGsDeclared are the distinct number of the 12 SDGs reported by LSEG with respect to the use of proceeds.
EnvContent_of_Total	Share of sentences in the report that are labeled as environmental (by ESG-BERT) over the total number of sentences in the document.
EnvClaims_of_Total	Share of sentences labeled as environmental claims (by ClimateBERT) over the total number of sentences in the document.
SDGContent_of_Total	Share of sentences containing SDG-related content (identified via the SDG dictionary) over the total number of sentences in the document.
SDGClaims_of_Total	Share of sentences classified as SDG-related claims over the total number of sentences in the document.
Sent_Opp_of_EnvContent	Proportion of environmental sentences whose sentiment is labeled as Opportunity (by ClimateBERT), relative to all environmental sentences.
Sent_Risk_of_EnvContent	Proportion of environmental sentences whose sentiment is labeled as Risk, relative to all environmental sentences.
SDG_Sent_Opp_of_SDGContent	Proportion of SDG-related sentences whose sentiment is labeled as Opportunity, relative to all SDG-related sentences.
SDG_Sent_Risk_of_SDGContent	Proportion of SDG-related sentences whose sentiment is labeled as Risk, relative to all SDG-related sentences.
Active_Bonds	Number of green bonds issued by the firm that are actively outstanding as of 31 December of year t .

TABLE 4 | Correlation between variables (SDGOI).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1)	1.00	0.71	0.50	0.48	-0.01	-0.13	0.10	0.04
(2)	0.71	1.00	0.43	0.51	0.37	-0.32	0.12	-0.08
(3)	0.50	0.43	1.00	0.87	0.18	-0.09	0.17	-0.04
(4)	0.48	0.51	0.87	1.00	0.16	0.08	0.20	-0.06
(5)	-0.01	0.37	0.18	0.16	1.00	-0.45	0.18	-0.12
(6)	-0.13	-0.32	-0.09	0.08	-0.45	1.00	-0.01	0.25
(7)	0.10	0.12	0.17	0.20	0.18	-0.01	1.00	-0.16
(8)	0.04	-0.08	-0.04	-0.06	-0.12	0.25	-0.16	1.00

Note: Variables: (1) EnvContent_of_Total = share of environmental content in total text; (2) EnvClaims_of_Total = share of environmental claims in total text; (3) SDGContent_of_Total = share of SDG content in total text; (4) SDGClaims_of_Total = share of SDG claims in total text; (5) Sent_Opp_of_EnvContent = opportunity sentiment within environmental content; (6) Sent_Risk_of_EnvContent = risk-related sentiment within environmental content; (7) SDG_Sent_Opp_of_SDGContent = opportunity sentiment within SDG content; (8) SDG_Sent_Risk_of_SDGContent = risk-related sentiment within SDG content.

transparency. This omission may be masked by generic environmental language or opportunity-oriented framing, while downplaying or avoiding discussion of climate-related risks. This observation leads to our first two hypotheses:

Hypothesis 1. (Target information): A greater use of SDG-themed sentences is associated with broader coverage of the SDGs declared by issuers.

Hypothesis 2. (Dilution of information): A greater use of environmental sentences is associated with less coverage of the SDGs declared by issuers.

We therefore test the effects, on the omission of SDG information, of the concentration of SDG-themed phrases and environmental phrases. The interest of this analysis lies in the fact that companies with a high level of cheap talk prioritizes maintaining a positive public perception over making meaningful changes to their business practices (Bingler et al. 2024).

This research not only studies the content of the disclosure but also the way it is presented. In particular, the focus is on the sentiment of SDG and environmental themed sentences. While previous studies, such as Tsang et al. (2023) and Koelbl (2019), have shown that positive and negative sentiment types carry

meaningful information, they do not extend their analysis to the specific context of European green bonds. Here, we aim to test whether there is an association between coverage of SDG information and type of sentiment, distinguishing it as opportunity-oriented, risky, or neutral. Therefore, we test two additional hypotheses:

Hypothesis 3. (*SDG sentiment importance*): A higher presence of opportunity-oriented sentiment within SDG-related statements is associated with more coverage of the SDGs declared by the issuers.

Hypothesis 4. (*Environmental sentiment importance*): A higher presence of opportunity-oriented sentiment within environmental-related statements is associated with less coverage of the SDGs declared by the issuers.

To operationalize SDG omission, we define the SDGs Omission Index (SDGOI) as follows:

$$\text{SDGOI} = 1 - \frac{\text{SDGs Found}}{\text{SDGs Declared}} \quad (1)$$

where SDGsFound are the number of unique SDGs identified by the NLP approach within the sustainability report (see Section 3.2), while SDGsDeclared are the number of unique SDGs reported by LSEG with respect to the use of proceeds. The ratio indicates the proportion of declared SDGs that are identified in the reporting phase. Subtracting this ratio from one highlights the share of SDGs that remain unmentioned in the reports. Crucially, the algorithm tracks only the specific SDGs explicitly declared by the issuer at the green bond issuance stage. By focusing on these self-selected targets, the index measures consistency in communication regarding specific commitments, rather than adherence to external materiality frameworks. Since the issuers themselves identified these goals as the primary focus on their proceeds, it is reasonable to expect their inclusion in the corporate narrative regardless of sectoral heterogeneity. This design grounds the index in the theory of selective disclosure: by measuring the gap between ex-ante commitments and ex-post narrative verification, the SDGOI captures the information asymmetry that prevents investors from monitoring the use of proceeds. Thus, a high index value signals a strategic break in the “promise-verification” loop rather than a mere mechanical discrepancy.

In the following analysis, only SDGs [1,2,3,6,7,8,9,11,12,13,14,15] are considered, as they map directly into GBPs. SDG-related sentences are extracted using the dictionary-based approach described in Section 3.2, which links keywords to each SDG by mapping the GBPs to the SDGs based on the ICMA SDG Mapping Report.²

4.2 | Empirical Results on SDGs Omission

To test the hypothesis described above, we use a fixed-effects panel regression with heteroskedasticity-robust standard errors clustered at the firm level:

$$\text{SDGOI}_{i,t} = \beta_0 + \beta_1 \text{EnvContent.of.Total}_{i,t} + \beta_2 \text{EnvClaims.of.Total}_{i,t} + \beta_3 \text{SDGContent.of.Total}_{i,t} + \beta_4 \text{SDGClaims.of.Total}_{i,t}$$

$$\begin{aligned} &+ \beta_5 \text{Sent.Opp.of.EnvContent}_{i,t} + \beta_6 \text{Sent.Risk.of.EnvContent}_{i,t} \\ &+ \beta_7 \text{SDG.Sent.Opp.of.SDGContent}_{i,t} \\ &+ \beta_8 \text{SDG.Sent.Risk.of.SDGContent}_{i,t} \\ &+ \beta_9 \text{Active.Bonds}_{i,t} + \epsilon_{i,t} \end{aligned} \quad (2)$$

for $i = 1, \dots$, number of issuers and $t = 2019, \dots, 2023$, where we control for firm, year and sector fixed effects. Including sector fixed effects accounts for unobserved heterogeneity across entities that may influence the dependent variable but remain time-invariant. These controls are essential to mitigate omitted variable bias, as they absorb systematic differences between sectors that could otherwise confound the estimates (see Section 3 and Table 3 for variables description).

