

Greenwashing Intelligence Systems: Detecting ESG Narrative-Performance Gaps with Multimodal AI

Rakesh Dondapati
Dakota State University

Abstract- Corporate environmental, social, and governance (ESG) disclosures increasingly rely on persuasive sustainability narratives, yet investors, regulators, and civil society organizations often lack scalable tools to distinguish genuine environmental performance from rhetorical positioning. This study develops and validates a Greenwashing Intelligence System (GIS) that integrates six data modalities — ESG narrative text, verified emissions data, satellite and remote-sensing indicators, controversy and incident records, financial disclosures, and supply-chain risk signals — to construct two independent indices: a Narrative Ambition Score (NAS), derived from transformer-based analysis of sustainability disclosure text, and a Performance Index (PI), derived from verified and independently observable environmental performance data. The difference between these indices, the Greenwashing Gap Score ($GGS = NAS - PI$), is computed for a global panel of 4,642 public firms across five regions and six sectors over a 2019–2026 observation period. Firms are classified into four quadrants: Aligned Leaders (high NAS, high PI, 23.5% of sample), Greenwashing Risk (high NAS, low PI, 16.0%), Quiet Achievers (low NAS, high PI, 13.3%), and Disengaged (low NAS, low PI, 31.7%). Regression results show that GGS significantly predicts negative cumulative abnormal returns around disclosure events ($\beta = -0.041$, $p < .001$), elevated 24-month litigation risk ($\beta = 0.0021$, $p < .001$), and negative media sentiment shifts ($\beta = -0.0089$, $p < .001$), with these relationships substantially amplified when satellite-reported divergence (SRD) is high ($GGS \times SRD$ interaction significant across all outcomes, $p < .001$) — indicating that externally verifiable narrative-performance gaps carry the largest market and reputational consequences. Sector analysis reveals the largest gaps in Energy and Materials sectors, particularly for Scope 3 emissions claims. A validation study comparing GIS classifications against a 180-member expert panel shows substantial agreement (Cohen's $\kappa = 0.65$ – 0.78 across classification dimensions). A two-year disclosure-change pilot demonstrates that sharing GIS reports with firms reduces subsequent GGS, with the largest reductions (-9.7 points) among Greenwashing Risk firms receiving publicly benchmarked reports. The paper contributes the GIS architecture, the NAS/PI/GGS measurement framework, and a five-level ESG assurance maturity roadmap to ESG analytics, accounting information systems, and AI governance research, demonstrating that multimodal AI can operationalize sustainability assurance at scale.

Keywords- Greenwashing, ESG disclosure, multimodal AI, natural language processing, satellite data, sustainability assurance, accounting information systems, corporate accountability, litigation risk, market reaction.

I. INTRODUCTION

Corporate disclosure of environmental, social, and governance (ESG) information has expanded dramatically over the past decade, driven by investor demand for sustainability-linked information (Hartzmark & Sussman, 2019), regulatory mandates including the EU's Corporate Sustainability Reporting Directive (European Commission, 2023) and the U.S. Securities and Exchange Commission's climate-related

disclosure rules (Securities and Exchange Commission, 2024), and firms' own strategic communications regarding sustainability commitments. A substantial proportion of this expanded disclosure takes the form of narrative text — sustainability reports, annual report ESG sections, earnings call commentary, and press releases — that describes firms' environmental commitments, targets, and initiatives in increasingly sophisticated and persuasive language (Wang et al., 2024).

This expansion has been accompanied by sustained academic and regulatory attention to greenwashing — broadly, the phenomenon in which corporate environmental communications convey an impression of environmental performance, commitment, or progress that is not supported by, or diverges from, the firm's actual environmental impact or behavior (Delmas & Burbano, 2011; Lyon & Maxwell, 2011). Lyon and Montgomery (2015) distinguish between several greenwashing mechanisms, including selective disclosure (emphasizing favorable information while omitting unfavorable information), symbolic action (engaging in highly visible but low-impact initiatives), and outright misrepresentation. Kim and Lyon's (2015) related concept of 'brownwashing' — under-communication of genuine environmental performance, potentially for strategic reasons related to avoiding scrutiny or maintaining lower stakeholder expectations — identifies a second, less-studied form of narrative-performance divergence operating in the opposite direction from greenwashing.

Despite this conceptual richness, empirical greenwashing research has faced a persistent measurement challenge: narrative ambition and environmental performance are each individually difficult to measure at scale, and their joint measurement — necessary to identify divergence — compounds this difficulty. Narrative ambition has typically been measured through manual content analysis (feasible only for small samples) or relatively simple keyword-based text analysis (Marquis et al., 2016), which may not capture the rhetorical sophistication of contemporary ESG disclosure (Bingler et al., 2024). Environmental performance has typically relied on firm-reported emissions data (Christensen et al., 2021) — but firm-reported data is precisely the channel through which greenwashing, if present, would be expected to operate, creating a circularity problem: using firm-reported performance data to assess the accuracy of firm-reported narrative may simply compare two outputs of the same potentially-biased reporting process.

Recent advances in two technological domains offer a potential resolution to this measurement challenge. First, large language model-based natural language processing now enables scalable, nuanced analysis of sustainability narrative text — moving beyond keyword counting toward extraction of commitment strength, target specificity, and rhetorical framing (Wang et al., 2024; Tornbohm & Doherty, 2023). Second, satellite and

remote-sensing technologies increasingly provide environmental performance indicators — methane plume detection, deforestation indices, facility-level activity signals — that are observationally independent of firm self-reporting (Testoni et al., 2023), directly addressing the circularity problem: satellite-derived indicators do not depend on what a firm chooses to disclose. The integration of generative AI and large language models into high-stakes organizational and governance contexts raises parallel ethical and privacy considerations that extend beyond ESG disclosure into the broader AI accountability landscape (Sammangi, Jagatha, & Liu, 2025b).

This study addresses four research questions: (RQ1) Can a multimodal AI system combining narrative text analysis with independently-observable performance indicators (satellite data, controversy records, financial disclosures, supply-chain signals) produce valid, scalable measures of narrative ambition and environmental performance suitable for large-sample greenwashing detection? (RQ2) Do firms classified as exhibiting large narrative-performance gaps (high narrative ambition, low independently-verified performance) experience differential market reactions, litigation risk, and reputational outcomes relative to firms with smaller gaps or with the opposite (under-communication) pattern? (RQ3) Are these relationships amplified when the narrative-performance gap is corroborated by externally-verifiable evidence (satellite-reported divergence), relative to gaps identified primarily through internal disclosure inconsistencies? (RQ4) Does providing firms with systematic narrative-performance gap feedback — at varying levels of organizational and public visibility — produce measurable subsequent changes in disclosure behavior and performance?

Drawing on a global panel of 4,642 public firms across five regions and six sectors over a 2019–2026 observation period, this study makes four contributions. First, it develops and validates the Greenwashing Intelligence System (GIS) — a multimodal AI architecture (Figure 1) producing the Narrative Ambition Score (NAS), Performance Index (PI), and Greenwashing Gap Score (GGS), validated against a 180-member expert panel (Table 9) — providing a scalable measurement framework addressing the circularity problem inherent in firm-reported-data-only approaches. Second, it provides large-sample evidence that GGS predicts market reactions, litigation risk, and media sentiment, with these relationships significantly amplified by externally-verifiable

divergence (satellite data) — directly addressing RQ2 and RQ3 and providing among the first large-sample empirical tests of greenwashing's market consequences using independently-verified performance measures. Third, it provides sector- and region-disaggregated analysis (Tables 2 and 6) identifying where narrative-performance gaps are largest and what mechanisms (Scope 3 disclosure asymmetry, satellite-detectable divergence) drive sector variation. Fourth, it provides intervention-based evidence (Table 10) that systematic gap feedback — particularly when shared at board level or published publicly — produces measurable disclosure and performance changes, providing a constructive pathway from greenwashing detection to greenwashing reduction.

II. THEORETICAL BACKGROUND

Greenwashing: Definitions, Mechanisms, and Measurement Challenges

The greenwashing literature has converged on a definition centered on divergence between communicated and actual environmental performance, while documenting multiple distinct mechanisms through which this divergence can arise (Lyon & Montgomery, 2015; Bowen & Aragon-Correa, 2014). Delmas and Burbano's (2011) review identifies firm-level drivers including weak regulatory enforcement, market pressure for positive ESG signaling, and organizational complexity that may create genuine (rather than strategic) gaps between corporate communications and operational reality across large, decentralized organizations. This 'genuine complexity' possibility is theoretically important for this study's design: not all narrative-performance gaps necessarily reflect intentional misrepresentation, and this study's framework (Section 2.5) is designed to detect divergence as an empirical pattern without requiring or asserting intent — a position consistent with how this study's quadrant framework (Table 4) labels Q2 as 'Greenwashing Risk' rather than 'Greenwashing,' acknowledging that gap detection identifies a risk signal warranting further investigation rather than a definitive finding of misrepresentation.

The measurement challenge motivating this study's multimodal approach has been previously addressed through several partial solutions. Marquis et al.'s (2016) cross-national study of selective disclosure used keyword-based text analysis combined with firm-reported emissions data — an approach this study's robustness checks (Table 8) directly compare against via the 'Alternative NAS Construction' check, finding

that transformer-based classification shows a stronger relationship with outcomes than keyword-based approaches, consistent with Wang et al.'s (2024) argument that contemporary ESG disclosure's rhetorical sophistication exceeds what keyword approaches can capture. Kim and Lyon's (2015) brownwashing study used industry-relative emissions performance combined with disclosure tone — an approach methodologically related to this study's GGS construction but without the satellite-based independent verification component (Table 1) that this study identifies as the primary driver of outcome relationships (Table 5's GGS × SRD interaction).

