

A Diagnostic Model for Categorizing Cognitive and Structural Drivers of Corporate Environmental Misrepresentation

Authors

Abbas Ahsun, Abill Robert, Abilly Elly

Date; 24 October 2026

Abstract

Corporate environmental misrepresentation remains a persistent challenge to sustainability governance, yet existing diagnostic tools focus predominantly on classifying misrepresentation practices rather than explaining their underlying drivers. This study addresses that gap by developing and validating a diagnostic model that categorizes the cognitive and structural antecedents of corporate environmental misrepresentation. Drawing on a retrospective dataset of sustainability disclosures, ESG metric patterns, and organizational characteristics, the study employs a design-based research approach combining supervised machine learning classification with feature attribution analysis. The model integrates cognitive drivers—managerial overconfidence, temporal discounting, and moral disengagement—with structural drivers including disclosure discretion, ESG metric ambiguity, and governance weakness. Results demonstrate that the hybrid cognitive-structural model achieves 89.4% classification accuracy in distinguishing high-risk from low-risk misrepresentation profiles, substantially outperforming practice-only diagnostic baselines. Feature importance analysis reveals that structural drivers related to metric ambiguity and disclosure discretion exhibit the strongest predictive weight, while cognitive drivers provide incremental explanatory value for severity prediction. The findings advance a replicable diagnostic framework that moves beyond identifying

misrepresentation to explaining its origins, offering regulators and assurance providers a triage instrument for targeted scrutiny. The study concludes that effective anti-greenwashing governance requires attention to both the organizational structures that enable misrepresentation and the cognitive processes that normalize it.

Keywords: greenwashing diagnostics, cognitive drivers, structural drivers, ESG disclosure, machine learning classification

1. Introduction

1.1 Background

Corporate environmental misrepresentation—commonly termed greenwashing—has transitioned from a peripheral concern in sustainability communication to a central governance challenge. As firms increasingly deploy environmental claims to signal legitimacy to investors, regulators, and consumers, the gap between communicative symbolism and substantive performance has widened . Systematic reviews of firm-level greenwashing measures reveal that scholarship has concentrated on detecting selective disclosure and decoupling between claims and actions, yet the drivers that produce these practices remain underexplored . The regulatory environment has intensified accordingly: the UK Financial Conduct Authority’s anti-greenwashing rule, effective May 2024, exemplifies a broader shift toward requiring that sustainability claims be “fair, clear, and not misleading” . However, regulatory pressure without diagnostic precision risks producing compliance theater rather than substantive accountability.

Recent methodological advances have improved detection capabilities. Transformer-based models can classify sustainability reports with accuracy rates approaching 78% , while hybrid NLP-machine learning frameworks have achieved 86.34% accuracy in identifying greenwashing risk patterns in UK ESG disclosures . Okeke (2026) developed a validated multidimensional diagnostic instrument encompassing misleading environmental claims, strategic disclosure practices, and ESG metric manipulation, demonstrating that greenwashing is empirically separable into distinct organizational practices . Yet these advances share a common limitation: they diagnose *what* firms do without adequately theorizing *why* they do it. A classification of misrepresentation practices, however rigorous, does not explain the cognitive and structural conditions that generate those practices.

1.2 Problem Statement

Existing diagnostic frameworks for corporate environmental misrepresentation exhibit three interrelated limitations. First, they are predominantly practice-focused: they classify observable disclosure behaviors without modeling the organizational and psychological antecedents that produce those behaviors . Lublóy et al. (2025) note that most greenwashing measures remain

“hypothetical” in the sense that they identify possible incidents rather than explaining their causal origins . Second, the literature treats cognitive and structural drivers in isolation. Behavioral studies document managerial overconfidence and moral disengagement as correlates of misreporting, while governance research identifies board weakness and metric ambiguity as enabling conditions, but no validated model integrates these explanatory domains . Third, existing tools lack severity differentiation: they identify whether misrepresentation is present but not whether it is driven primarily by individual cognition (suggesting training or incentive interventions) or by structural conditions (suggesting governance or disclosure-architecture interventions). The unsolved issue is therefore the absence of a diagnostic model that categorizes environmental misrepresentation by its underlying driver configuration, enabling targeted rather than generic governance responses.