Empirical results are shown in Table 5. The proportion of SDG sentences (SDGContent_Of_Total), as well as the proportion of SDG claims (SDG-Claims_of_Total), displays a strong and consistently negative relationship with the SDGOI across all model specifications ($p < 0.01$). This indicates, as expected, that a greater proportion of SDG content within a report is associated with more complete coverage of the declared SDGs. The negative coefficient suggests that integrating SDG content might serve as a mechanism for mitigating omission and selective disclosure, which is consistent with Hypothesis 1.

Conversely, Hypothesis 2 cannot be substantiated, as the share of environmental sentences (EnvContent_Of_Total) and the share of environmental claims (EnvClaims_Of_Total) are not statistically significant.

Hypotheses 3 and 4 cannot be rejected. The proportion of environmental sentences with opportunity-oriented sentiment (Sent_Opp_Of_EnvContent) exhibits a consistently significant positive relationship ($p < 0.1$ or $p < 0.05$) with the dependent variable across all models. We acknowledge that opportunity-oriented language is inherent to financial communication, where emphasizing value creation and transition potential is standard practice. However, within the specific context of sustainability reporting, an excessive reliance on opportunity-laden rhetoric, when not supported by comprehensive disclosure, can serve as a mechanism of impression management. In this context, opportunity-oriented sentiment appears to act as a ‘positivity bias’ that masks information gaps rather than a neutral disclosure strategy. Conversely, the opportunity-oriented sentiment of SDG sentences (SDG_Sent_Opp_Of_SDGContent) demonstrates a consistent and highly significant negative relationship ($p < 0.01$) with the SDGs Omission Index across all models. This implies that an opportunity-oriented tone in SDG-related communication is associated with broader disclosure of the declared SDGs.

Regarding risky sentiment, the coefficients of SDG-related sentences with risky sentiment (SDG_Sent_Risk_Of_SDGContent) as well as those of environmental sentences with risky sentiment (Sent_Risk_Of_EnvContent) are not significant.

Results remain robust with different specifications (see Tables S1, S2, and S3).

Overall, the results point to a nuanced link between disclosure content, tone, and the selective omission of SDG-related

TABLE 5 | Fixed effects panel regression estimates of the impact of environmental and SDG-related content, claims and sentiment on SDGOI.

	(1) SDGOI	(2) SDGOI	(3) SDGOI	(4) SDGOI
EnvContent_of_Total	0.040	—	—	—
EnvClaims_of_Total	(0.119)	0.199	—	—
SDGContent_of_Total	—	(0.375)	-4.672***	—
SDGClaims_of_Total	—	—	(1.008)	-6.250***
	—	—	—	(2.163)
Sent_Opp_of_EnvContent	0.302**	0.282*	0.325**	0.328**
	(0.144)	(0.150)	(0.128)	(0.132)
Sent_Risk_of_EnvContent	-0.176	-0.166	-0.235	-0.191
	(0.311)	(0.311)	(0.295)	(0.242)
SDG_Sent_Opp_of_SDGContent	-0.426***	-0.427***	-0.412***	-0.410***
	(0.066)	(0.066)	(0.065)	(0.035)
SDG_Sent_Risk_of_SDGContent	-0.130	-0.125	-0.145	-0.148
	(0.132)	(0.132)	(0.133)	(0.096)
Active_Bonds	0.001*	0.001**	0.001	0.001
	(0.000)	(0.000)	(0.000)	(0.000)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Sector FE	Yes	Yes	Yes	Yes
Observations	667	667	667	667
R ²	0.287	0.288	0.320	0.300
F statistic	19.079***	19.103***	22.308***	20.233***

Note: Standard errors clustered at firm level in parentheses. Standard errors are heteroskedasticity-robust and clustered at the firm level. Year and sector fixed effects are included in all specifications. Robustness test including firm fixed effects w.r.t. also domicile and type of report provide qualitatively similar results (see Table S1). The linear fixed-effects (OLS) estimator is presented as the baseline specification because it delivers directly interpretable marginal effects and accommodates a rich set of fixed effects. By contrast, the correlated random-effects fractional logit estimator is used as a robustness check, as it ensures correct handling of the fractional nature of the dependent variable while preserving comparability with the linear FE structure. Findings on opportunity-oriented sentiment remain robust when estimating a fractional logit model with correlated random effects (Papke and Wooldridge 2008) (see Table S2).

* $p < 0.1$. ** $p < 0.05$. *** $p < 0.01$.

information in sustainability reports. This motivates a new metric that isolates opportunity-oriented environmental sentiment *outside* SDG-referenced statements and examines whether it covaries with external benchmarks such as ESG controversies and greenwashing accusations. Any positive association should be interpreted as convergent evidence that this tone-based signal is consistent with greenwashing risk, rather than as definitive detection.

5 | A Sentiment-Based Metric

The analysis of SDGOI shows that higher SDG omission tends to co-occur with a more opportunity-oriented tone in environmental statements. This pattern suggests that textual features can be informative about *commitment–disclosure misalignment* that is, situations in which SDGs declared ex ante in green bond documentation are less visibly reflected in subsequent sustainability reporting. Importantly, we interpret this result as associational and consistent with selective disclosure risk, rather than as conclusive evidence of opportunistic intent. Opportunity-oriented language is common in financial communication; it becomes potentially informative in our setting primarily when it coincides with weaker SDG-specific substantiation (higher SDGOI), a

configuration consistent with “narrative substitution” under scrutiny.

5.1 | The Environmental Sentiment Metric

We introduce the Environmental Sentiment Metric (ESM), which is calculated as the standardized count of environment-related sentences exhibiting opportunity-oriented sentiment, following min–max normalization:

$$ESM_{i,t} = \frac{Env_Sent_Opp_{i,t} - \overline{Env_Sent_Opp}_t}{\sigma_{Env_Sent_Opp_t}} \quad (3)$$

where $Env_Sent_Opp_{i,t}$ is the number of environmental labeled sentences with opportunity-oriented sentiment for issuer i in year t after min–max normalization, $\overline{Env_Sent_Opp}_t$ is the mean of the number of environmental labeled sentences with opportunity-oriented sentiment across all issuers in year t , and $\sigma_{Env_Sent_Opp_t}$ is the corresponding standard deviation.

The ESM addresses critical challenges in textual analysis of sustainability disclosures by integrating min–max normalization and annual standardization. Min–max normalization accounts

for stylistic differences (Sebastiani 2002) and verbosity while preserving ordinality, avoiding the instability inherent in ratio-based metrics that conflate sentiment intensity with textual volume (Gentzkow and Shapiro 2010; Loughran and McDonald 2011). Standardization within each year ensures cross-sectional and temporal comparability, anchoring scores to evolving reporting norms (e.g., post-2015 SDGs adoption trends) and isolating firm-level deviations from industry-wide baselines (Hoberg and Phillips 2016). Crucially, the ESM design accounts for the prevalence of transition-related language by standardizing scores annually (see Equation 3). This process centers the metric around the market average, effectively neutralizing the ‘standard’ level of opportunity signaling expected in financial reporting. Consequently, high ESM values reflect a relative excess of promotional rhetoric rather than its mere presence, distinguishing strategic overemphasis from normal market communication. By retaining absolute counts rather than ratios, the metric preserves the magnitude of opportunity-oriented sentences, distinguishing firms with substantively different disclosure scales—critical for greenwashing alerts. This methodology aligns with NLP best practices, mirroring TF-IDF weighting for corpus-adjusted relevance (Ramos 2003) and Z-score standardization for robust feature engineering in document classification (Sebastiani 2002).