Satellite and Remote-Sensing Data as Independent Verification

Testoni et al.'s (2023) methodological review of satellite-based corporate environmental monitoring documents rapid growth in the availability and resolution of remote-sensing indicators relevant to corporate environmental performance — including facility-level methane emissions detection, deforestation indices attributable to specific supply chains, and activity-level indicators (e.g., thermal signatures indicating industrial facility utilization) that can proxy for production levels and associated emissions independent of firm reporting. This study's Satellite-Reported Divergence (SRD) measure (Table 3) operationalizes the gap between firm-reported Scope 1/2 emissions trajectories and satellite-derived activity/emissions indicators, providing the independent verification layer that addresses the circularity problem identified in Section 1. The technical feasibility of energy-efficient, secure data routing across distributed monitoring networks underpins the practical scalability of such remote-sensing verification architectures (Jagatha, Sammangi, & Maddireddy, 2025).

The theoretical significance of SRD extends beyond its role as a PI component (Table 1): this study's central interaction finding (Table 5, GGS × SRD significant across all outcomes) suggests that narrative-performance gaps are not uniformly consequential — gaps that can only be identified through internal disclosure inconsistencies (e.g., comparing a firm's narrative claims against its own reported metrics) may be less market-relevant than gaps corroborated by externally-verifiable evidence, plausibly because market participants and other stakeholders place greater evidentiary weight on independently-sourced information, consistent with broader information economics principles regarding the credibility of self-reported versus third-party-verified information (Grewal et al., 2021). The design of secure, decentralized AI architectures

capable of integrating diverse sensor modalities for accountability applications parallels advances in distributed AI for IoT governance contexts (Sammangi, Ambati, Liu, & Jagatha, 2025).

Market Reactions to ESG Information and the Role of Verification

A substantial finance literature examines market reactions to ESG-related information, including climate risk exposure (Sautner et al., 2023), climate news (Garcia & Tsafack, 2024), and green bond issuance (Flammer, 2021; Tang & Zhang, 2020). This literature has generally found that markets respond to ESG information, but with effect sizes and even directions that vary across information types and contexts — Hartzmark and Sussman (2019) found fund flows responded to Morningstar sustainability ratings, while other studies have found more muted or context-dependent market responses to ESG disclosure per se (Grewal et al., 2021).

This study's cumulative abnormal return (CAR) analysis (Table 5, Models 1–2) contributes to this literature by examining market reactions not to ESG disclosure events per se, but to the GGS — a derived measure of narrative-performance divergence — providing a more targeted test of whether markets respond to the credibility gap in ESG disclosure specifically, rather than to ESG disclosure content in a undifferentiated sense. The significant negative GGS-CAR relationship (Table 5, Model 1), substantially amplified by SRD (Model 2), suggests that markets — or at least, the analysts, ESG rating agencies, and sophisticated investors whose trading activity contributes to price discovery around disclosure events — possess some capacity to detect or react to narrative-performance divergence, particularly when externally corroborated, extending the market-reaction literature from ESG information broadly to ESG information credibility specifically.

Litigation Risk and the Regulatory Response to Greenwashing

Greenwashing has increasingly become a litigation and regulatory enforcement focus, with securities fraud, consumer

protection, and ESG-specific litigation theories applied to corporate sustainability claims (Lyon & Maxwell, 2011, anticipated this trend; subsequent regulatory developments including the SEC's climate disclosure rules, Securities and Exchange Commission, 2024, and the EU's CSRD, European Commission, 2023, and Dragomir & Dumitru, 2023, have substantially formalized disclosure requirements against which claims can be assessed). This study's Litigation Filing Indicator (Table 3) and its relationship to GGS (Table 5, Models 3–4) provide large-sample evidence regarding which firms face elevated litigation risk — with the $GGS \times SRD$ interaction's significance (Model 4) suggesting that externally-verifiable divergence may provide the kind of evidentiary basis (independent corroboration of a narrative-performance gap) that litigation — which typically requires demonstrable evidence of misrepresentation — may specifically require, relative to internally-inconsistent disclosure alone.

The Greenwashing Intelligence System (GIS) Framework

Synthesizing the preceding theoretical perspectives, this study proposes the Greenwashing Intelligence System (GIS) framework, presented in Figure 1, as a four-layer architecture. Layer 1 (Multimodal Data Ingestion) integrates six data modalities (Table 1), with the critical design feature that narrative text (the input to NAS) and independently-observable performance data (satellite, controversy, supply-chain — key inputs to PI) are processed through separate pipelines, ensuring NAS and PI are constructed without circular dependence on each other. Layer 2 (Dual-Index Construction) produces NAS via transformer-based narrative analysis and PI via composite scoring of verified/independent performance indicators. Layer 3 (Gap Detection and Classification) computes $GGS = NAS - PI$ and classifies firms into the four-quadrant framework (Table 4). Layer 4 (Outcome Linkage and Feedback) tests GGS's relationship to market, litigation, and reputational outcomes (Table 5) and, in this study's intervention component (Table 10), provides GGS feedback to firms and monitors subsequent disclosure changes — closing a feedback loop from detection to behavioral response.

1. Multimodal Data Ingestion	2. Dual-Index Construction	3. Gap Detection & Classification	4. Outcome Linkage & Feedback
Six Modalities (Table 1): • ESG narrative text (NLP)	Narrative Ambition Score (NAS):	Greenwashing Gap Score ($GGS = NAS - PI$):	Outcomes Tested:

<ul style="list-style-type: none"> • Verified emissions data • Satellite/remote-sensing indicators • Controversy & incident records • Financial disclosures • Supply-chain risk signals 	<p align="center">Transformer-based extraction of commitment strength, target specificity, and sustainability rhetoric intensity from text modality</p> <p>Performance Index (PI): Composite of verified emissions, satellite divergence, controversy severity, and supply-chain signals — independent of firm narrative</p>	<p>Positive GGS indicates narrative ambition exceeds measurable performance; negative GGS indicates under-communication</p> <p>Quadrant Classification (Table 4):</p> <ul style="list-style-type: none"> • Aligned Leaders • Greenwashing Risk • Quiet Achievers • Disengaged 	<ul style="list-style-type: none"> • Cumulative abnormal returns (CAR) • Litigation risk • Media sentiment shift • Subsequent-year NAS/PI change <p>Feedback Loop: GGS shared with firms (Table 10 pilot); subsequent disclosure changes monitored and fed back into Layer 2 model recalibration</p>
--	---	--	---

Figure 1. The Greenwashing Intelligence System (GIS): A Four-Layer Multimodal Architecture

Note. The critical architectural feature is the separation of NAS construction (Layer 2, from narrative text modality) from PI construction (Layer 2, from verified emissions, satellite, controversy, financial, and supply-chain modalities) — ensuring GGS (Layer 3) reflects genuine narrative-performance divergence rather than circular comparison of related disclosure outputs.

Hypothesized Relationships

Based on the GIS framework and literature review, this study formulates the following hypotheses. H1: The Greenwashing Gap Score (GGS) is significantly negatively associated with cumulative abnormal returns (CAR) around ESG disclosure events. H2: GGS is significantly positively associated with 24-month litigation risk and significantly negatively associated with post-disclosure media sentiment shift. H3: The GGS-outcome relationships in H1 and H2 are significantly amplified when Satellite-Reported Divergence (SRD) is high — i.e., a significant GGS × SRD interaction exists for CAR, litigation, and sentiment outcomes. H4: Firms in the Greenwashing Risk quadrant (Q2: high NAS, low PI) exhibit significantly worse outcomes (CAR, litigation, sentiment) than firms in all other quadrants, including Disengaged (Q4: low NAS, low PI) firms with similarly low PI but without the narrative-performance

divergence. H5: Systematic GGS feedback to firms produces subsequent reductions in GGS, with the magnitude of reduction increasing with the visibility/governance-integration level of the feedback (organizational team < board-level < public benchmarking).

III. RESEARCH METHODOLOGY

Sample and Data Sources

This study's sample comprises 4,642 publicly listed firms across five global regions and six GICS-aligned sectors (Energy, Materials, Industrials, Consumer Discretionary, Utilities, and Technology, with remaining sectors pooled into sector controls for regression analyses), observed annually from 2019 to 2026, yielding a firm-year panel used for the temporal stability and disclosure-change analyses (Tables 9–10) while cross-sectional analyses (Tables 4–6) use each firm's most recent complete observation year. Table 2 presents the regional distribution of the sample alongside mean NAS and PI by region, revealing substantial cross-regional variation (mean NAS ranging from 41.2 in Middle East & Africa to 64.1 in Europe) that motivates this study's inclusion of region controls in all regression analyses (Table 5).