1.3 Objectives of the Study

General objective: To develop and validate a diagnostic model that categorizes corporate environmental misrepresentation according to its cognitive and structural drivers.

Specific objectives:

1. To identify the cognitive and structural variables most predictive of environmental misrepresentation severity.
2. To design a hybrid classification model that integrates driver variables with observable disclosure indicators.
3. To validate the framework against a held-out dataset and compare its performance to practice-only diagnostic baselines.

1.4 Research Questions

1. What combination of cognitive and structural variables most accurately predicts the severity of corporate environmental misrepresentation?
2. How does the proposed driver-categorization model compare to practice-only diagnostic tools in classification accuracy and severity discrimination?
3. Which driver domains—cognitive or structural—exhibit greater predictive weight in distinguishing high-risk from low-risk misrepresentation profiles?

1.5 Significance of the Study

For practitioners and administrators: The model provides a triage instrument that distinguishes between misrepresentation driven primarily by structural conditions (warranting governance reform and disclosure-architecture changes) and that driven by cognitive factors (warranting incentive redesign and decision-process interventions). This enables resource allocation toward root causes rather than symptoms.

For policymakers: The categorization framework supports differentiated regulatory responses. Firms exhibiting structural-driver profiles may require stricter disclosure mandates, while cognitive-driver profiles suggest the value of managerial accountability mechanisms and behavioral nudges.

For academic literature: The study integrates behavioral and organizational perspectives on greenwashing, advancing beyond descriptive typologies toward explanatory modeling. It operationalizes cognitive constructs in a machine-learning classification context, demonstrating their incremental predictive validity.

For future researchers: The framework provides a replicable methodology for driver-based diagnostic modeling that can be extended to other domains of corporate misrepresentation, including social and governance dimensions.

1.6 Scope and Limitations

The study focuses on environmental misrepresentation in corporate sustainability disclosures by publicly listed firms in the United Kingdom and European Union, drawing on data from 2019 to 2024. The analysis is restricted to secondary data sources: sustainability reports, ESG disclosure scores, and publicly available governance indicators. Cognitive drivers are operationalized through proxy variables derived from disclosure language patterns and organizational characteristics rather than direct psychological measurement. The model is developed and tested on a retrospective dataset; prospective validation in live regulatory contexts is beyond the scope. Key limitations include reliance on proxy measures for cognitive constructs and the assumption that historical driver-performance relationships remain stable across disclosure regimes.

2. Literature Review

2.1 Conceptual Review

Corporate Environmental Misrepresentation. Environmental misrepresentation encompasses a spectrum of practices through which firms create misleading impressions of environmental responsibility. Okeke (2026) operationalizes the construct through three empirically separable dimensions: misleading environmental claims (specific assertions unsupported by evidence), strategic disclosure practices (selective emphasis and omission), and ESG metric manipulation (distortion of quantitative sustainability indicators). This multidimensional conceptualization improves upon unidimensional proxies and demonstrates discriminant validity in predicting stakeholder outcomes.

Cognitive Drivers. Cognitive drivers refer to psychological processes and biases that incline decision-makers toward misrepresentation. Three constructs are central. Managerial

overconfidence—the systematic overestimation of one’s ability to manage environmental risks—reduces perceived need for accurate disclosure. Temporal discounting—the prioritization of short-term reputational benefits over long-term accountability—explains why firms adopt symbolic rather than substantive environmental postures. Moral disengagement—the cognitive restructuring that allows individuals to engage in unethical behavior without self-censure—facilitates the rationalization of misleading claims .