5.2 | ESG Controversies Score and Greenwashing Accusations

In the absence of actual greenwashing cases, empirical testing relies on issuers’ ESG Controversy Scores (ESGCs) from LSEG and a self-developed Greenwashing Accusations (GWAs) metric. In line with prior literature, we define a controversy as any negative ESG-related event identified through media monitoring (Delmas and Burbano 2011; Lyon and Maxwell 2011), whereas an accusation refers to explicit claims of misleading or deceptive sustainability communication, whether in legal documents, regulatory actions, NGO reports, or investigative journalism. Furthermore, we also consider ESG Scores and the logarithm of the market capitalization as a proxy for issuer size, both provided by LSEG.

The ESG Score, as defined by LSEG, represents an overall company score based on the self-reported information in the environmental, social, and corporate governance pillars, with values from 0 (poorest ESG performance) to 100 (best grade). The ESGCs score, as defined by LSEG, is a measure of a company’s exposure to environmental, social, and governance controversies and negative events related to sustainability issues which are tracked by global media. In the following analysis we construct a binary variable, still called ESGCs, which is equal to 0 if the firm had no controversies and 1 otherwise.

Figure 1, in the top panel, reports the relationship between ESG Scores and the presence of ESG controversies in the sample period. The boxplots suggest that firms facing controversies tend to exhibit higher ESG Scores. This result is in line with the literature (e.g., Kathan et al. 2025; Kim and Lyon 2015; Marquis et al. 2016) and could be explained by considering that firms with better scores receive more media and stakeholder

attention and, as a consequence, could be more likely subject to disputes.

We also introduce the Greenwashing Accusations (GWAs) variable, defined as a binary indicator that identifies whether a firm was subject to formal legal allegations or public accusations of greenwashing between 2019 and 2023. To account for potential late disclosures in sustainability reports, the coverage of 2023 accusations was extended through March 2024.

Following established literature on corporate environmental misconduct (Delmas and Burbano 2011; Lyon and Montgomery 2015), the variable includes both formal legal actions (e.g., regulatory sanctions or lawsuits) and non-legal but publicly documented allegations (e.g., NGO reports or media investigations), recognizing that reputational damage can arise even without litigation (Seele and Gatti 2017).

We constructed the variable through a two-phase methodology, integrating AI-assisted data collection with manual verification. Initial data gathering employed large language models (i.e., DeepSeek and ChatGPT), which were systematically prompted to identify legal and non-legal allegations against each target company between January 2019 and March 2024. To ensure traceability, each query explicitly requested source citations and verifiable URLs.

As recommended in recent LLM auditing research (Tonmoy et al. 2024), all AI-collected outputs were subject to strict manual source verification to avoid hallucinations and ensure provenance. Rigorous manual validation was conducted, involving both direct examination of the AI-provided sources and supplementary searches using identical parameters on Google.

This dual approach, leveraging computational efficiency while maintaining human oversight, was implemented to mitigate potential hallucinations or omissions inherent in automated retrieval systems, thereby enhancing the dataset’s reliability.

While this approach provides a useful measure of greenwashing risk, it may underreport accusations against private firms (Kim and Lyon 2015). Nevertheless, the variable aligns with reputational risk theory (Walker and Wan 2012), capturing both regulatory and societal dimensions of greenwashing, and is particularly relevant for green bond issuers given the region’s stringent disclosure standards (EU Taxonomy Regulation, 2020).

The boxplots in the bottom panel of Figure 1 highlight a positive relationship also between ESG Scores and GWAs. Both ESGCs and GWAs, as a consequence, appear to be coherent indicators of reputational risk exposure, suggesting that they capture overlapping dimensions of ESG-related scrutiny, albeit with differing scope and intensity. GWA concentration among high-ESG firms mirrors the broader controversy trend, reinforcing the idea that superior ESG performance attracts heightened attention by stakeholders that may uncover both generalized controversies and targeted allegations. Both ESGCs and GWAs serve as indicators of potential issues related to greenwashing among green bond issuers. ESGCs typically cover a broad range of incidents and apply to a larger set of companies, with 208 reported cases.

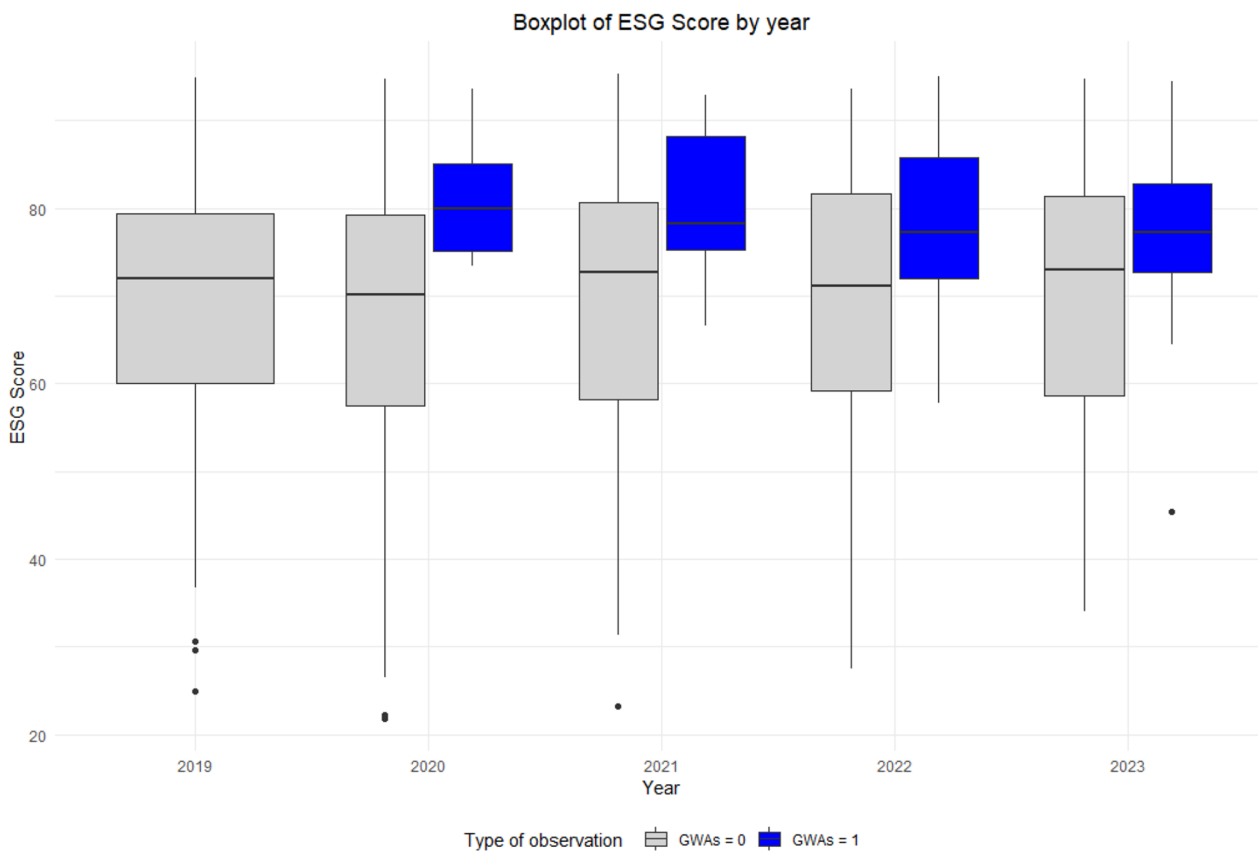



FIGURE 1 | Boxplots of ESG scores. Top panel: issuers with ESGCs = 0 (gray) and ESGCs = 1 (red). Bottom panel: issuers with GWAs = 0 (gray) and GWAs = 1 (blue).