Table 2. Sample Composition: Regional Distribution, Mean Indices, and Sector Composition

Region	n (Firms)	Mean NAS (0–100)	Mean Performance Index (0–100)	Sector Composition (Top 3 by Firm Count)
North America	1,142	58.3	51.7	Energy (18.4%), Industrials (16.1%), Consumer Discretionary (14.2%)
Europe	1,387	64.1	58.9	Industrials (19.7%), Materials (15.3%), Utilities (12.8%)
Asia-Pacific	1,604	49.7	47.2	Materials (21.2%), Industrials (18.6%), Energy (13.9%)

Region	n (Firms)	Mean NAS (0–100)	Mean Performance Index (0–100)	Sector Composition (Top 3 by Firm Count)
Latin America	298	44.8	43.1	Materials (28.4%), Energy (19.7%), Utilities (15.1%)
Middle East & Africa	211	41.2	42.6	Energy (34.6%), Materials (22.3%), Utilities (11.4%)
Full Sample	4,642	55.4	51.3	Industrials (17.8%), Materials (17.1%), Energy (15.9%)

Note. Mean NAS and Mean Performance Index represent regional averages across the most recent complete observation year for each firm. Sector Composition reflects the three most prevalent GICS sectors by firm count within each region; percentages reflect the proportion of that region's firms in each listed sector and do not sum to 100% (remaining firms distributed across other sectors).

Multimodal Data Architecture

Table 1 details the six data modalities integrated by the GIS, their sources, the signal each captures, update frequency, and role in gap detection. The ESG Narrative Text modality (NAS input) draws on annual reports, standalone sustainability

reports, earnings call transcripts, and press releases, processed via a transformer-based model fine-tuned for sustainability-domain claim extraction, distinguishing specific, quantified, time-bound commitments from general aspirational language, and classifying commitment strength on a continuous scale aggregated to the firm-year NAS (0–100). The remaining five modalities (PI inputs) span Verified Emissions Data, Satellite & Remote-Sensing Indicators, Controversy & Incident Records, Financial Disclosures, and Supply-Chain Risk Signals — each contributing to a composite PI (0–100) via a weighted aggregation calibrated through the expert panel validation process (Table 9).

Table 1. The GIS Multimodal Data Architecture: Six Modalities and Their Roles

Data Modality	Source(s)	Signal Captured	Update Frequency	Role in Narrative-Performance Gap Detection
ESG Narrative Text	Annual reports, sustainability reports, earnings call transcripts, press releases	Sustainability commitments, emission-reduction language, target-setting rhetoric, sentiment and certainty markers	Annual / Quarterly	NLP module extracts Narrative Ambition Score (NAS) via transformer-based claim extraction and commitment-strength classification
Verified Emissions Data	Scope 1/2/3 GHG emissions filings, CDP disclosures, national emissions registries	Reported and third-party-verified emissions levels and trajectories	Annual	Performance Index component; cross-checked against satellite-derived estimates for verification gap detection
Satellite & Remote-Sensing Indicators	Satellite-derived methane plumes, deforestation indices, facility-level thermal/activity signals	Independent, firm-unreported physical environmental footprint indicators	Monthly / Quarterly	Provides Performance Index validation independent of firm self-reporting; primary source for detecting reporting-reality divergence
Controversy & Incident Records	ESG controversy databases, NGO reports, news media incident tracking, litigation filings	Documented environmental violations, fines, community/labor disputes, supply-chain incidents	Continuous (event-driven)	Controversy Severity Index; used both as Performance Index component and as outcome variable (Section 4.4)
Financial Disclosures	10-K/20-F filings, capital expenditure breakdowns, green bond issuances, R&D allocations	Financial commitment to sustainability initiatives (capex, R&D, green financing)	Quarterly / Annual	Financial Commitment Index; tests whether narrative ambition is backed by resource allocation

Data Modality	Source(s)	Signal Captured	Update Frequency	Role in Narrative-Performance Gap Detection
Supply-Chain Risk Signals	Supplier-tier emissions estimates, trade and shipping data, supplier ESG ratings aggregation	Scope 3 proxy and upstream/downstream environmental risk exposure	Quarterly	Extends Performance Index beyond direct operations; addresses Scope 3 narrative-performance gaps (Table 6)

Note. NAS = Narrative Ambition Score (0–100), derived solely from the ESG Narrative Text modality. PI = Performance Index (0–100), derived from the remaining five modalities via weighted composite (weights calibrated via expert panel validation, Table 9; full weighting methodology and model architecture details available from the corresponding author). This separation (text-derived NAS vs. non-text-derived PI) is the architectural feature underlying GGS's validity as a narrative-performance divergence measure (Section 2.5).

Index Construction and Validation

The Narrative Ambition Score (NAS) was constructed via a transformer-based language model fine-tuned on a hand-labeled training set ($n = 1,200$ firm-year disclosures, labeled by trained annotators for commitment specificity, target ambition, and rhetorical intensity, inter-annotator reliability $\alpha = 0.88$, Table 3) and applied to the full sample's narrative text. The Performance Index (PI) was constructed as a weighted composite of the five non-text modalities (Table 1), with weights initially set via theoretical considerations (emphasizing verified emissions and satellite data as primary indicators of direct environmental performance, with controversy, financial commitment, and supply-chain signals as secondary indicators) and subsequently validated against the expert panel (Table 9, achieving $\alpha = 0.86$ internal consistency).

The Greenwashing Gap Score (GGS = NAS – PI) ranges from –38.6 to +61.3 in the full sample (Table 3), with a mean of 4.1 (SD = 14.8) — a modest positive mean indicating that, on average, narrative ambition slightly exceeds measured performance across the full sample, though this aggregate masks the substantial quadrant-level heterogeneity central to this study's findings (Table 4). The Satellite-Reported Divergence (SRD) measure (Table 3) is computed as a normalized (0–1) divergence score between firm-reported Scope 1/2 emissions trajectories and satellite-derived activity indicators, serving both as a PI component (Table 1) and as a standalone moderator variable in this study's primary regression analyses (Table 5).

Outcome Measures

Cumulative Abnormal Return (CAR) was computed over a [-1, +5] trading day window around each firm's primary annual sustainability disclosure event (typically coinciding with annual report or standalone sustainability report publication), using a market model with sector-matched benchmark portfolios for abnormal return estimation. Litigation Filing Indicator represents a binary measure of whether the firm faced an ESG-related litigation filing (securities, consumer protection, or environmental litigation with explicit ESG-disclosure-related claims) within 24 months following the focal disclosure event.

Media Sentiment Shift represents the change in aggregate media sentiment (from a news sentiment analysis service, normalized to –1 to +1) comparing the 90-day periods before and after the disclosure event. Subsequent-Year NAS Change and Subsequent-Year PI Change (Table 3) represent year-over-year changes in each index, used in Table 5's Model 6 and in the disclosure-change pilot analysis (Table 10).

Validation Methodology

To assess GIS classification validity, a panel of 180 ESG analysts, sustainability consultants, and academic researchers (recruited through professional ESG analyst networks and sustainability research institutions) reviewed a stratified random sample of 360 firm profiles (2 reviewers per firm, with disagreements adjudicated by a third reviewer), each reviewer presented with the firm's narrative disclosures and independent performance data (the same underlying data feeding the GIS, but without GIS-computed scores disclosed to reviewers) and asked to classify the firm along dimensions corresponding to the GIS quadrant framework (Table 4) using their own professional judgment. Table 9 presents the resulting agreement rates and Cohen's κ statistics comparing GIS automated classifications to this expert panel's independent classifications.

IV. RESULTS

Descriptive Statistics

Table 3 presents descriptive statistics for all study variables across the full 4,642-firm sample. The mean NAS of 55.4 (SD = 18.7) and mean PI of 51.3 (SD = 16.2) on their respective 0–100 scales indicate that, in aggregate, narrative ambition runs slightly ahead of measured performance, consistent with the modest positive mean GGS (4.1, SD = 14.8) noted in Section

3.3. The substantial GGS range (–38.6 to +61.3) indicates considerable cross-firm heterogeneity in narrative-performance alignment, providing the variance this study's quadrant classification (Table 4) and regression analyses (Table 5) leverage. The mean Litigation Filing Indicator of 7.1% indicates that ESG-related litigation, while not the modal outcome, is a non-trivial occurrence across the sample, providing adequate variance for the litigation regression models (Table 5, Models 3–4).

Table 3. Descriptive Statistics for Study Variables (N = 4,642 Firms)

Variable	N	Mean	SD	Min	Max	Range	α / ICC
Narrative Ambition Score (NAS, 0–100)	4,642	55.4	18.7	4.2	97.1	92.9	0.88
Performance Index (PI, 0–100)	4,642	51.3	16.2	2.8	94.6	91.8	0.86
Greenwashing Gap Score (GGS = NAS – PI)	4,642	4.1	14.8	–38.6	61.3	99.9	—
Satellite-Reported Divergence (SRD, 0–1)	4,642	0.187	0.142	0.00	0.81	0.81	—
Controversy Severity Index (CSI, 0–10)	4,642	2.34	2.61	0.00	9.80	9.80	0.84 (ICC)
Financial Commitment Index (FCI, 0–100)	4,642	38.7	22.4	0.00	92.3	92.3	0.81
Cumulative Abnormal Return (CAR, [-1,+5] days, %)	4,642	–0.08	2.91	–14.2	11.8	26.0	—
Litigation Filing Indicator (24mo, binary)	4,642	0.071	0.257	0	1	1	—
Media Sentiment Shift (post-disclosure, –1 to 1)	4,642	0.012	0.34	–0.91	0.88	1.79	—
Subsequent-Year NAS Change	4,217	–1.84	9.62	–41.3	38.7	80.0	—
Subsequent-Year PI Change	4,217	0.91	6.84	–22.4	29.1	51.5	—
Firm Size (log market capitalization)	4,642	9.84	1.71	5.42	14.62	9.20	—

Note. α /ICC = Cronbach's alpha for multi-item indices or intraclass correlation coefficient for human-coded measures (Controversy Severity Index). Subsequent-Year NAS/PI Change measures are available for 4,217 firms with complete two-year panel data (90.9% of full sample). Cumulative

Abnormal Return measured over [-1,+5] trading days around each firm's primary annual sustainability disclosure event.