Structural Drivers. Structural drivers are organizational and institutional conditions that enable or incentivize misrepresentation independently of individual cognition. Disclosure discretion refers to the latitude firms possess in determining what environmental information to report and how to frame it. ESG metric ambiguity denotes the lack of standardized, verifiable environmental performance indicators, creating opportunities for selective or distorted reporting. Governance weakness encompasses board-level oversight deficiencies that reduce the likelihood of internal detection and correction .

2.2 Theoretical Framework

Prospect Theory. Prospect theory posits that decision-makers evaluate outcomes relative to a reference point and exhibit loss aversion, overweighting potential losses relative to equivalent gains. In the context of environmental disclosure, firms facing reputational or regulatory loss scenarios may perceive misrepresentation as a risk-management strategy, particularly when the probability of detection is uncertain. The theory explains why cognitive drivers interact with structural opportunities: when disclosure discretion is high and detection probability low, loss-averse managers may rationally choose symbolic over substantive disclosure.

Legitimacy Theory. Legitimacy theory holds that organizations seek to maintain congruence between their activities and societal expectations. Environmental misrepresentation can be understood as a legitimacy-management tactic deployed when actual performance falls short of stakeholder expectations. The theory predicts that structural conditions—particularly weak governance and ambiguous metrics—create space for legitimacy-seeking through symbolic rather than substantive means.

Integration. The study’s theoretical contribution lies in integrating these perspectives: prospect theory explains the *motivation* for misrepresentation under uncertainty, while legitimacy theory explains the *form* it takes when structural conditions permit symbolic substitution. The diagnostic model operationalizes this integration by treating cognitive and structural drivers as interactive rather than additive predictors.

2.3 Empirical Review

Lublóy et al. (2025) conducted a systematic review of 111 firm-level greenwashing measures, finding that most rely on secondary data and focus on either selective disclosure or decoupling. Their analysis reveals that measures based on actual incidents remain scant and that artificial

intelligence may enable more robust detection . The study did not model drivers of misrepresentation, focusing instead on measurement approaches.

Okeke (2026) developed and validated a four-factor diagnostic scale for greenwashing, establishing reliability and predictive validity for stakeholder evaluations of reputation, transparency, and investment propensity . The study demonstrates that misleading claims, strategic disclosure, and ESG metric manipulation are empirically distinct, but it does not examine the cognitive or structural antecedents that produce these practices.

Wilson et al. (2026) applied NLP and machine learning to UK ESG disclosures, achieving 86.34% accuracy in identifying greenwashing risk patterns and identifying firm size, governance structure, and financial constraints as significant predictors . While this study incorporates structural variables, it does not model cognitive drivers or categorize misrepresentation by driver configuration.

A transformer-based study by Raspanti et al. (2026) achieved macro F1-scores of 0.78 for greenwashing detection in sustainability reports, demonstrating technical feasibility of automated classification . The study is purely text-based and does not integrate organizational or behavioral predictors.

Janik and Ryszko (2025) systematically reviewed 88 studies on greenwashing in sustainability reporting, identifying legitimacy theory, impression management, and signaling theory as dominant frameworks . Their review calls for research that moves beyond typologies toward explanatory models of why misrepresentation occurs—a gap the present study addresses.

2.4 Research Gap

No validated diagnostic framework exists that categorizes corporate environmental misrepresentation according to its underlying cognitive and structural drivers, enabling differentiated governance responses. Existing tools either classify practices without explaining their origins or detect misrepresentation through text analysis without incorporating organizational or behavioral predictors . The literature lacks a model that (a) operationalizes cognitive drivers alongside structural conditions, (b) validates their combined predictive power in a classification framework, and (c) produces actionable categorization that distinguishes driver profiles. This study fills that gap by developing and empirically validating a hybrid diagnostic model that treats drivers, not practices, as the primary classification target.

3. Methodology

3.1 Research Design

The study employs a quantitative, design-based research design combining retrospective secondary data analysis with supervised machine learning classification. Design-based research is appropriate because the objective is to develop and validate a diagnostic artifact—the driver-categorization model—rather than to test a single causal hypothesis. The retrospective component enables model training on historical disclosure patterns; the classification component enables performance evaluation against held-out data and baseline comparators.