TABLE 6 | Univariate analysis results comparing firms with ESGCs equal to 1 or 0.

Variable	Mean	Mean	Mean	<i>t</i> test	Wilcoxon
	(ESGCs = 0)	(ESGCs = 1)	Diff.		
ESM _{<i>t</i>}	−0.13	0.26	0.39	***	***
ESM _{<i>t</i>+1}	−0.13	0.32	0.45	***	***
ESG score _{<i>t</i>}	65.88	76.90	11.02	***	***
ESG score _{<i>t</i>+1}	67.28	77.82	10.54	***	***
ΔESG score	2.07	0.77	−1.30	**	**
Size _{<i>t</i>}	23.00	24.10	1.10	***	***

Note: Differences in means and medians for the ESM and ESG Scores are evaluated using *t* tests and Wilcoxon rank-sum tests.

p* < 0.05. *p* < 0.01.

In contrast, GWAs (coded as 1) are more narrowly defined and pertain to a smaller subset of issuers (50 in total), often highlighting specific cases in which firms are suspected of misleadingly portraying their sustainability efforts. As expected, there is strong overlap between the two, with 34 out of the 50 GWAs also captured within ESGCs.

Furthermore, if we consider the distribution of firm size by ESGCs and GWAs, we can notice a positive dependence between those variables. These findings are consistent with prior literature which suggests that larger firms have broader operations and higher ESG values, and that greater visibility and disclosure by large firms increases the probability of detection and reporting of controversies (e.g., Gregory 2024; Kathan et al. 2025; Lins et al. 2017).

We acknowledge that ESG controversies and greenwashing accusations are noisy and potentially visibility-biased benchmarks: ESG controversies may capture events unrelated to disclosure, and accusations tend to be rare and concentrated among large, high-profile firms. For this reason, we do not interpret the observed associations between the ESM and these benchmarks as confirmatory evidence of greenwashing. Rather, we treat them as convergent (concurrent) validity checks, i.e., supportive evidence that the proposed metrics co-vary with external signals plausibly related to greenwashing risk.

5.3 | Empirical Analysis

In the following, we conduct both univariate and multivariate analyses to empirically test the ESM.

5.3.1 | Univariate Analysis

We begin with a univariate analysis to compare group means between firms with ESGCs or GWAs equal to 1 and those equal to 0, in order to test whether the two groups differ significantly across key variables. To assess statistical significance, we applied two tests: an independent samples *t* test, which evaluates whether the means of each variable differ significantly between groups, and a Wilcoxon rank-sum test, which assesses differences in medians.

Table 6 presents group differences between observations with and without ESG controversies (ESGCs equal to 1 vs. ESGCs equal

to 0) at time *t*, regarding the ESM, ESG Scores, the difference in ESG scores from time *t* + 1 to *t* (i.e., ΔESG Score), and Size at time *t*.

The analysis reveals statistically significant differences between firms with and without ESGCs. Results related to the ESM show that being a firm with controversies tends to be associated with a higher frequency of opportunity-oriented environmental sentences in sustainability reports.

As expected from inspecting also Figure 1, firms under dispute exhibit significantly higher ESG Scores both in *t* and *t* + 1.

Several strands of academic research have sought to explain the counterintuitive finding that firms with higher ESG scores may also be more frequently associated with ESG-related controversies (e.g., Kathan et al. 2025; Kim and Lyon 2015). First, firms with elevated ESG ratings are often large, publicly traded entities that actively promote their sustainability credentials. This heightened visibility could make them more susceptible to scrutiny from stakeholders, media outlets, and regulatory bodies. Consequently, any deviation from stated ESG commitments is more likely to be detected and publicized, resulting in a greater incidence of reported controversies. Furthermore, firms with robust ESG programs often engage in multifaceted initiatives across diverse geographies and sectors. The complexity inherent in managing such programs increases the likelihood of operational lapses or ethical breaches.

Lastly, it is important to underline the temporal granularity of ESG scores provided by LSEG. ESG performance is typically reported on an annual basis, reflecting year-end values across a range of environmental, social, and governance indicators. This annual frequency implies that ESG scores may not fully capture intra-year developments, including short-term controversies or corrective actions taken by firms in response to stakeholder pressure. The lagged nature of ESG data can contribute to a disconnect between a firm's current ESG performance and its most recently published ESG score. For instance, a firm may receive a high ESG rating based on the previous year's disclosures, while simultaneously facing reputational challenges or controversies that emerge in the current reporting period. This temporal misalignment may partially explain the observed positive correlation between ESG scores and ESG controversies in empirical studies: ESG scores provided by LSEG may incorporate information regarding controversies over time. Indeed, the improvement in

TABLE 7 | Univariate analysis results comparing firms with GWAs equal to 1 or 0.

Variable	Mean (GWAs = 0)	Mean (GWAs = 1)	Mean diff.	<i>t</i> test	Wilcoxon
ESM _{<i>t</i>}	-0.04	0.42	0.46	**	***
ESM _{<i>t+1</i>}	0.00	0.45	0.45	*	***
ESG score _{<i>t</i>}	68.80	78.29	9.49	***	***
ESG score _{<i>t+1</i>}	70.41	80.60	10.19	***	***
ΔESG score	1.69	0.37	-1.32	—	—
Size _{<i>t</i>}	23.30	24.20	0.90	***	***

Note: Differences in means and medians for the ESM and ESG scores are evaluated using *t* tests and Wilcoxon rank-sum tests.

p* < 0.1. *p* < 0.05. ****p* < 0.01.

ESG Score (ΔESG Score) is less pronounced for controversial firms (0.77) compared to non-controversial ones (2.07). While ESG Scores tend to improve over time, those of controversial firms improve less, suggesting that the score incorporates information regarding disputes only with some delay.

The same group comparison analysis was replicated using GWAs as the classification variable, yielding qualitatively consistent results, which are shown in Table 7. The direction and statistical significance of key findings, including higher opportunity-oriented sentiment, larger size, elevated ESG scores, and attenuated post-event ESG improvement, closely mirrored those observed for ESGCs. This coherence across the two related measures strengthens the inference that both variables capture underlying firm-level traits associated with reputational risk and strategic disclosure behavior. Such consistency supports the use of either variable, ESGC or GWA, or both complementarily in analyses of ESG-related disclosure strategies.