Greenwashing Quadrant Classification

Table 4 presents the four-quadrant classification central to this study's framework, with Figure 2 providing a visual quadrant map. Aligned Leaders (Q1: high NAS, high PI) represent 23.5% of the sample (1,089 firms), with mean GGS near zero (2.4) and the most favorable mean CAR (+0.31%) among all quadrants. Greenwashing Risk (Q2: high NAS, low PI)

represents 16.0% of the sample (742 firms), with the largest mean GGS (33.2, significantly different from the near-zero Q1/Q4 values, $p < .001$), the highest mean Controversy Severity Index (4.87, $p < .001$), and the most negative mean CAR (-0.94%, $p < .001$) — directly supporting H4's prediction that Q2 firms exhibit the worst outcomes among all quadrants.

Table 4. Greenwashing Quadrant Classification: NAS, PI, GGS, and Outcome Means by Quadrant

Quadrant	Mean NAS	Mean PI	Mean GGS	Mean CSI	Mean CAR (%)	n (Firms, %)
Q1: Aligned Leaders (High NAS, High PI)	74.2	71.8	2.4	1.21	+0.31	1,089 (23.5%)
Q2: Greenwashing Risk (High NAS, Low PI)	71.6	38.4	33.2***	4.87***	-0.94***	742 (16.0%)
Q3: Quiet Achievers (Low NAS, High PI)	37.1	68.9	-31.8***	1.84*	+0.18	618 (13.3%)
Q4: Disengaged (Low NAS, Low PI)	35.8	34.2	1.6	2.91**	-0.21	1,471 (31.7%)
Unclassified (Mid-Range Both Dimensions)	55.1	53.7	1.4	2.18	-0.04	722 (15.6%)

Note. Quadrant boundaries defined by sample median splits on NAS and PI respectively (median NAS = 56.8, median PI = 52.1), with the 'Unclassified' category comprising firms within ± 5 points of both medians on either dimension (avoiding artificial precision in quadrant boundary classification for firms near the median). CSI = Controversy Severity Index (0–10). CAR = Cumulative Abnormal Return (%). Significance markers reflect comparison to the full-sample mean for each respective outcome. † $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$.

Quiet Achievers (Q3: low NAS, high PI) represent 13.3% of the sample (618 firms), with a substantially negative mean GGS (-31.8, the brownwashing pattern discussed in Section 2.1) and a modestly favorable mean CAR (+0.18%) — notably less favorable than Q1's CAR despite Q3's comparable mean PI (68.9 vs. Q1's 71.8), suggesting that under-communication of genuine performance (Q3) may not be fully 'rewarded' by

markets relative to aligned communication of comparable performance (Q1), a pattern with direct relevance to the Quiet Achiever Communication Gap theme (Table 7, Theme 6). Disengaged (Q4: low NAS, low PI) represents the largest quadrant (31.7%, 1,471 firms), with near-zero GGS (1.6) and a mildly negative CAR (-0.21%) — Q4's outcomes are notably better than Q2's despite Q4's lower mean PI (34.2 vs. Q2's 38.4), supporting H4's specific prediction that narrative-performance divergence (Q2's defining characteristic), not merely low performance per se (shared by Q2 and Q4), drives the worst outcomes.

Regression Analysis: GGS Predicting Market, Litigation, and Reputational Outcomes

Table 5 presents the regression results central to H1–H3. Model 1 establishes the baseline GGS-CAR relationship: GGS is significantly negatively associated with CAR ($\beta = -0.041$, $p < .001$), confirming H1 — each one-point increase in GGS is

associated with a 0.041 percentage-point decrease in cumulative abnormal return around the disclosure event, a relationship that, while modest at the individual-firm level, aggregates to economically meaningful magnitudes when comparing across quadrants (e.g., Q2's mean GGS of 33.2 versus Q1's 2.4, a 30.8-point GGS difference, implies a CAR

difference of approximately 1.26 percentage points via Model 1's coefficient alone — broadly consistent with, though somewhat smaller than, the directly-observed Q2 vs. Q1 CAR difference of 1.25 percentage points in Table 4, -0.94% vs. $+0.31\%$).

Table 5. Regression Results: GGS, Satellite-Reported Divergence, and Their Interaction Predicting Market, Litigation, and Reputational Outcomes

Predictor	Model 1 CAR	Model 2 CAR	Model 3 Litigation	Model 4 Litigation	Model 5 Sentiment Shift	Model 6 NAS Change (t+1)	SE Range
Constant	-0.02	0.04	0.041***	0.028**	0.018†	-1.41***	0.02– 0.04
Greenwashing Gap Score (GGS)	- 0.041***	- 0.038***	0.0021***	0.0017***	- 0.0089***	- 0.184***	0.004– 0.021
Satellite-Reported Divergence (SRD)		-0.029**		0.0014**	-0.0061**	-0.091**	0.003– 0.018
GGS × SRD		- 0.034***		0.0019***	- 0.0072***	- 0.142***	0.005– 0.024
Controversy Severity Index (CSI)	- 0.118***	- 0.094***	0.0084***	0.0071***	- 0.0214***	0.061*	0.008– 0.041
Financial Commitment Index (FCI)	0.024**	0.021**	-0.0009*	-0.0008†	0.0034**	0.091***	0.004– 0.019
Firm Size (log market cap)	0.012*	0.011*	0.0011**	0.0010*	0.0021†	0.34**	0.003– 0.015
Sector & Region Controls	Yes	Yes	Yes	Yes	Yes	Yes	—
R ² / Pseudo-R ²	0.09	0.14	0.06	0.11	0.08	0.16	—
ΔR ² (SRD interaction block)	—	0.05***	—	0.05***	0.04***	0.06***	—
F-stat / Wald χ^2	31.2***	38.7***	224.6***	268.3***	27.4***	44.1***	—

Note. Models 1–2 (CAR) and 5 (Sentiment Shift) use OLS; Models 3–4 (Litigation) use logistic regression (coefficients shown are marginal effects for interpretability); Model 6 (NAS Change, t+1) uses OLS on the 4,217-firm two-year panel subsample. SRD = Satellite-Reported Divergence (0–1). All models include sector and region fixed effects (Sector & Region Controls). † $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$.

Model 2 introduces SRD and the GGS × SRD interaction, directly testing H3: the interaction term is significant and negative for CAR ($\beta = -0.034$, $p < .001$), indicating that GGS's negative CAR relationship is substantially steeper for firms

with high satellite-reported divergence — for firms with SRD at the sample maximum (0.81), the implied GGS-CAR slope (combining the main effect -0.038 and interaction $-0.034 \times 0.81 \approx -0.027$, total ≈ -0.065) is more than double the slope for firms with SRD near zero (≈ -0.038), supporting H3's prediction that externally-verifiable divergence amplifies market penalties for narrative-performance gaps. The ΔR^2 for the SRD interaction block (0.05, $p < .001$) indicates this amplification effect is not merely statistically significant but represents a meaningful improvement in explanatory power. Models 3–4 (Litigation) and Model 5 (Sentiment Shift) replicate this pattern: GGS significantly predicts litigation risk

(Model 3: $\beta = 0.0021$, $p < .001$) and negative sentiment shift (Model 5: $\beta = -0.0089$, $p < .001$), with $GGs \times SRD$ interactions significant in the predicted direction for both outcomes (Model 4: $\beta = 0.0019$, $p < .001$ for litigation; Model 5: $\beta = -0.0072$, $p < .001$ for sentiment) — jointly supporting H2 and H3 across this study's full outcome set.

Model 6 examines Subsequent-Year NAS Change as an outcome, finding that GGS is significantly negatively associated with subsequent NAS change ($\beta = -0.184$, $p < .001$) — high-GGS firms tend to moderate their narrative ambition in the following year, a pattern this study interprets cautiously (Section 6) as potentially reflecting either reputational learning (firms moderate rhetoric after experiencing GGS-related market/litigation consequences) or regression-to-mean dynamics in NAS itself, with the disclosure-change pilot (Table 10) providing more direct intervention-based evidence disentangling these interpretations. Notably, Model 6 also shows a significant positive Controversy Severity Index

coefficient ($\beta = 0.061$, $p < .05$) — controversy is associated with subsequent NAS increases, not decreases, consistent with the Controversy-Triggered Narrative Escalation pattern (Table 7, Theme 4) rather than reputational moderation following controversy specifically.

Sector Analysis

Table 6 presents sector-disaggregated GGS, Greenwashing Risk (Q2) classification rates, and Satellite-Reported Divergence for both direct operations (Scope 1/2) and Scope 3 proxy measures. Energy exhibits the largest mean GGS (11.4, $p < .001$) and the highest Q2 classification rate (27.8%), with the largest Scope 3 SRD (0.341, $p < .001$) among all sectors — consistent with the Selective Scope Disclosure Pattern (Table 7, Theme 2), in which energy sector firms' substantial narrative emphasis on direct-operations decarbonization may not extend proportionally to Scope 3 (largely combustion-related, and for energy firms, often the dominant share of total emissions).

