



School of Economics and Management

Finance

LLM-Based Assessment of ESG Disclosure Quality in China: Greenwashing Determinants and Regulatory Effects

The ESG Disclosure Problem

ESG Disclosure creates transparency but also opportunities for **strategic communication**

Strategic ESG Rhetoric

- *Aspirational claims*
- *Selective narratives*
- *ESG branding*
- *Volume of disclosure*



Information
Asymmetry /
Greenwashing

Substantive ESG Practices

- *Verifiable metrics*
- *Operational changes*
- *Quantitative evidence*
- *Auditable outcomes*

Greenwashing is the **strategic amplification** of ESG rhetoric without a commensurate improvement in **substantive environmental or social practices**

Greenwashing as a Rhetoric-Substance Gap

Greenwashing is not simply false disclosure; it is a mismatch between communication and verifiable reality


Passive vs. Active Mechanisms

Narrative Focus vs Structural Focus	Omission - selective disclosure - withholding adverse ESG information	Commission - impression management - decoupling
	Symbolic Layer - positive framing - policies without enforcement	Substantive Layer - verifiable metrics - operational change

“Greenwashing = Rhetorical Appearance – Verifiable Substance”

China as a Regulatory Laboratory

ESG disclosure is shaped by the overlapping market and state-led legitimacy pressures




Theoretical Framework: Legitimacy Theory & The Legitimacy Budget

Why Do Firms Greenwash?

Disclosure is a resource-constrained mode of legitimacy production



1. *Legitimacy Theory in China*: Firms face dense, overlapping pressures: state policy, exchange templates, and investor expectations. Disclosure signals alignment
2. *The Legitimacy Budget Constraint*: Because organizational resources are finite, firms optimize their legitimacy production by decoupling external rhetoric from substantive alignment



Theoretical Mechanism: Rhetoric Migration

What happens when regulation forces standardized disclosure? Mandatory disclosure does not mechanically eliminate greenwashing; rather, it alters the internal composition of symbolic compliance

Post-2022 Mandatory Standardization



Rhetoric Migration:
When regulation tightens in highly verifiable domains (*Environmental*), symbolic compliance does not disappear — it migrates to less verifiable domains (*Social*)


Research Questions & Hypotheses

RQ1: Does **systematic greenwashing** (*positive mean GWI*) exist among Chinese blue-chip firms in the 2018–2024 panel?



H1: Chinese listed firms are expected to exhibit a **systematic rhetoric–substance gap** in ESG disclosure (*supported by Delmas & Burbano, 2011; Lyon & Montgomery, 2015*)

RQ2: Do **extractive-sector firms** and **state-owned enterprises** exhibit higher GWI than non-extractive and private firms respectively?



H2: The rhetoric–substance gap should be more pronounced in firms facing **heavier** and more **overlapping legitimacy pressures**, particularly firms with high environmental salience and firms with strong political embeddedness (*supported by Suchman, 1995; Patten, 2002*)

RQ3: Did the **post-2022** phase of intensified ESG disclosure regulation and exchange-led standardization change the within-firm rhetoric-substance gap, and **through which pillar**?



H3: Under conditions of mandatory ESG disclosure and **uneven cross-pillar standardization**, symbolic compliance should **migrate toward less verifiable domains** rather than disappear (*supported by Meyer & Rowan, 1977; Bromley & Powell, 2012*)



Identifying the Research Gap

Despite rapid advancements, current LLM-based sustainable finance research exhibits significant structural gaps

GAP 1: WESTERN & ENGLISH BIAS

Most LLM studies focus exclusively on US-listed firms (e.g., NASDAQ 100) and English prospectuses. The largest emerging market (China) and native-language processing remain critically underexplored


GAP 2: MONOLITHIC ESG MEASUREMENT

Previous computational studies rarely decompose ESG text into distinct Rhetorical vs. Substantive dimensions. They measure disclosure quantity (how much a firm talks) rather than disclosure quality (the rhetoric-substance gap)

GAP 3: UNEXAMINED REGULATORY EFFECTS

The causal impact of the 2022 SSE/SZSE mandatory disclosure shock on the internal, pillar-level composition of greenwashing remains entirely absent from the empirical literature

The Evolution of Textual Analysis in Sustainable Finance



Why LLMs for ESG

Semantic Interpretation: Distinguishes confident, aspirational tone from evidential, factual substance

Long-Document Processing: Ingests 100+ page unstructured annual/ESG reports via advanced chunking and embedding

Multilingual Capability: Processes native Chinese texts directly, avoiding systemic translation distortions

Structured Evaluation: Follows explicit, prompt-based rubrics (e.g., separating verifiable Scope 1 emissions figures from vague net-zero promises)



Data Base and Sample Construction

Key Sectors	No. of firms	Extractive Status	Typical SOE Share
Oil, gas, coal, energy	17	Yes	High (>50%)
Metals & mining	14	Yes	Medium–High
Financials (banks, insurers)	23	No	High
Utilities	9	No	High
Industrials	18	No	Medium
Consumer & retail	13	No	Low
Technology & telecoms	12	No	Mixed
Other	26	No	Low