From an economic perspective, these univariate patterns are consistent with a disclosure-and-scrutiny channel. Larger and more visible issuers have stronger incentives to emphasize transition opportunities to attract investor demand and sustain favorable ESG perceptions, especially in markets where sustainability labels may lower financing costs and broaden the investor base. At the same time, visibility increases monitoring intensity by media, NGOs, investors, and regulators, raising the likelihood that contentious events are detected and recorded as ESG controversies or accusations. In this environment, opportunity-oriented disclosure can function as a reputational instrument: it is not inherently misleading, but when it is disproportionately prominent, it may accompany higher exposure to reputational risk and closer scrutiny. This interpretation supports the role of the ESM as greenwashing alert indicators that help prioritize issuers for further review, rather than as tools that identify or detect greenwashing.

5.3.2 | Multivariate Analysis

Building on the univariate analysis, we estimate two separate logistic regression models to examine how key explanatory variables affect the likelihood of ESGCs and GWAs. Each model includes fixed effects for year and sector to control for unobserved heterogeneity and to isolate the effect of the independent variables on the outcomes of interest. In particular, we estimate

$$\text{logit}(Pr[Y_{i,t}]) = \beta_0 + \beta_1 \text{ESM}_{i,t} + \beta_2 \text{ESG score}_{i,t} + \beta_3 \text{Size}_{i,t} + FE_t, \quad (4)$$

$$\text{logit}(Pr[Y_{i,t+1}]) = \beta_0 + \beta_1 \text{ESM}_{i,t} + \beta_2 \text{ESG score}_{i,t} + \beta_3 \text{Size}_{i,t} + FE_t, \quad (5)$$

where *Y* is either ESGCs or GWAs, and *FE_t* refers to year and sector fixed effects.

Estimates with ESGCs and GWAs at time *t* are reported in Table 8, while Table 9 shows the results when considering ESGCs and GWAs at time *t* + 1. Counterintuitively, higher ESG Scores show a positive association both with ESGCs at time *t* and *t* + 1, and with GWAs at time *t* and *t* + 1, potentially reflecting “ESG overclaiming” (Lyon and Montgomery 2015) or measurement limitations in capturing genuine sustainability efforts, as noted in critiques of ESG score divergence (Berg et al. 2022).

The analysis identifies that ESM is persistently linked with ESGCs both contemporaneously and with a one-period lag. The ESM is also positively linked with GWAs within the same year when the ESG Score is not explicitly accounted for. The positive association aligns with theoretical frameworks on opportunity-oriented environmental disclosure, where firms engage in symbolic sustainability practices to manipulate stakeholder perceptions while avoiding substantive environmental commitments (Lyon and Montgomery 2015; Testa et al. 2018).

It is important to clarify the nature of the relationship between opportunity-oriented sentiment and controversies. We caution against a strictly causal interpretation where sentiment ‘triggers’ controversies. It is plausible that disclosure tone and controversy exposure are jointly determined by firm strategy and scrutiny levels: firms facing higher pressure may strategically adopt a more positive, opportunity-driven tone to manage legitimacy. Consequently, our analysis is intended to be associational. Regardless of the direction of causality, the key finding remains: an excess of opportunity-oriented language, particularly consistently co-occurs with higher greenwashing risk profiles.

When ESM and ESG Score are included jointly in the full models, the results confirm that both predictors maintain their statistical significance in explaining ESGCs. This suggests that ESM provides information not fully subsumed by ESG scores, consistent with prior research showing that sentiment-based indicators capture dimensions of market perception distinct from rating-based measures (Loughran and McDonald 2016). The

TABLE 8 | Logistic regression assessing the effect of ESM, ESG Scores and Size on the probability of ESG controversies (ESGCs) and greenwashing accusations (GWAs) at time t .

Dependent variable	ESGCs _{<i>t</i>}			GWAs _{<i>t</i>}		
Variables						
ESM _{<i>t</i>}	0.600*** (0.202)	—	0.347** (0.165)	0.327** (0.129)	—	0.179 (0.134)
ESG score _{<i>t</i>}		0.065*** (0.012)	0.059*** (0.012)		0.057*** (0.014)	0.053*** (0.014)
Size _{<i>t</i>}	0.463*** (0.117)	0.368*** (0.112)	0.367*** (0.112)	0.309*** (0.113)	0.271** (0.116)	0.267** (0.118)
Fixed-effects						
Year	Yes	Yes	Yes	Yes	Yes	Yes
Sector	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	623	623	623	502	502	502
Squared correlation	0.228	0.277	0.289	0.110	0.127	0.129
Pseudo R^2	0.156	0.203	0.211	0.069	0.103	0.100

Note: The models control for year and sector fixed effects and account for clustering at the issuer level. Robustness checks confirm that results are qualitatively similar after including total sentence counts or its logarithmic transformation as control variable. Results are available upon request. Clustered (IssuerID) standard errors in parentheses signif. Codes: ***, 0.01, **, 0.05, *, 0.1.

TABLE 9 | Logistic regression assessing the effect of ESM, ESG Scores and Size on the probability of ESG controversies (ESGCs) and greenwashing accusations (GWAs) at time $t + 1$.

Dependent variable	ESGCs _{<i>t+1</i>}			GWAs _{<i>t+1</i>}		
Variables						
ESM _{<i>t</i>}	0.604*** (0.192)	—	0.362** (0.154)	0.086 (0.201)	—	-0.135 (0.205)
ESG Score _{<i>t</i>}		0.059*** (0.012)	0.053*** (0.012)		0.060*** (0.014)	0.063*** (0.015)
Size _{<i>t</i>}	0.396*** (0.123)	0.300*** (0.114)	0.301*** (0.114)	0.262** (0.107)	0.202* (0.115)	0.203* (0.115)
Fixed-effects						
Year	Yes	Yes	Yes	Yes	Yes	Yes
Sector	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	482	482	482	464	464	464
Squared correlation	0.205	0.258	0.270	0.092	0.128	0.130
Pseudo R^2	0.126	0.167	0.176	0.037	0.093	0.088

Note: The models control for year and sector fixed effects and account for clustering at the issuer level. Robustness checks confirm that these associations are not driven by document length; coefficients remain stable after including total sentence counts and their logarithmic transformation as controls. Clustered (IssuerID) standard errors in parentheses Signif. codes: ***, 0.01, **, 0.05, *, 0.1.

attenuation of ESM's significance in some full specifications is not related to multicollinearity issues (see Tables 10 and 11), but rather is consistent with the partial conceptual overlap between sentiment-based and score-based ESG proxies. This dynamic has been documented in related work showing that alternative sustainability metrics often share explanatory content but contribute complementary information depending on the outcome of interest (Chatterji et al. 2016; Dorfleitner et al. 2015).

Firm size emerges as a positive and highly significant variable across nearly all specifications. Larger firms have a higher probability of facing both ESG controversies and greenwashing accusations, both in the current and subsequent year. This finding provides robust support for the "visibility hypothesis," which posits that larger firms attract greater scrutiny from media, activists, and regulators, making them more susceptible to public controversies (Brammer et al. 2006; Clarkson et al. 2008).

TABLE 10 | Pearson correlation matrix.

Variable	Size	ESG score	ESM
Size	1.00	—	—
ESG score	0.34***	1.00	—
ESM	0.10**	0.31***	1.00

Note: Significance levels.