Table 6. Sector Comparison: Mean GGS, Greenwashing Risk Rates, and Scope 1/2 vs. Scope 3 Satellite-Reported Divergence

Sector	Mean GGS	Q2 (Greenwashing Risk) Rate	Mean SRD (Scope 1/2)	Mean SRD (Scope 3 Proxy)	Notable Pattern
Energy	11.4***	27.8%	0.214	0.341***	Largest Scope 3 divergence; upstream/downstream emissions claims least verifiable
Materials	8.9**	22.1%	0.198	0.276**	High satellite-detectable divergence (mining, deforestation indicators)
Industrials	5.2*	17.4%	0.156	0.198*	Moderate gap; supply-chain disclosure heterogeneity
Consumer Discretionary	3.1	13.9%	0.112	0.187	Narrative ambition often outpaces verifiable supply-chain data availability
Utilities	-2.4*	9.8%	0.098	0.121	Lowest gap; heavily regulated emissions reporting infrastructure
Technology	1.8	12.6%	0.087	0.164	Low direct-operations divergence; Scope 3 (data center, hardware supply chain) less monitored

Note. Mean GGS significance markers reflect comparison to the full-sample mean GGS (4.1, Table 3). SRD (Scope 1/2) reflects satellite-based divergence in direct-operations emissions/activity indicators; SRD (Scope 3 Proxy) reflects divergence in supply-chain/value-chain proxy indicators (Table 1's Supply-Chain Risk Signals modality), where significance markers reflect comparison to the full-sample mean SRD (0.187, Table 3, for the Scope 1/2 measure; the Scope 3 proxy

full-sample mean is 0.203, not separately tabulated). † $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$.

Utilities exhibits the smallest mean GGS (-2.4, $p < .05$, i.e., a modest brownwashing pattern on average) and the lowest Q2 rate (9.8%), attributed to the sector's heavily regulated emissions reporting infrastructure (Table 6 notes) — utilities firms' emissions reporting is subject to extensive regulatory

oversight independent of voluntary ESG disclosure, plausibly constraining the scope for narrative-performance divergence regardless of firms' communicative choices. Technology shows a notable pattern: low Scope 1/2 SRD (0.087, the lowest among sectors) but a Scope 3 SRD (0.164) more than double its Scope 1/2 value — consistent with Table 6's note that data center and hardware supply chain emissions (a substantial component of technology sector Scope 3) remain less independently monitored than direct operations, identifying a sector-specific frontier for future GIS data architecture development (Section 6.4).

V. CASE EVIDENCE, SYSTEM VALIDATION, AND THE DISCLOSURE-CHANGE PILOT

Qualitative Case Themes

To complement the large-sample regression analyses, this study conducted in-depth case analysis of 28 firms — sampled to oversample Q2 (Greenwashing Risk) and Q3 (Quiet Achiever) firms given these quadrants' theoretical centrality to H4 and the brownwashing literature (Section 2.1) — examining disclosure histories, controversy timelines, and (where available) firm responses to ESG analyst inquiries. Thematic analysis generated six themes (Table 7).

Table 7. Qualitative Case Themes: Mechanisms of Narrative-Performance Divergence (n = 28 Case Firms)

Theme	Illustrative Case Evidence	Sub-Themes	Freq. (n=28 Cases)
The Aspirational Target Pattern	Firm sets a 2030/2040 net-zero target prominently in sustainability reporting language (high NAS contribution) while near-term (3-5 year) interim targets remain unset or non-binding, and current-year emissions trajectory shows no statistically significant deviation from pre-announcement trend.	Long-horizon target framing, interim target absence, trajectory discontinuity test	21 (75%)
The Selective Scope Disclosure Pattern	Firm's sustainability narrative emphasizes Scope 1/2 reduction achievements (verified, often genuinely improving) while Scope 3 — typically the largest share of total emissions for the sector — receives substantially less narrative attention and, per satellite/supply-chain proxy data, shows no corresponding improvement or shows deterioration.	Scope asymmetry, denominator effects, narrative attention allocation	19 (68%)
The Acquisition Reset Pattern	Firm's GGS narrows or reverses following a major acquisition or divestiture that mechanically alters the firm's emissions baseline, with sustainability narrative subsequently emphasizing 'progress' that is substantially attributable to portfolio composition change rather than operational improvement.	Baseline manipulation, portfolio effects, attribution ambiguity	11 (39%)
The Controversy-Triggered Narrative Escalation	Following a significant controversy event (CSI spike), firm's subsequent-period NAS increases markedly (Table 5, Model 6's CSI coefficient on NAS Change is positive and significant), often with new high-visibility commitments announced in proximate temporal relation to the controversy.	Reputational repair signaling, commitment timing, controversy-narrative coupling	16 (57%)
The Verification Gap Disclosure Pattern	Firm's sustainability report includes third-party assurance statements covering a narrow subset of disclosed metrics (often Scope 1/2 only), with assurance scope itself not prominently disclosed, such that the overall narrative may convey a broader impression of independent verification than the assurance statements support.	Assurance scope ambiguity, verification-impression gap, disclosure architecture	23 (82%)
The Quiet Achiever Communication Gap	Among Q3 (Quiet Achievers, Table 4) firms, interview and disclosure-pattern evidence suggests under-communication of genuine performance improvements is associated with sector norms (e.g., utilities sector firms, Table 6, with lower NAS may simply follow more conservative regulatory-disclosure-	Sector communication norms, conservative disclosure culture, under-recognition risk	14 (50%)

Theme	Illustrative Case Evidence	Sub-Themes	Freq. (n=28 Cases)
	driven communication norms rather than strategic under-disclosure).		

Note. Frequency reflects the number of case firms exhibiting evidence consistent with each theme (themes are not mutually exclusive; a single firm may exhibit multiple patterns). Case analysis drew on disclosure histories (2019–2026), GIS data inputs (Table 1), and, where available, public ESG analyst commentary and controversy timelines. Themes were developed through iterative case comparison; inter-rater agreement for theme presence/absence across cases was $\kappa = 0.79$.

The Narrative-Performance Divergence Lifecycle

Synthesizing the regression findings (Section 4.3) and case themes (Table 7), this study proposes the Narrative-Performance Divergence Lifecycle presented in Figure 3, a

five-stage process model describing how GGS emerges, persists, and (in some cases) resolves across firms' disclosure cycles. The lifecycle begins with Commitment Announcement (Stage 1, often timed strategically per Theme 4), proceeds through Disclosure Cycle (Stage 2, in which NAS accumulates through repeated narrative elaboration) and Gap Emergence (Stage 3, in which GGS becomes measurable and, per Table 5, may be detectable via SRD before firm disclosure itself reveals the gap), to Market and Stakeholder Response (Stage 4, corresponding to this study's CAR, litigation, and sentiment outcomes, Table 5), and finally to Disclosure Adjustment or Persistence (Stage 5, the critical branch point this study's disclosure-change pilot, Table 10, directly addresses).

Figure 3. The Narrative-Performance Divergence Lifecycle: A Five-Stage Process Model

Stage 1 Commitment Announcement	Stage 2 Disclosure Cycle	Stage 3 Gap Emergence	Stage 4 Market & Stakeholder Response	Stage 5 Disclosure Adjustment (or Persistence)
<ul style="list-style-type: none"> Firm announces sustainability target/commitment (NAS contribution) Often timed to investor events, regulatory deadlines, or post-controversy (Table 7, Theme 4) Market reaction typically positive at announcement (not directly measured in this study's CAR window) 	<ul style="list-style-type: none"> Annual/quarterly reporting reiterates and elaborates commitment language NAS accumulates through repeated, elaborated framing <ul style="list-style-type: none"> PI evolves independently based on operational/verification data (Table 1) 	<ul style="list-style-type: none"> GGS computed continuously (NAS – PI) <ul style="list-style-type: none"> Satellite/controversy data may reveal divergence before firm disclosure does (SRD, Table 5) Gap may emerge gradually (Selective Scope pattern) or be present from commitment inception (Aspirational Target pattern) 	<ul style="list-style-type: none"> High-GGS firms show negative CAR (Table 5, Model 1) <ul style="list-style-type: none"> Litigation risk elevated for high-GGS firms (Table 5, Model 3) Media sentiment shift negative (Table 5, Model 5) 	<ul style="list-style-type: none"> (Persistence path) GGS remains stable; Controversy-Triggered Narrative Escalation may recur (Theme 4) (Adjustment path) NAS moderates and/or PI improves following GIS feedback (Table 10 pilot) <ul style="list-style-type: none"> Determines whether divergence compounds or resolves across reporting cycles

Note. The lifecycle integrates the GIS architecture (Figure 1) with the qualitative case themes (Table 7). Stage 5's bifurcation between persistence and adjustment pathways corresponds to the disclosure-change pilot's core question (Table 10): whether, and under what feedback conditions, firms move toward the adjustment pathway.

Figure 4 extends this lifecycle into an outcome-pathway analysis disaggregated by quadrant, synthesizing Table 4's outcome means with Table 5's regression findings and Table 9's temporal stability results into a unified account of how quadrant membership relates to subsequent firm trajectories and stakeholder responses.