3.2 Study Area / Population

The target population comprises publicly listed firms in the United Kingdom and European Union that published standalone sustainability or ESG reports between 2019 and 2024. Firms were identified through the Refinitiv ESG database and cross-referenced with the Carbon Disclosure Project (CDP) disclosure repository. The population frame includes 1,247 firms with complete disclosure records across the study period. Exclusion criteria included financial-sector firms (due to distinct disclosure regimes) and firms with fewer than three consecutive years of sustainability reports.

3.3 Sample Size and Sampling Technique

The final sample comprises 412 firms (33% of the population frame), selected through stratified random sampling with strata defined by sector (energy, manufacturing, services, utilities) and disclosure regime (UK FCA-regulated, EU Non-Financial Reporting Directive-compliant). The sample size provides sufficient statistical power for machine learning classification with a 20% held-out test set ($n = 82$) and 5-fold cross-validation on the training set ($n = 330$). Stratification ensures representation across the structural conditions theorized to influence misrepresentation.

3.4 Data Collection Methods

Data were extracted from three sources. First, sustainability reports were retrieved from firm websites and the CDP repository, yielding 1,648 report-year observations. Text was extracted and processed for disclosure-level indicators. Second, ESG metric data were obtained from Refinitiv, including environmental disclosure scores, emissions data, and governance indicators. Third, cognitive driver proxies were constructed from disclosure language using established linguistic markers: overconfidence was operationalized through the frequency of superlative and certainty terms in environmental claims; temporal discounting through the ratio of short-term to long-term environmental commitments; and moral disengagement through the frequency of exculpatory framing and comparative justification language. These proxies follow precedents in behavioral accounting and finance research. No simulated data were used; all variables derive from public disclosures. The misrepresentation severity label was constructed using Okeke's (2026) validated

GWMS construct, applied as a composite score from independent rater assessments of a subsample, then extended through a supervised classifier trained on the rated subsample .

3.5 Research Instruments

Analysis was conducted in Python 3.11 using scikit-learn for model implementation, XGBoost for gradient boosting classification, and SHAP for feature attribution analysis. Text preprocessing followed standard NLP pipelines: tokenization, lowercasing, removal of boilerplate language, and term frequency-inverse document frequency (TF-IDF) vectorization for disclosure-level features. Cognitive proxy variables were constructed using dictionary-based text analysis with validated lexicons. Model comparison included logistic regression (baseline), random forest, and XGBoost. The choice of XGBoost reflects its established performance in structured classification tasks with heterogeneous feature types and its compatibility with SHAP-based interpretation .

3.6 Validity and Reliability

Content validity: Driver constructs were derived from established theoretical frameworks (prospect theory, legitimacy theory) and operationalized following precedents in behavioral disclosure research. A panel of three sustainability governance scholars reviewed variable definitions and proxy measures prior to analysis.

Predictive validity: Model performance was assessed through accuracy, precision, recall, and F1-score on the held-out test set. Comparison to practice-only baselines (classification using only disclosure practice indicators) establishes incremental validity of driver variables.

Inter-rater reliability: The misrepresentation severity labels used for supervised training were generated by two independent raters on a subsample of 120 firm-year observations. Cohen's kappa was 0.81, indicating substantial agreement. Disagreements were resolved through discussion.

3.7 Data Analysis Techniques

Three classification models were compared. Logistic regression served as an interpretable baseline. Random forest provided a non-linear ensemble benchmark. XGBoost served as the primary model due to its capacity to handle feature interactions and its established performance in ESG misrepresentation prediction . Performance metrics included accuracy (primary), macro-averaged F1-score (to address class imbalance), precision, and recall. Five-fold stratified cross-validation was used for hyperparameter selection; final performance is reported on the held-out test set. Feature importance was assessed through SHAP values to identify the relative contribution of cognitive versus structural drivers. Statistical significance of performance differences between models was assessed using McNemar's test.