** $p < 0.01$. *** $p < 0.001$.

TABLE 11 | Multicollinearity diagnostics.

Variable	VIF	Tolerance
Size	1.13	0.88
ESG score	1.24	0.81
ESM	1.11	0.90

Note: Variance inflation factors computed from full logistic models (Tables 4 and 10). All values remain substantially below the conservative threshold of 2.5 (Allison 2012).

Results remain robust with different specifications (see Tables S4 and S5).

Overall, the fact that ESM significantly correlates with external 'red flags' (i.e., ESG Controversies and greenwashing accusations) underscores the importance of distinguishing between authentic environmental stewardship and strategic posturing.

6 | Discussion and Conclusion

Green bonds have emerged as a cornerstone of sustainable finance, enabling companies and governments to raise capital for environmentally beneficial projects. However, the rapid growth of this market has raised concerns about the credibility and verifiability of environmental claims, making it essential to develop metrics for issuing potential alerts on greenwashing. This paper proposes two disclosure-based alert metrics for greenwashing in the European corporate green bond market. The SDGs Omission Index (SDGOI) captures the gap between SDG commitments declared at issuance and their subsequent coverage in sustainability reports; the Environmental Sentiment Metric (ESM) quantifies the relative excess of opportunity-oriented language in environmental disclosures. Together, they provide a scalable, interpretable screening infrastructure grounded in firms' own textual output.

By introducing the SDGOI and linking it to text-based variables, we find that when the SDGs announced at issuance of green bonds are not properly reflected in subsequent sustainability-disclosures, companies tend to adopt and compensate with a more opportunity-driven environmental narrative. In this sense, the SDGOI provides a scalable screening metric that can help identify cases where the alignment between ex ante commitments and ex post disclosure deserves closer attention.

Since the SDGOI rely primarily on issuers' initial SDG declaration, we develop a second, more broadly applicable metric based on environmental sentiment: the ESM. Its strength lies in its wide

applicability and its value as an independent screening tool for supervisory authorities. In our empirical analysis, higher values of ESM are associated with ESG controversies and greenwashing accusations, suggesting that it captures dimensions of disclosure risk that are relevant for ex post monitoring. When resources are scarce and competent authorities must target their efforts efficiently, these metrics offer a practical way to identify issuers whose disclosures merit closer inspection.

The above indicators should not be read as a substitute for legal assessment or evidentiary review, but rather as an alert screening tool capable of helping supervisors prioritize files, thematic reviews, or request clarification. In environments characterized by high volumes of sustainability disclosures and finite enforcement capacity, a tool of this kind may support a more proportionate and data-informed allocation of supervisory resources.

Furthermore, for market participants, the metrics can complement external reviews and second-party opinions by highlighting potential commitment–disclosure misalignment (SDGOI) and unusually opportunity-heavy environmental framing (ESM) that may warrant deeper due diligence. More broadly, the approach is consistent with the direction of European sustainable-finance regulation toward stronger transparency and comparability: rather than replacing qualitative assessment, these indicators provide a structured way to triage cases, improve monitoring efficiency, and enhance accountability around stated use of proceeds and impact narratives.

Some limitations should be acknowledged. Our study focuses exclusively on the European green bond market, although the methodology itself is highly adaptable and could be extended to other contexts by, for example, modifying the SDG dictionary to reflect different thematic priorities. Additional design variations could also be explored, such as weighting SDGs according to sector-level materiality maps or the prominence they receive within an issuer's framework. Another important avenue for future development involves extending the analysis to documents written in languages other than English, which would likely require language-specific models and careful treatment of the additional noise typically introduced by multilingual datasets.

At the same time, these limitations open several promising directions for future research. Refining SDGOI and ESM through AI and machine-learning techniques could improve their contextual sensitivity and strengthen the detection of omission patterns and sentiment-driven misrepresentation. Applying these metrics to a broader and more diverse pool of issuers would further enhance their generalizability and robustness, while a wider availability of documented greenwashing cases would be especially valuable for testing their predictive power. Future work should also examine the adaptability of the proposed model to evolving regulatory frameworks, assessing whether it can remain effective as sustainability disclosure requirements are revised and expanded.

More broadly, the flexibility of the methodology suggests that it could be calibrated to different user needs, supervisory settings, disclosure regimes, and financial products. Over time, embedding SDGOI and ESM within the evolving infrastructure of sustainable finance, whether in disclosure standards, independent assurance practices, or ESG data platforms, could significantly

enhance their practical utility. Such integration would not only provide investors and regulators with more reliable tools for assessing the credibility of sustainability disclosures but also help foster more harmonized monitoring practices and greater convergence in sustainability-related communications.

Endnotes

¹See: <https://eur-lex.europa.eu/legal-content/en/txt/pdf/?uri=celex:52018dc0097>.

²See: <https://www.icmagroup.org/assets/documents/regulatory/green-bonds/june-2020/mapping-sdgs-to-green-social-and-sustainability-bonds-2020-june-2020-090620.pdf>.

Acknowledgments

Open access publishing facilitated by Universita degli Studi di Trento, as part of the Wiley - CRUI-CARE agreement.