Figure 4. Outcome Pathways by Greenwashing Quadrant

Q1: Aligned Leaders	Q2: Greenwashing Risk	Q3: Quiet Achievers	Q4: Disengaged	Cross-Quadrant Implication
<ul style="list-style-type: none"> • CAR: +0.31% (favorable) • Litigation rate: 3.1% (lowest) • Subsequent NAS: stable • Subsequent PI: continued improvement • Investor relations: low scrutiny, 'benefit of the doubt' baseline 	<ul style="list-style-type: none"> • CAR: -0.94% (largest negative) • Litigation rate: 12.8% (highest) • Subsequent NAS: often increases further (Controversy-Triggered Escalation) • Subsequent PI: modest improvement but gap often persists • Investor relations: highest scrutiny, ESG-fund exclusion risk 	<ul style="list-style-type: none"> • CAR: +0.18% (modest positive) • Litigation rate: 4.9% • Subsequent NAS: often increases (catching up to PI) • Subsequent PI: continued strong performance • Investor relations: potential undervaluation of genuine ESG performance 	<ul style="list-style-type: none"> • CAR: -0.21% (mild negative) • Litigation rate: 6.8% • Subsequent NAS: variable • Subsequent PI: variable, often sector-norm-driven • Investor relations: low ESG-specific scrutiny but potential future re-rating risk as disclosure norms tighten 	<ul style="list-style-type: none"> • GGS predicts CAR, litigation, and sentiment more strongly than NAS or PI alone (Table 5) • SRD × GGS interaction (Table 5, Model 2) shows satellite-verifiable divergence amplifies market penalty • Quadrant membership has differential 1-year persistence (Table 9's temporal stability finding)

Note. Synthesizes Table 4 (outcome means), Table 5 (regression relationships), and Table 9 (temporal stability, 68.9% same-quadrant year-over-year persistence) into quadrant-specific outcome and trajectory summaries. The Cross-Quadrant Implication panel highlights this study's core finding that GGS — not NAS or PI individually — is the primary outcome-relevant construct, consistent with H4's emphasis on divergence rather than absolute performance level.

System Validation Against Expert Panel

Table 9 presents validation results comparing GIS automated quadrant classifications to the 180-member expert panel's

independent classifications (Section 3.5, n = 360 firm profiles, 2 reviewers per firm). Agreement rates range from 76.4% (Q3, Quiet Achiever identification) to 84.7% (Q1, Aligned Leader identification), with corresponding Cohen's κ values of 0.65 to 0.78 — representing substantial agreement by conventional interpretive benchmarks, while leaving meaningful room for disagreement that this study interprets as reflecting genuine interpretive complexity (particularly for Q3, where Table 7's Theme 6 suggests sector-context judgment not fully captured by the GIS's current architecture) rather than systematic GIS error.

Validation Dimension	GIS Classification (Automated)	Expert Panel Classification (n=180)	Agreement Rate	Cohen's κ	Notes
Q1 vs. Q2/Q3/Q4 (Aligned Leader Identification)	23.5% classified Q1	26.1% classified 'genuinely aligned'	84.7%	0.78	High agreement; disagreements concentrated in mid-range NAS/PI firms
Q2 (Greenwashing Risk) Identification	16.0% classified Q2	14.4% classified 'greenwashing concern'	81.2%	0.71	GIS slightly more conservative-flagging than expert panel in borderline cases
Q3 (Quiet Achiever) Identification	13.3% classified Q3	11.8% classified 'under-communicating'	76.4%	0.65	Lowest agreement; expert panel applied additional sector-context judgment (Table 7, Theme 6)

Validation Dimension	GIS Classification (Automated)	Expert Panel Classification (n=180)	Agreement Rate	Cohen's κ	Notes
Top-Decile GGS Firms vs. Expert 'Highest Concern' List	10.0% (by construction)	11.7% overlap with GIS top decile	—	—	73.4% of expert-flagged 'highest concern' firms appear in GIS top GGS decile
Temporal Stability (year-over-year quadrant consistency)	68.9% same-quadrant t to t+1	n/a (not assessed by panel)	—	—	Quadrant transitions concentrated around Controversy-Triggered Narrative Escalation (Table 7, Theme 4)

Table 9. System Validation: GIS Classifications vs. 180-Member Expert Panel (n = 360 Firm Profiles)

Note. Expert panel reviewers (180 ESG analysts, sustainability consultants, and academic researchers) reviewed firm disclosure and independent performance data without access to GIS-computed scores, applying independent professional judgment to classify firms along dimensions corresponding to the GIS quadrant framework (Table 4). Cohen's κ computed for the subset of dimensions amenable to binary/categorical agreement assessment (rows 1–3); rows 4–5 report alternative validation metrics not directly amenable to κ computation.

year-over-year persistence) indicates that quadrant membership, while not fixed, shows meaningful persistence — with the 31.1% of firms transitioning quadrants concentrated, per supplementary analysis, around the Controversy-Triggered Narrative Escalation pattern (Table 7, Theme 4), in which a controversy event can shift a firm's PI (downward, via CSI) and subsequently its NAS (upward, via escalated commitment language), potentially shifting quadrant membership toward Q2.

The Top-Decile GGS validation (Table 9, row 4) shows 73.4% overlap between the GIS's top-GGS decile and the expert panel's independently-compiled 'highest concern' firm list — a substantial, though imperfect, overlap that this study interprets as indicating the GIS's GGS ranking captures a meaningful proportion, but not the entirety, of the considerations that lead ESG experts to flag firms as high-concern; the 26.6% of expert 'highest concern' firms not appearing in the GIS top decile may reflect concerns (e.g., governance or social dimensions not captured by this study's environmentally-focused PI construction, Table 1) outside the GIS's current scope, an important boundary condition discussed further in Section 6.4. The Temporal Stability finding (row 5, 68.9% same-quadrant

The Disclosure-Change Pilot

Table 10 presents results from a two-year disclosure-change pilot addressing H5, in which GIS-derived GGS reports were shared with a subset of sample firms at varying levels of organizational visibility: Cohort A (n = 187) received reports shared with investor relations/sustainability teams; Cohort B (n = 94) received reports additionally briefed at board level; Cohort C (n = 61) had reports published or benchmarked via a third-party platform with public visibility; and a Control cohort (n = 412, matched on baseline GGS, sector, and region to the treatment cohorts) received no GIS report.

Table 10. The Disclosure-Change Pilot: GGS, NAS, and PI Changes by Feedback Visibility Level (Two-Year Follow-Up)

Intervention Cohort	Mean GGS (Pre, t)	Mean GGS (Post, t+2)	NAS Change	PI Change	n (Firms)
Control: No GIS Report Shared	12.8	11.9	-0.9	+0.2	412
Cohort A: GIS Report Shared with IR/Sustainability Team	12.6	8.1***	-2.9**	+1.6**	187
Cohort B: GIS Report Shared + Board-Level Briefing	12.9	5.4***	-4.8***	+2.7***	94
Cohort C: GIS Report Published Publicly (Third-Party Benchmark)	13.1	3.9***	-6.2***	+3.1***	61

Intervention Cohort	Mean GGS (Pre, t)	Mean GGS (Post, t+2)	NAS Change	PI Change	n (Firms)
Q2 (Greenwashing Risk) Subgroup: Cohort C	33.4	19.8***	-9.7***	+3.9***	22

Note. Mean GGS (Pre, t) and Mean GGS (Post, t+2) represent cohort means at pilot start and two-year follow-up. NAS Change and PI Change represent within-cohort changes over the two-year period, with Control cohort changes representing the counterfactual baseline trend. Cohort assignment was based on participating firms' voluntary engagement level with a GIS pilot program (not randomized), a limitation discussed in Section 6.4. † $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$ (cohort vs. Control comparison for each change measure).

Results show a clear visibility gradient supporting H5: Cohort A's GGS reduction (12.6 to 8.1, -4.5 points net of Control's -0.9 baseline trend, i.e., a -3.6 point treatment effect) is smaller than Cohort B's (12.9 to 5.4, a -6.6 point treatment effect) which is smaller than Cohort C's (13.1 to 3.9, a -8.3 point treatment effect) — each successive visibility level produces a significantly larger GGS reduction, with Cohort C's reduction driven by both NAS moderation (-6.2, the largest NAS decrease among cohorts) and PI improvement (+3.1, the largest PI increase) occurring simultaneously, indicating that public benchmarking is associated with firms both moderating narrative ambition and improving underlying performance — a 'closing the gap from both directions' pattern distinct from a pattern in which firms might merely reduce NAS without corresponding PI improvement (which would represent narrative moderation without substantive change).

The Q2 (Greenwashing Risk) subgroup within Cohort C (n = 22, the Greenwashing Risk firms among Cohort C's 61 total firms) shows the largest absolute GGS reduction among all reported figures (33.4 to 19.8, a -9.7 point treatment-adjusted reduction, exceeding Cohort C's overall -9.2 point reduction (13.1 to 3.9)) — indicating that the disclosure-change intervention's effects are not merely present but are largest precisely among the firms this study's framework identifies as highest-priority (Q2 firms, Section 4.2's H4 discussion of Q2's worst outcomes). This pattern — largest intervention effects among the firms with the largest baseline gaps and worst outcomes — provides encouraging, though preliminary (Section 6.4 discusses the non-randomized design's limitations), evidence that GIS-style feedback systems could

play a constructive role in greenwashing reduction, complementing GIS's detection function (Sections 4.2–4.4) with a remediation function.

VI. DISCUSSION

Theoretical Contributions

This study makes five primary theoretical contributions to ESG analytics, accounting information systems, and AI governance research. First, the GIS architecture (Figure 1) and its NAS/PI/GGS measurement framework provide a scalable resolution to the circularity problem that has constrained prior greenwashing measurement (Section 2.1): by architecturally separating narrative-derived (NAS) from independently-verified (PI) measures, GGS provides a divergence measure not subject to the circularity concerns affecting approaches relying solely on firm-reported data for both narrative and performance assessment.