TOTAL

132

Step 1:
Universe
Selection

Step 2:
Inclusion
Criteria

Step 3: Final
Panel

132 blue-chip listed firms (*Intersection of CSI 300 & SSE 50 indices*)

Data Sources: CNINFO, Wind, CSMAR, SASAC, IR websites

- Continuous listing across the **2018–2024 period**
- Availability of **machine-readable PDF** standalone ESG/CSR or Annual reports
- Minimum threshold of **2,000 non-financial words** per report


924 balanced firm-year observations

The LLM Processing Pipeline Architecture

The pipeline **transforms** *unstructured ESG disclosures* into *structured firm-level greenwashing measures*

Reports are extracted, cleaned, segmented, vectorized, and filtered through **ESG-specific semantic retrieval** before being evaluated by *GPT-4o*

The architecture combines **Retrieval-Augmented Generation (RAG)** with **dual LLM scoring** to distinguish *rhetorical* ESG communication from verifiable *substantive* disclosure



Operationalizing the Greenwashing Index (GWI)

RES anchors:

0 → No ESG rhetoric

50 → General commitments


100 → Comprehensive multi-year strategy

SES anchors:

0 → Purely qualitative disclosure

50 → Partial numerical disclosure

100 → Fully auditable metrics



The **Greenwashing Index (GWI)**

operationalizes ESG disclosure quality as *the gap* between rhetorical appearance and substantive reality

Positive values indicate that ESG narratives *exceed* verifiable evidence, while **lower or negative values** reflect *closer alignment* between communication and measurable ESG performance

The framework *enables* pillar-level and firm-level assessment of **symbolic compliance**

Econometric Framework and Hypothesis Testing

1. Existence & Time Trends (H1)

H1a: One-sample t-test and **Wilcoxon signed-rank test** are applied to the full pooled panel under $H_0: \mu(\text{GWI}) = 0$, separately for each pillar

H1b:

$$M1: \text{GWI}_{it} = \alpha + \beta_1 \cdot \text{Year_centred}_t + \varepsilon_{it}$$

$$M2: \text{GWI}_{it} = \alpha + \beta_2 \cdot \text{Post2022}_t + \varepsilon_{it}$$

2. Structural Determinants (H2)

$$\text{GWI}_{it} = \gamma_t + \beta_1 \cdot \text{Extractive}_i + \beta_2 \cdot \text{SOE}_i + \beta_3 \cdot (\text{Extractive}_i \times \text{SOE}_i) + \varepsilon_{it}$$

Panel OLS with year fixed effects (γ_t)

Firm fixed effects omitted to isolate time invariant traits

3. Regulatory Effect (H3)

$$\text{GWI}_{it} = \alpha_i + \beta \cdot \text{Post2022}_t + \varepsilon_{it}$$

Two-way fixed effects specification

Firm fixed effects (α_i) absorb time-invariant heterogeneity




Finding 1: Systematic Greenwashing is Robust

The Chinese blue-chip panel exhibits a highly significant rhetoric-substance gap, driven by environmental claims

Mean Composite GWI: 8.39

A structural gap exists across the panel



Pooled Panel Summary (2018–2024, N=924)

Metric	Mean	Std. Dev.
Rhetoric (RES)	62.55	11.08
Substance (SES)	54.16	16.43
Composite GWI	8.39	10.56
Environmental GWI	19.04	16.66
Social GWI	-1.01	18.24
Governance GWI	7.14	12.05

Key Insight 1: The Composite GWI is highly positive (8.39), proving a baseline rhetoric-substance gap

Key Insight 2: The gap is overwhelmingly concentrated in the Environmental pillar (19.04), where reputational stakes are highest

Key Insight 3: The Social pillar appears balanced (-1.01) in aggregate, but this pooled figure masks severe temporal volatility


Finding 2: Greenwashing is Stronger in Extractive and State-Controlled Firms

Key Insight 1: Environmental salience increases compensatory ESG rhetoric

Key Insight 2: Political embeddedness intensifies symbolic alignment with national ESG priorities

Key Insight 3: Extractive × SOE interaction is not statistically significant. Legitimacy pressures appear additive rather than synergistic

STRUCTURAL DETERMINANTS OF COMPOSITE GWI




Extractive: +5.70 | SOE: +2.40

Both legitimacy pressures significantly increase the rhetoric–substance gap

Firms facing stronger environmental and political legitimacy pressures exhibit systematically higher rhetoric–substance gaps

Finding 3: Greenwashing Intensifies After 2022

ANNUAL EVOLUTION OF PILLAR-LEVEL GWI, 2018–2024




1. GWI_E is essentially flat across the panel, ranging from 17.42 to 20.53

2. GWI_S evolves dramatically: from clearly negative values in 2018–2020 to clearly positive values from 2021 onwards

Takeaway: Environmental greenwashing remains structurally persistent, while Social greenwashing reverses sharply after 2021