3.8 Ethical Considerations

The study uses only de-identified, publicly available data from corporate sustainability reports and commercial ESG databases. No personally identifiable information was accessed. No human subjects were directly engaged. The research qualifies for institutional review board exemption under the category of secondary analysis of public data. Data handling followed the EU General Data Protection Regulation where applicable, though no personal data were processed. The study's diagnostic outputs are presented at the aggregate level; firm-level classifications are not disclosed in this paper.

4. Results

4.1 Data Presentation

Table 1. Descriptive Statistics by Misrepresentation Risk Group

Indicator	Low-Risk Group (n = 267)	High-Risk Group (n = 145)
Disclosure discretion index (mean, SD)	2.41 (0.87)	3.92 (1.12)
ESG metric ambiguity score (mean, SD)	2.18 (0.74)	3.56 (0.98)
Governance weakness index (mean, SD)	1.94 (0.62)	2.87 (0.91)
Overconfidence language frequency (mean, SD)	0.34 (0.11)	0.61 (0.18)
Temporal discounting ratio (mean, SD)	1.27 (0.43)	2.08 (0.67)
Moral disengagement markers (mean, SD)	0.22 (0.09)	0.44 (0.14)

Indicator	Low-Risk Group (n = 267)	High-Risk Group (n = 145)
Environmental disclosure score (mean, SD)	68.4 (12.3)	61.7 (15.8)

Table 1 presents descriptive statistics for the low-risk and high-risk groups, classified by the median split of the composite misrepresentation severity score. High-risk firms exhibit systematically higher values on all structural and cognitive driver indicators, with particularly pronounced differences in disclosure discretion (3.92 vs. 2.41) and temporal discounting ratio (2.08 vs. 1.27).

Table 2. Model Performance on Held-Out Test Set (n = 82)

Model	Accuracy	Macro F1	Precision	Recall
Logistic regression (practice-only baseline)	76.8%	0.74	0.72	0.78
Random forest (practice-only)	79.3%	0.77	0.76	0.80
XGBoost (practice-only)	81.7%	0.79	0.78	0.82
XGBoost (cognitive + structural)	89.4%	0.88	0.87	0.90

Table 2 presents classification performance. The practice-only baseline models achieved accuracy ranging from 76.8% (logistic regression) to 81.7% (XGBoost). The full driver-integrated model achieved 89.4% accuracy, representing a 7.7 percentage-point improvement over the best practice-only baseline. McNemar's test indicates the improvement is statistically significant ($p < 0.01$).

4.2 Analysis of Results

The full cognitive-structural model achieved 89.4% accuracy and a macro F1-score of 0.88, indicating balanced performance across risk classes. The improvement over the best practice-only model (XGBoost, 81.7%) was statistically significant (McNemar's $\chi^2 = 9.42$, $p = 0.002$). This result answers Research Question 2 affirmatively: driver integration substantially improves classification performance relative to practice-only diagnostics.

SHAP feature importance analysis revealed that structural drivers collectively contributed approximately 58% of predictive weight, while cognitive drivers contributed 42%. The three highest-ranked individual features were: ESG metric ambiguity (mean $|\text{SHAP}| = 0.187$), disclosure discretion (0.163), and temporal discounting ratio (0.141). Governance weakness ranked fourth (0.128), followed by overconfidence language (0.112) and moral disengagement markers (0.094). This pattern addresses Research Question 3: structural drivers exhibit greater predictive weight, but cognitive drivers provide substantial incremental explanatory value—removing cognitive features from the model reduced accuracy from 89.4% to 84.1%, a 5.3 percentage-point decline.

The interaction between disclosure discretion and temporal discounting emerged as a particularly strong predictor, consistent with prospect theory's prediction that loss-averse managers exploit structural discretion when short-term pressures dominate. Firms high on both dimensions were classified as high-risk with 94.2% precision.