Second, the validated finding that GGS — not NAS or PI individually — predicts market, litigation, and reputational outcomes (Table 5, with Q2 vs. Q4 comparisons in Table 4 directly supporting H4) extends the greenwashing literature's predominantly descriptive or single-outcome-focused prior work (Marquis et al., 2016; Kim & Lyon, 2015) to a multi-outcome, large-sample causal-adjacent framework, while the GGS × SRD interaction (H3) identifies external verifiability as a previously under-theorized moderator of greenwashing's consequences — narrative-performance gaps are not uniformly consequential, but become substantially more consequential when independently corroborated, a finding with direct implications for how the growing availability of satellite and remote-sensing data (Testoni et al., 2023) may reshape the greenwashing landscape going forward.

Third, the brownwashing-relevant finding that Q3 (Quiet Achiever) firms' CAR (+0.18%), while positive, is notably smaller than Q1 (Aligned Leader) firms' CAR (+0.31%) despite comparable PI levels, extends Kim and Lyon's (2015) brownwashing concept with large-sample market-reaction evidence suggesting that under-communication of genuine

performance may carry a modest market cost — though this study's design cannot fully disentangle whether this reflects markets failing to recognize Q3 firms' genuine performance (an information/communication problem) or other unobserved differences between Q1 and Q3 firms correlated with their communication choices (Section 6.4). Fourth, the Narrative-Performance Divergence Lifecycle (Figure 3) provides a process-level theoretical account integrating this study's cross-sectional findings (Sections 4.2–4.4) with temporal dynamics (Table 9's stability findings, Table 10's intervention findings), extending greenwashing theory from a primarily static (point-in-time disclosure assessment) to a dynamic (lifecycle and feedback) framing.

Fifth, the disclosure-change pilot's visibility gradient finding (Table 10, H5) — that GGS reduction magnitude increases monotonically with feedback visibility/governance-integration level — extends AI governance and ethics-based auditing research (Mökander & Floridi, 2021) from internal organizational AI governance contexts to external, multi-stakeholder accountability contexts: the finding suggests that AI-generated accountability signals' behavioral effects may depend substantially on the governance architecture through

which those signals are channeled — a 'governance-mediated AI accountability' framing with implications extending beyond the greenwashing context to other domains in which AI systems generate accountability-relevant signals about organizational behavior. Blockchain-based frameworks offer a complementary approach to preserving audit-trail integrity and data provenance in multi-stakeholder governance pipelines, paralleling FACI and GIS audit-layer requirements in regulated AI contexts (Sammangi, Jagatha, & Liu, 2025c).

The ESG Assurance Maturity Roadmap

Figure 5 synthesizes this study's findings — particularly the disclosure-change pilot's visibility gradient (Table 10) and the sector analysis's identification of Scope 3 and supply-chain monitoring gaps (Table 6) — into a five-level ESG assurance maturity roadmap, characterizing organizational and ecosystem-level progression from Narrative-Only Reporting (Level 1, approximate GGS range +12 to +33, corresponding to this study's Q2/Greenwashing Risk firms, Table 4) to Publicly Benchmarked (Level 5, approximate GGS range –4 to +4, corresponding to Cohort C's post-intervention state, Table 10).

Figure 5. The ESG Assurance Maturity Roadmap: Five Levels from Narrative-Only Reporting to Publicly Benchmarked

Level 1 Narrative-Only Reporting	Level 2 Metric Tracking	Level 3 External Benchmarking	Level 4 Governance-Integrated	Level 5 Publicly Benchmarked
GGS ≈ +12 to +33 • Sustainability reporting driven primarily by narrative/marketing function • No systematic cross-check against physical/satellite data • Assurance, if present, scoped narrowly (Table 7, Theme 5)	GGS ≈ +5 to +12 • Scope 1/2 emissions tracked and disclosed • No Scope 3 / supply-chain systematic monitoring • No external divergence benchmarking	GGS ≈ +2 to +8 • Third-party ESG ratings consulted • Satellite/independent data reviewed periodically • GIS-style gap analysis conducted but not board-integrated (Table 10, Cohort A)	GGS ≈ 0 to +5 • GIS-style reporting integrated into board ESG oversight (Table 10, Cohort B) • Scope 3 / supply-chain signals systematically monitored • NAS-PI alignment a tracked governance metric	GGS ≈ –4 to +4 • GIS-style reports published or third-party benchmarked (Table 10, Cohort C) • Continuous feedback loop (Figure 1, Layer 4) operational • Q2 (Greenwashing Risk) firms show largest GGS reduction at this level

Note. GGS ranges represent approximate benchmarks derived from this study's quadrant classification (Table 4: Q2's mean GGS of 33.2 anchors Level 1's upper bound; Table 10: Cohort C's post-intervention GGS of 3.9, and the Q2-within-Cohort-C subgroup's post-intervention GGS of 19.8, jointly inform Level 5's range) and sector comparison (Table 6: Utilities' mean GGS of –2.4 informs the lower bound of Level 4–5 ranges). Organizations and ecosystem participants (regulators, ESG

data providers, investors) should validate these benchmarks against sector- and region-specific norms (Table 2, Table 6).

Practical Implications for Investors, Regulators, and Firms
 For investors and ESG data providers, this study's findings suggest that GGS — and particularly SRD-amplified GGS (Table 5, Models 2/4/5) — provides information content beyond what NAS or PI individually, or conventional ESG

ratings (which often emphasize disclosure comprehensiveness, correlating more closely with NAS than with PI), would capture. The CAR results (Table 5, Model 1–2) suggest that GGS-informed screening could identify firms facing elevated risk of negative abnormal returns around future disclosure events — though this study notes that if GGS-based screening becomes widely adopted, the documented CAR relationship could itself attenuate as the information becomes priced in more efficiently, an important caveat for any information-based investment strategy.

For regulators, the $\text{GGS} \times \text{SRD}$ interaction findings (H3) suggest that satellite and remote-sensing data — increasingly available per Testoni et al. (2023) — could play a structured role in regulatory disclosure review, potentially providing a scalable mechanism for identifying disclosures warranting closer regulatory scrutiny without requiring regulators to independently develop the full multimodal architecture this study's GIS represents.

The sector findings (Table 6) — particularly Energy and Materials sectors' elevated Scope 3 SRD — suggest that regulatory frameworks requiring Scope 3 disclosure (which the EU's CSRD, European Commission, 2023, and evolving SEC requirements, Securities and Exchange Commission, 2024, increasingly do) may benefit from sector-specific guidance regarding the independent verification mechanisms (e.g., supply-chain satellite monitoring) most relevant to each sector's Scope 3 profile.

For firms, the disclosure-change pilot results (Table 10) suggest a constructive framing: rather than viewing GIS-style gap analysis solely as an external accountability mechanism, firms could proactively adopt similar internal analysis (corresponding to Level 3–4 of the maturity roadmap, Figure 5) as a governance tool — Cohort B's board-level-briefing result (–6.6 point GGS reduction) suggests that internal governance integration, even without public disclosure of gap analysis results, produces meaningful disclosure and performance changes, potentially allowing firms to address narrative-performance gaps proactively before they become externally visible (and, per Table 5's $\text{GGS} \times \text{SRD}$ findings, potentially more consequential when externally corroborated).

Limitations and Future Research

Several limitations merit acknowledgment. First, the Performance Index (PI), while constructed from modalities

independent of firm narrative (Table 1), remains a composite measure whose weighting (Section 3.3) involves judgment calls; while validated against the expert panel (Table 9, achieving substantial though imperfect agreement), alternative PI constructions might yield somewhat different GGS values and quadrant classifications — future research exploring PI construction sensitivity, potentially through multiple alternative weighting schemes compared against the expert panel benchmark, would strengthen confidence in this study's specific quadrant assignments, though the core GGS-outcome relationships (Table 5) are likely robust to modest PI construction variation given their consistency across the robustness checks (Table 8).

Second, the Top-Decile GGS validation's 73.4% overlap with expert 'highest concern' classifications (Table 9) leaves 26.6% divergence that this study attributes to PI's environmental focus potentially missing governance or social dimensions relevant to experts' holistic 'concern' assessments; future research extending the GIS architecture (Figure 1) to incorporate social and governance dimension data — analogous to but distinct from the environmental modalities in Table 1 — represents a natural and important extension, particularly given that 'ESG' encompasses dimensions beyond the environmental focus of this study's PI construction.

Third, the disclosure-change pilot's (Table 10) cohort assignment was based on voluntary engagement rather than randomization (Section 5.4's note) — firms that voluntarily engaged with a GIS pilot program, particularly at the board-briefing (Cohort B) or public-benchmarking (Cohort C) levels, may differ systematically from non-engaging firms in ways that independently predict GGS reduction (e.g., firms already inclined toward improved ESG governance may be more likely to engage with such pilots, and might have reduced GGS even absent the pilot). While the Control cohort's modest GGS reduction (12.8 to 11.9, a –0.9 point baseline trend that treatment effects are computed net of) provides some reassurance that secular trends do not fully explain the treatment cohorts' larger reductions, a randomized design — assigning firms to feedback visibility levels independent of voluntary engagement — would provide substantially stronger causal evidence for H5, representing a priority direction for future research.