References

- Allison, P. D. 2012. *Logistic Regression Using SAS: Theory and Application*. SAS Institute.
- Arena, C., S. Bozzolan, and G. Michelon. 2015. "Environmental Reporting: Transparency to Stakeholders or Stakeholder Manipulation? An Analysis of Disclosure Tone and the Role of the Board of Directors." *Corporate Social Responsibility and Environmental Management* 22, no. 6: 346–361.
- Berg, F., J. F. Kölbel, and R. Rigobon. 2022. "Aggregate Confusion: The Divergence of Esg Ratings." *Review of Finance* 26, no. 6: 1315–1344.
- Bernini, F., M. Giuliani, and F. La Rosa. 2024. "Measuring Greenwashing: A Systematic Methodological Literature Review." *Business Ethics, the Environment & Responsibility* 33, no. 4: 649–667.
- Bingler, J. A., M. Kraus, M. Leippold, and N. Webersinke. 2022. "Cheap Talk and Cherry-Picking: What Climatebert Has to Say on Corporate Climate Risk Disclosures." *Finance Research Letters* 47: 102776.
- Bingler, J. A., M. Kraus, M. Leippold, and N. Webersinke. 2024. "How Cheap Talk in Climate Disclosures Relates to Climate Initiatives, Corporate Emissions, and Reputation Risk." *Journal of Banking & Finance* 164: 107191.
- Brammer, S., C. Brooks, and S. Pavelin. 2006. "Corporate Social Performance and Stock Returns: UK Evidence From Disaggregate Measures." *Financial Management* 35, no. 3: 97–116.
- Chatterji, A. K., R. Durand, D. I. Levine, and S. Touboul. 2016. "Do Ratings of Firms Converge? Implications for Managers, Investors and Strategy Researchers." *Strategic Management Journal* 37, no. 8: 1597–1614.
- Cho, C. H., R. P. Guidry, A. M. Hageman, and D. M. Patten. 2012. "Do Actions Speak Louder Than Words? An Empirical Investigation of Corporate Environmental Reputation." *Accounting, Organizations and Society* 37, no. 1: 14–25.
- Cho, C. H., M. Laine, R. W. Roberts, and M. Rodrigue. 2015. "Organized Hypocrisy, Organizational Façades, and Sustainability Reporting." *Accounting, Organizations and Society* 40: 78–94.
- Cho, C. H., G. Michelon, and D. M. Patten. 2012. "Impression Management in Sustainability Reports: An Empirical Investigation of the Use of Graphs." *Accounting and the Public Interest* 12, no. 1: 16–37. <https://doi.org/10.2308/apin-10249>.
- Christensen, H. B., L. Hail, and C. Leuz. 2021. "Mandatory Csr and Sustainability Reporting: Economic Analysis and Literature Review." *Review of Accounting Studies* 26, no. 3: 1176–1248.
- Clarkson, P. M., Y. Li, G. D. Richardson, and F. P. Vasvari. 2008. "Revisiting the Relation Between Environmental Performance and Environmental Disclosure: An Empirical Analysis." *Accounting, Organizations and Society* 33, no. 4–5: 303–327.
- Delmas, M. A., and V. C. Burbano. 2011. "The Drivers of Greenwashing." *California Management Review* 54, no. 1: 64–87.
- Dorflleitner, G., G. Halbritter, and M. Nguyen. 2015. "Measuring the Level and Risk of Corporate Responsibility—An Empirical Comparison of Different Esg Rating Approaches." *Journal of Asset Management* 16, no. 7: 450–466.
- Dorflleitner, G., and S. Utz. 2024. "Green, Green, It's Green They Say: A Conceptual Framework for Measuring Greenwashing on Firm Level." *Review of Managerial Science* 18: 3463–3486.
- Drempetic, S., C. Klein, and B. Zwergel. 2020. "The Influence of Firm Size on the ESG Score: Corporate Sustainability Ratings Under Review." *Journal of Business Ethics* 167: 333–360.
- Dyck, A., K. V. Lins, L. Roth, and H. F. Wagner. 2019. "Do Institutional Investors Drive Corporate Social Responsibility? International Evidence." *Journal of Financial Economics* 131, no. 3: 693–714.
- European Banking Authority. 2024. *Greenwashing Monitoring and Supervision: Final Report (EBA/REP/2024/09)*.
- Feghali, K., R. Najem, and B. D. Metcalfe. 2025. "Greenwashing in the Era of Sustainability: A Systematic Literature Review." *Corporate Governance and Sustainability Review* 9, no. 1: 18–31.
- Flammer, C. 2021. "Corporate Green Bonds." *Journal of Financial Economics* 142, no. 2: 499–516.
- Ge, P., Y. Liu, C. Tang, and R. Zhu. 2025. "Green Bonds and Corporate Environmental Social and Governance Performance: Innovative Approaches to Identifying Greenwashing in Green Bond Markets." *Corporate Social Responsibility and Environmental Management* 32, no. 1: 1060–1078.
- Gentzkow, M., and J. M. Shapiro. 2010. "What Drives Media Slant? Evidence From Us Daily Newspapers." *Econometrica* 78, no. 1: 35–71.
- Gorovaia, N., and M. Makrominas. 2024. "Identifying Greenwashing in Corporate-Social Responsibility Reports Using Natural-Language Processing." *European Financial Management* 31, no. 1: 427–462.
- Gregory, R. P. 2024. "How Greenwashing Affects Firm Risk: An International Perspective." *Journal of Risk and Financial Management* 17, no. 11: 526.
- Hoberg, G., and G. Phillips. 2016. "Text-Based Network Industries and Endogenous Product Differentiation." *Journal of Political Economy* 124, no. 5: 1423–1465.
- International Capital Market Association. 2022. "Green Bond Principles: Voluntary Process Guidelines for Issuing Green Bonds." <https://www.icmagroup.org/sustainable-finance/the-principles-guidelines-and-handbooks/green-bond-principles-gbp/>.
- Kathan, M. C., S. Utz, G. Dorflleitner, J. Eckberg, and L. Chmel. 2025. "What You See Is Not What You Get: ESG Scores and Greenwashing Risk." *Finance Research Letters* 74: 106710.
- Kim, E.-H., and T. P. Lyon. 2015. "Greenwash vs. Brownwash: Exaggeration and Undue Modesty in Corporate Sustainability Disclosure." *Organization Science* 26, no. 3: 705–723.
- Koelbl, M. 2019. *MD&A Disclosure and Performance of US REITs: The Information Content of Textual Tone (Tech. Rep.)*. European Real Estate Society (ERES).
- Lagasio, V. 2024. "ESG-Washing Detection in Corporate Sustainability Reports." *International Review of Financial Analysis* 96: 103742.
- Lins, K. V., H. Servaes, and A. Tamayo. 2017. "Social Capital, Trust, and Firm Performance: The Value of Corporate Social Responsibility During the Financial Crisis." *Journal of Finance* 72, no. 4: 1785–1824.