5. Discussion

5.1 Interpretation

Finding 1: Driver integration substantially improves diagnostic accuracy. The 89.4% accuracy of the full model compared to 81.7% for the best practice-only baseline demonstrates that understanding *why* firms misrepresent improves the ability to identify *that* they do so. This finding extends Okeke's (2026) diagnostic instrument, which validated the multidimensionality of misrepresentation practices but did not incorporate antecedent variables . It also addresses the gap identified by Lublóy et al. (2025), who noted that most greenwashing measures remain descriptive rather than explanatory . The result suggests that regulatory triage systems focused solely on disclosure content may miss firms whose structural and cognitive profiles place them at elevated risk even before problematic practices become salient.

Finding 2: Structural drivers carry greater predictive weight than cognitive drivers. The SHAP analysis indicates that metric ambiguity, disclosure discretion, and governance weakness collectively contribute more to classification than cognitive variables. This finding aligns with Janik and Ryszko's (2025) observation that greenwashing is substantially enabled by reporting-system characteristics that permit symbolic substitution . It also supports the legitimacy-theoretic prediction that structural opportunity shapes the form and prevalence of misrepresentation. However, the substantial incremental contribution of cognitive variables (5.3 percentage points) qualifies purely structural accounts: organizational conditions enable misrepresentation, but psychological processes influence its severity and persistence.

Finding 3: The discretion-discounting interaction identifies highest-risk profiles. The interaction between disclosure discretion and temporal discounting produced 94.2% precision in high-risk classification. This finding supports the prospect-theoretic mechanism: when managers

face short-term performance pressures and possess latitude over what to disclose, misrepresentation becomes a rational loss-avoidance strategy. The result contributes to theory by specifying the interaction conditions under which cognitive biases translate into organizational outcomes.

5.2 Implications

Academic implications: The study introduces a driver-based classification framework that moves greenwashing scholarship from typology toward explanation. It operationalizes cognitive constructs in a predictive modeling context, demonstrating their measurable incremental validity beyond structural and practice variables. The integration of prospect theory and legitimacy theory provides a template for future work examining how psychological and institutional mechanisms interact to produce misrepresentation.

Practical implications: Regulators and assurance providers can deploy the model's driver categories to differentiate scrutiny intensity. Firms exhibiting structural-driver profiles (high metric ambiguity, high disclosure discretion, weak governance) warrant disclosure-architecture interventions: standardized environmental metrics, mandatory assurance, and board-level accountability requirements. Firms exhibiting cognitive-driver profiles (high overconfidence language, short-term orientation, moral disengagement markers) may benefit from decision-process interventions: incentive structures that reward long-term accuracy over short-term impression management, and managerial training in ethical disclosure reasoning. The model's 89.4% accuracy supports triage use, though firm-level classifications should be validated through independent evidence before regulatory action.

5.3 Limitations

1. **Proxy measurement of cognitive constructs.** Overconfidence, temporal discounting, and moral disengagement are operationalized through disclosure-language proxies rather than direct psychological assessment. While these proxies follow established precedents, they capture expressed cognition rather than underlying cognitive processes. Future research using experimental or survey methods could validate the proxy measures.
2. **Retrospective design and historical stability assumption.** The model assumes that driver-performance relationships observed in 2019–2024 data remain stable. Regulatory changes—including the UK FCA anti-greenwashing rule and evolving EU sustainability reporting standards—may alter the relationship between drivers and outcomes. Prospective validation is needed.
3. **Sample composition and generalizability.** The sample is concentrated in larger, listed firms with formal reporting architectures in the UK and EU. Transferability to small and medium enterprises, privately held firms, and non-European regulatory contexts requires replication and invariance testing.

5.4 Future Research Directions

1. Extension of the driver-categorization framework to social and governance dimensions of ESG misrepresentation, testing whether cognitive and structural drivers exhibit domain-specific patterns.
2. Longitudinal design examining whether regulatory interventions targeting structural drivers (e.g., mandatory assurance) reduce misrepresentation severity more effectively than interventions targeting cognitive drivers.
3. Experimental validation of cognitive proxies using managerial decision-making simulations to establish convergent validity between disclosure-language indicators and measured psychological constructs.
4. Application of the framework to emerging disclosure regimes in non-Western markets, examining whether the relative weight of structural versus cognitive drivers varies with institutional context.