Fourth, this study's 2019–2026 observation period, while substantial, spans a period of rapidly evolving ESG disclosure

regulation (Table 6 notes reference the EU's CSRD and U.S. SEC climate disclosure rules, both implemented during or near this study's observation window); the GGS-outcome relationships documented in this study may themselves evolve as mandatory disclosure requirements reduce the discretionary narrative latitude that contributes to NAS variance — future research examining whether GGS's distribution, predictive validity, and the disclosure-change pilot's intervention effects persist, attenuate, or change character as disclosure regulation matures represents an important longitudinal extension. Fifth, this study's sample, while global (Table 2), exhibits regional variation in data availability (particularly satellite and supply-chain modalities, Table 1, which may have uneven global coverage) that could affect PI construction comparability across regions — future research explicitly examining and, where necessary, adjusting for regional data availability differences in PI construction would strengthen cross-regional comparisons (Table 2).

Conclusion

This study has developed and validated the Greenwashing Intelligence System (GIS), demonstrating that a multimodal AI architecture — combining transformer-based narrative analysis with satellite, controversy, financial, and supply-chain data independent of firm self-reporting — can produce a Greenwashing Gap Score (GGS) that significantly predicts market reactions, litigation risk, and media sentiment shifts, with these relationships substantially amplified when narrative-performance divergence is corroborated by externally-verifiable evidence. The four-quadrant classification framework's central finding — that Greenwashing Risk (Q2) firms experience meaningfully worse outcomes than Disengaged (Q4) firms despite comparable underlying performance levels — demonstrates that narrative-performance divergence itself, not merely poor environmental performance, carries distinct market and reputational consequences, a finding with direct implications for how ESG information is processed by markets and stakeholders.

The disclosure-change pilot's visibility gradient finding — that GGS feedback produces progressively larger reductions as feedback visibility and governance integration increase, with the largest effects concentrated among the highest-priority (Greenwashing Risk) firms — provides a constructive complement to the GIS's detection capabilities, suggesting that AI-based greenwashing detection systems, when integrated into appropriate governance and disclosure architectures, may

contribute not only to identifying narrative-performance gaps but to reducing them. The broader potential of deep learning and AI-based forecasting systems for classification and prediction tasks in high-stakes organizational contexts continues to be demonstrated across domains, underscoring the transferability of the multimodal, multi-objective evaluation approach central to this study (Sharma, Singh, Sammangi, Sharma, Pandey, Srivastava, Agarwal, & Singh, 2025a).

As ESG disclosure continues to expand in volume and regulatory significance, and as the independent verification data sources (satellite, remote-sensing, supply-chain monitoring) underlying this study's PI construction continue to improve in coverage and resolution (Testoni et al., 2023), the GIS framework, NAS/PI/GGS measurement architecture, and ESG assurance maturity roadmap developed in this study provide investors, regulators, and firms with both a diagnostic capability and a constructive pathway toward sustainability disclosure that more closely reflects — and, per the disclosure-change pilot, may help drive — genuine environmental performance.

REFERENCES

1. Bingler, J. A., Kraus, M., Leippold, M., & Webersinke, N. (2024). How cheap talk in climate disclosures relates to climate initiatives, corporate emissions, and reputation risk. *Journal of Banking & Finance*, 158, 107057.
2. Bowen, F., & Aragon-Correa, J. A. (2014). Greenwashing in corporate environmentalism research and practice: The importance of what we say and do. *Organization & Environment*, 27(2), 107–112.
3. Christensen, H. B., Hail, L., & Leuz, C. (2021). Mandatory CSR and sustainability reporting: Economic analysis and literature review. *Review of Accounting Studies*, 26(3), 1176–1248.
4. Delmas, M. A., & Burbano, V. C. (2011). The drivers of greenwashing. *California Management Review*, 54(1), 64–87.
5. Dragomir, V. D., & Dumitru, M. (2023). Sustainability reporting regulation: Reflections on the EU Corporate Sustainability Reporting Directive. *Sustainability*, 15(15), 11770.
6. European Commission. (2023). Corporate Sustainability Reporting Directive (CSRD): Implementation guidance for European Sustainability Reporting Standards (ESRS). Official Journal of the European Union.

7. Flammer, C. (2021). Corporate green bonds. *Journal of Financial Economics*, 142(2), 499–516.
8. Garcia, D., & Tsafack, G. (2024). Climate news and stock returns: International evidence. *Journal of Financial and Quantitative Analysis*, 59(2), 521–561.
9. Grewal, J., Hauptmann, C., & Serafeim, G. (2021). Material sustainability information and stock price informativeness. *Journal of Business Ethics*, 171(3), 513–544.
10. Hartzmark, S. M., & Sussman, A. B. (2019). Do investors value sustainability? A natural experiment examining ranking and fund flows. *The Journal of Finance*, 74(6), 2789–2837.
11. Hsueh, L. (2024). Voluntary corporate climate disclosure and the role of analyst pressure. *Strategic Management Journal*, 45(3), 612–639.
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. Jagatha, A., Sammangi, H., & Maddireddy, H. G. (2025). Decentralized multi-hop federated reinforcement learning for energy-efficient and secure routing in LoRaWAN-based smart city infrastructure. *Engineering: Open Access*, 3(6), 1–8.
14. Kölbel, J. F., Leippold, M., Rillaerts, J., & Wang, Q. (2024). Ask BERT: How regulatory disclosure of transition and physical climate risks affects the CDS term structure. *Journal of Financial Econometrics*, 22(1), 30–69.
15. Lashitew, A. A. (2024). Corporate uptake of the Sustainable Development Goals: Mere greenwashing or an advent of institutional change? *Journal of International Business Policy*, 7(1), 78–107.
16. Lyon, T. P., & Maxwell, J. W. (2011). Greenwash: Corporate environmental disclosure under threat of audit. *Journal of Economics & Management Strategy*, 20(1), 3–41.
17. Lyon, T. P., & Montgomery, A. W. (2015). The means and end of greenwash. *Organization & Environment*, 28(2), 223–249.
18. Marquis, C., Toffel, M. W., & Zhou, Y. (2016). Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, 27(2), 483–504.
19. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
20. Mökander, J., & Floridi, L. (2021). Ethics-based auditing to develop trustworthy AI. *Minds and Machines*, 31(2), 323–327.
21. Reber, B., Gold, A., & Gold, S. (2022). ESG disclosure and idiosyncratic risk in initial public offerings. *Journal of Business Ethics*, 179(3), 867–886.
22. Sammangi, H., Ambati, L. S., Liu, J., & Jagatha, A. (2025). AI-driven decentralized IoT for secure and scalable healthcare. *AMCIS 2025 Proceedings*. https://aisel.aisnet.org/amcis2025/health_it/sig_health/3
23. Sammangi, H., Jagatha, A., & Liu, J. (2025b). Harnessing generative AI and large language models for revolutionizing cybersecurity in the Internet of Things: Ethical and privacy implications. *Engineering: Open Access*, 3(6), 1–12.
24. Sammangi, H., Jagatha, A., & Liu, J. (2025c). Integrating blockchain technology into telemedicine: A framework for enhancing data privacy and security. *Engineering: Open Access*, 3(6), 1–7.
25. Sautner, Z., van Lent, L., Vilkov, G., & Zhang, R. (2023). Firm-level climate change exposure. *The Journal of Finance*, 78(3), 1449–1498.
26. Securities and Exchange Commission. (2024). Final rule: The enhancement and standardization of climate-related disclosures for investors. 17 CFR Parts 210, 229, 232, 239, and 249.
27. Sharma, G., Singh, J., Sammangi, H., Sharma, M., Pandey, R., Srivastava, S., Agarwal, G., & Singh, I. (2025a). A comprehensive assessment of developing a forecasting model for kidney stone formation using deep learning approaches. In H. Sharma, A. Chakravorty, S. Hussain, & R. Kumari (Eds.), *Artificial Intelligence: Theory and Applications* (Vol. 5588, pp. 121–132). Springer Nature Singapore. https://doi.org/10.1007/978-981-96-1918-4_9
28. Seetharaman, A., Niranjana, I., Saravanan, A. S., & Balaji, D. (2015). A study of the issues affecting ESG investments. *Journal of Corporate Accounting & Finance*, 26(6), 19–25.
29. Tang, D. Y., & Zhang, Y. (2020). Do shareholders benefit from green bonds? *Journal of Corporate Finance*, 61, 101427.
30. Testoni, M., Ramnath, S., & Vyas, S. (2023). Satellite-based monitoring of corporate environmental performance: A methodological review. *Environmental Research Letters*, 18(4), 043001.

31. Tornbohm, C., & Doherty, S. (2023). Generative AI for ESG: Risks and opportunities in sustainability reporting. Gartner Research.
32. Walker, K., & Wan, F. (2012). The harm of symbolic actions and green-washing: Corporate actions and communications on environmental performance and their financial implications. *Journal of Business Ethics*, 109(2), 227–242.
33. Wang, Q., Zheng, W., & Zheng, Y. (2024). Large language models and ESG disclosure analysis: A systematic review. *Journal of Accounting Literature*, 46(1), 1–34.
34. Yu, E. P., Van Luu, B., & Chen, C. H. (2020). Greenwashing in environmental, social and governance disclosures. *Research in International Business and Finance*, 52, 101192.
35. Zhang, D. (2022). Are firms motivated to greenwash by financial constraints? Evidence from global firms' data. *Journal of International Financial Management & Accounting*, 33(1), 80–97.