- Liu, C., J. Wang, Q. Ji, and D. Zhang. 2024. "To be Green or Not to be: How Governmental Regulation Shapes Financial Institutions' Greenwashing Behaviors in Green Finance." *International Review of Financial Analysis* 93: 103225.
- Liu, Y., S. Fu, G. Liu, and X. Huang. 2025. "Green Bond Issuance and Corporate Environmental Disclosure: Signaling or Greenwashing?" *Journal of Business Ethics* 204: 825–854.
- Loughran, T., and B. McDonald. 2011. "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks." *Journal of Finance* 66, no. 1: 35–65.
- Loughran, T., and B. McDonald. 2016. "Textual Analysis in Accounting and Finance: A Survey." *Journal of Accounting Research* 54, no. 4: 1187–1230.
- Lublov, Á., J. L. Keresztúri, and E. Berlinger. 2025. "Quantifying Firm-Level Greenwashing: A Systematic Literature Review." *Journal of Environmental Management* 373: 123399.
- Luccioni, A., E. Baylor, and N. Duchene. 2020. Analyzing Sustainability Reports Using Natural Language Processing arXiv Preprint arXiv:2011.08073.
- Lyon, T. P., and J. W. Maxwell. 2011. "Greenwash: Corporate Environmental Disclosure Under Threat of Audit." *Journal of Economics and Management Strategy* 20, no. 1: 3–41.
- Lyon, T. P., and A. W. Montgomery. 2015. "The Means and End of Greenwash." *Organization & Environment* 28, no. 2: 223–249.
- Marquis, C., M. W. Toffel, and Y. Zhou. 2016. "Scrutiny, Norms, and Selective Disclosure: A Global Study of Greenwashing." *Organization Science* 27, no. 2: 483–504.
- Moodaley, W., and A. Telukdarie. 2023. "Greenwashing, Sustainability Reporting and Artificial Intelligence: A Systematic Literature Review." *Sustainability* 15, no. 2: 1481.
- OECD. 2025. *Global Debt Report 2025: Financing Growth in a Challenging Debt Market Environment*. OECD Publishing. <https://doi.org/10.1787/8ee42b13-en>.
- Papke, L. E., and J. M. Wooldridge. 2008. "Panel Data Methods for Fractional Response Variables With an Application to Test Pass Rates." *Journal of Econometrics* 145, no. 1–2: 121–133.
- Paterlini, S., A. Nicolodi, M. Gentile, V. Manzillo, M. Sancilio, and P. Deriu. 2025. "Greenwashing Alert System for EU Green Bonds. The CONSOB-University of Trento Prototype (CONSOB Fintech Series No. 14)." CONSOB. <https://doi.org/10.2139/ssrn.5379964>.
- Peng, Q., and Y. Xie. 2024. "Esg Greenwashing and Corporate Debt Financing Costs." *Finance Research Letters* 69: 106012.
- Persakis, A., T. Nikolopoulos, I. C. Negkakis, and A. Pavlopoulos. 2025. "Greenwashing in Marketing: A Systematic Literature Review and Bibliometric Analysis." *International Review on Public and Nonprofit Marketing* 22: 957–992.
- Poiriazzi, E., G. Zournatzidou, G. Konteos, and N. Sariannidis. 2025. "Analyzing the Interconnection Between Environmental, Social, and Governance (ESG) Criteria and Corporate Corruption: Revealing the Significant Impact of Greenwashing." *Administrative Sciences* 15, no. 3: 100.
- Ramos, J. 2003. "Using TF-IDF to Determine Word Relevance in Document Queries." *Proceedings of the First Instructional Conference on Machine Learning* 242, no. 1: 29–48.
- Roggi, O., L. Bellardini, and S. Conticelli. 2024. "Effects of Esg Performance and Sustainability Disclosure on Gss Bonds' Yields and Spreads: A Global Analysis." *Finance Research Letters* 68: 105988.
- Schimanski, T., A. Reding, N. Reding, J. Bingler, M. Kraus, and M. Leippold. 2024. "Bridging the Gap in Esg Measurement: Using NLP to Quantify Environmental, Social, and Governance Communication." *Finance Research Letters* 61: 104979.
- Sebastiani, F. 2002. "Machine Learning in Automated Text Categorization." *ACM Computing Surveys* 34, no. 1: 1–47.
- Seele, P., and L. Gatti. 2017. "Greenwashing Revisited: In Search of a Typology and Accusation-Based Definition Incorporating Legitimacy Strategies." *Business Strategy and the Environment* 26, no. 2: 239–252.
- Sneideriene, A., and R. Legenzova. 2025. "Greenwashing Prevention in Environmental, Social, and Governance (ESG) Disclosures: A Bibliometric Analysis." *Research in International Business and Finance* 74: 102720.
- Stammach, D., N. Webersinke, J. Bingler, M. Kraus, and M. Leippold. 2023. "Environmental Claim Detection." In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, vol. 2, 1051–1066. Association for Computational Linguistics.
- Tang, D. Y., and Y. Zhang. 2020. "Do Shareholders Benefit From Green Bonds?" *Journal of Corporate Finance* 61: 101427.
- Testa, F., O. Boiral, and F. Iraldo. 2018. "Internalization of Environmental Practices and Institutional Complexity: Can Stakeholders Pressures Encourage Greenwashing?" *Journal of Business Ethics* 147: 287–307.
- Tillé, Y., and M. Wilhelm. 2017. "Probability Sampling Designs: Principles for Choice of Design and Balancing." *Statistical Science* 32: 176–189.
- Todaro, D. L., and R. Torelli. 2024. "From Greenwashing to ESG-Washing: A Focus on the Circular Economy Field." *Corporate Social Responsibility and Environmental Management* 31, no. 5: 4034–4046.
- Tonmoy, S., S. Zaman, V. Jain, et al. 2024. A Comprehensive Survey of Hallucination Mitigation Techniques in Large Language Models. arXiv Preprint arXiv:2401.01313.
- Treepongkaruna, S., H. H. Au Yong, S. Thomsen, and K. Kyaw. 2024. "Greenwashing, Carbon Emission, and ESG." *Business Strategy and the Environment* 33, no. 8: 8526–8539.
- Tsang, R. C., A. A. Baldwin, J. F. Hair, E. Affuso, and K. D. Lahtinen. 2023. "The Informativeness of Sentiment Types in Risk Factor Disclosures: Evidence From Firms With Cybersecurity Breaches." *Journal of Information Systems* 37, no. 3: 157–190.
- Walker, K., and F. Wan. 2012. "The Harm of Symbolic Actions and Greenwashing: Corporate Actions and Communications on Environmental Performance and Their Financial Implications." *Journal of Business Ethics* 109: 227–242.
- Webersinke, N., M. Kraus, J. Bingler, and M. Leippold. 2022. Climate-BERT: A Pretrained Language Model for Climate-Related Text. Proceedings of AAAI 2022 Fall Symposium: The Role of AI in Responding to Climate Challenges.
- Yu, E. P.-y., B. Van Luu, and C. H. Chen. 2020. "Greenwashing in Environmental, Social and Governance Disclosures." *Research in International Business and Finance* 52: 101192.
- Zioło, M., I. Bąk, and A. Spoz. 2024. "Literature Review of Greenwashing Research: State of the Art." *Corporate Social Responsibility and Environmental Management* 31, no. 6: 5343–5356.

Supporting Information

Additional supporting information can be found online in the Supporting Information section. **Table S1:** Firm-level fixed effects estimates. Standard errors clustered at firm level. **Table S2:** Fixed effects correlated random effects (CRE) fractional logit estimates. Standard errors clustered at issuer level. **Table S3:** Robustness: Fixed-Effects Estimates with Winsorised Regressors (OLS). Findings remain robust when also including Domicile and Type of Reports as additional control variables. **Table S4:** Placebo test using randomized environmental sentiment. Logistic regressions of ESG controversies (ESGCs) and greenwashing accusations (GWAs) at time t . The opportunity-oriented sentiment component of the ESM (i.e., `Rand_ESM`) is randomized across environmental sentences, and the ESM formula is recomputed accordingly. The models include year and sector fixed effects and standard errors clustered at the issuer

level. **Table S5:** This table reports the results of logistic regressions estimating the probability of ESG controversies (ESGC) and greenwashing accusations (GWA) at time t and $t + 1$. The main independent variable is the Environmental Sentiment Metric (ESM). The model includes an interaction term with Post2022, a dummy variable equal to 1 for years ≥ 2022 (post-SFDR/Taxonomy implementation) and 0 otherwise. Control variables include firm size, ESG score, and total word count. All models include Year and Sector fixed effects. Standard errors clustered at the issuer level are re-reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix

Acronym	Description
AI	Artificial Intelligence
BERT	Bidirectional Encoder Representations from Transformers
CBI	Climate Bond Initiative
CONSOB	Commissione Nazionale per le Società e la Borsa (CONSOB) is the public authority responsible for regulating the Italian financial markets.
EBA	European Banking Authority
ESG	Environmental, Social & Governance
ESGCs	ESG Controversies (see in Section 5.1)
ESM	Environmental Sentiment Metric (see Section 5)
ESMA	European Securities and Markets Authority
GBPs	Green Bond Principles
GWAs	Greenwashing Accusation (see Section 5.1)
ICMA	International Capital Market Association
LLM	Large Language Model
NGO	Non-Governmental Organization
NLP	Natural Language Processing
OECD	Organization for Economic Co-operation and Development
SDGOI	Sustainable Development Goals Omission Index (see Section 4.1)
SDGs	Sustainable Development Goals
TF-IDF	Term Frequency–Inverse Document Frequency