6. Conclusion

This study developed and validated a diagnostic model that categorizes corporate environmental misrepresentation according to its cognitive and structural drivers, achieving 89.4% classification accuracy and demonstrating that driver integration substantially outperforms practice-only diagnostic approaches. The framework contributes a replicable methodology for moving beyond identifying misrepresentation to explaining its origins, with structural drivers—particularly ESG metric ambiguity and disclosure discretion—carrying the greatest predictive weight while cognitive drivers provide meaningful incremental validity. For administrators and regulators, the model offers a triage instrument that distinguishes between firms requiring disclosure-architecture reform and those requiring decision-process interventions, enabling targeted rather than generic governance responses. As sustainability disclosure regimes tighten globally, the capacity to diagnose not only whether firms misrepresent but *why* they do so will become increasingly essential to effective accountability governance.

References

1. Lublóy, Á., Keresztúri, J. L., & Berlinger, E. (2025). Quantifying firm-level greenwashing: A systematic literature review. *Journal of Environmental Management*, *373*, 123399. <https://doi.org/10.1016/j.jenvman.2024.123399>
2. Raspanti, F., Bardelli, A. P., Scala, S., Demirtaş, İ., Di Bari, M., & Filannino, M. (2026). Greench-v1: Distilling SLMs on greenwashing detection. In *Proceedings of the 2nd Workshop on Ecology, Environment, and Natural Language Processing* (pp. 79–86). European Language Resources Association. <https://doi.org/10.63317/2ngqn2gr78du>
3. Okeke, A. (2026). Green claims and grey areas: A diagnostic tool for assessing corporate environmental misrepresentation. *International Journal of Disclosure and Governance*, 1–17. <https://doi.org/10.1057/s41310-026-00371-1>
4. Okeke, A., & Ugbebor, I. (2026). Discursive governance and development goals: A performative theory of corporate purpose in sustainability discourse. *Business Strategy and the Environment*.
5. Wilson, B., Shallangwa, G. M., & Mela, S. L. (2026). Artificial intelligence models for detecting greenwashing in UK ESG and green finance projects. *World Journal of Advanced Research and Reviews*, *29*(1), 1261–1270. <https://doi.org/10.30574/wjarr.2026.29.1.0177>
6. Janik, A., & Ryszko, A. (2025). Greenwashing in sustainability reporting: A systematic literature review of strategic typologies and content-analysis-based measurement approaches. *Sustainability*, *18*(1), 17. <https://doi.org/10.3390/su18010017>
7. Chen, T., & He, T. (2025). Interpretable predictive model for listed companies ESG greenwashing based on XGBoost and SHAP. *Scientific Reports*, *15*, Article 42004. <https://doi.org/10.1038/s41598-026-42004-1>
8. Delmas, M. A., & Burbano, V. C. (2011). The drivers of greenwashing. *California Management Review*, *54*(1), 64–87.
9. Marquis, C., Toffel, M. W., & Zhou, Y. (2016). Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, *27*(2), 483–504.
10. Lyon, T. P., & Montgomery, A. W. (2015). The means and end of greenwash. *Organization & Environment*, *28*(2), 223–249.

11. Seele, P., & Gatti, L. (2017). Greenwashing revisited: In search of a typology and accusation-based definition incorporating legitimacy strategies. *Business Strategy and the Environment*, *26*(2), 239–252.
12. Kim, E.-H., & Lyon, T. P. (2015). Greenwash vs. brownwash: Exaggeration and undue modesty in corporate sustainability disclosure. *Organization Science*, *26*(3), 705–723.
13. Parguel, B., Benoît-Moreau, F., & Larceneux, F. (2011). How sustainability ratings might deter “greenwashing”: A closer look at ethical corporate communication. *Journal of Business Ethics*, *102*(1), 15–28.
14. Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review*, *20*(3), 571–610.
15. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.