Finding 4: Post-2022 Regulation Shifts Increases Greenwashing

In April 2022, mandatory structured ESG reporting was enforced. It **did not reduce** the rhetoric-substance gap



The Mandate


Required standardized reporting on E, S, and G indicators for major index constituents

The Reality

Assumption: Mandatory disclosure uniformly reduces greenwashing
Reality: Mandatory disclosure expanded the scope for unsupported rhetoric

Finding 5: Evidence of Rhetoric Migration

PRE/POST PILLAR MIGRATION



Forced to report across all pillars, firms shifted symbolic compliance toward the Social pillar

The Mechanism Proven

- New templates demanded Social narratives (diversity, training, community)
- Firms generated massive Social rhetoric to comply
- However, verifiable Social metrics (LTIFR, training spend) could not keep pace, creating a massive new gap



Robustness, Validity, and Limitations

Metric	Original Prompt	Alt1 (Vocab Sub)	Alt2 (Rubric Swap)
Mean GWI	7.67	7.67	7.50
H1a t-statistic	> 6.7	> 6.8	> 6.9
Result	Stable	Stable	Stable

Prompt Sensitivity

While firm-level ranks shift slightly under different prompts, aggregate distributional moments are highly stable. GWI is validated as a macro-level policy instrument

Anchor Clustering

LLM outputs naturally cluster at prompted 0, 25, 50, 75, 100 markers, limiting minor variance capture

Scope Limitation

Analysis strictly captures textual greenwashing but inherently misses visual greenwashing (e.g., infographics, selective photography)

Validation of Hypotheses Against Prior Literature

HYPOTHESIS	RESULT	LINK TO PRIOR LITERATURE
H1: Chinese listed firms are expected to exhibit a systematic rhetoric–substance gap in ESG disclosure	 Supported	Consistent with Delmas & Burbano (2011) and Cho et al. (2015) , who argue that corporate ESG disclosure often contains symbolic rather than substantive commitments
H2: The rhetoric–substance gap is stronger in firms facing heavier and more overlapping legitimacy pressures , (extractive firms and SOEs)	 Supported	Supports legitimacy-based explanations of greenwashing proposed by Suchman (1995) and Marquis & Toffel (2012)
H3: Under conditions of mandatory ESG disclosure and uneven cross-pillar standardization , symbolic compliance should migrate toward less verifiable domains rather than disappear	 Supported	Extends disclosure-regulation literature (Christensen et al., 2021) by identifying a new mechanism: <i>Rhetoric Migration</i>

Policy Implications of Rhetoric Migration

Disclosure templates that mandate narrative without anchoring it to metrics inadvertently expand greenwashing

For Policymakers

- Uneven standardization creates hiding spots
- Every qualitative pillar requirement must be explicitly tethered to standardized, auditable quantitative indicators

Warning: Mandating broad text boxes without metrics invites rhetoric migration

For Investors

- Single-number, composite ESG scores obscure risk
- Analysts must decompose ESG ratings to identify rhetoric migration and cross-pillar substitution patterns

Warning: A rising ESG score may just indicate a shift in the corporate legitimacy budget

Conclusion: Greenwashing is a dynamic organizational response to resource constraints, not a static state



Conclusion

Greenwashing is not a static lie, but a dynamic reallocation of corporate resources
Partial standardization causes symbolic compliance to migrate, not disappear

Cross-Market Validation

Extending the **RAG-LLM framework** to **other emerging markets** with **mandatory regimes** (India, Brazil)

Multimodal AI Integration

Incorporating **computer vision models** to **detect visual greenwashing** in charts and photography

Market Outcomes

Linking the GW to **future stock mispricing** or **regulatory fines**

Textual Baseline
(2024)







- Ali, H., Zafar, M. B., & Aysan, A. F. (2025). Generative AI in finance: Replicability, methodological contingencies, and future research directions. *Finance Research Letters*, 86, 108797. <https://doi.org/10.1016/j.frl.2025.108797>
- Bao, R., & Wei, L. (2024). The impact of companies disclosing ESG reports in multiple languages on the enthusiasm of foreign investors for holding shares. *PLOS ONE*, 19(3), e0299643.
- Berg, F., Kölbel, J. F., & Rigobon, R. (2022). Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, 26(6), 1315–1344. <https://doi.org/10.1093/rof/rfac033>
- Bingler, J. A., Kraus, M., Leippold, M., & Webersinke, N. (2022). Cheap talk and cherry-picking: What ClimateBERT has to say on corporate climate risk disclosures. *Finance Research Letters*, 47, 102776. <https://doi.org/10.1016/j.frl.2022.102776>
- 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*, 164, 107191.
- Boiral, O. (2013). Sustainability reports as simulacra? A counter-account of A and A+ GRI reports. *Accounting, Auditing & Accountability Journal*, 26(7), 1036–1071.
- Bromley, P., & Powell, W. W. (2012). From Smoke and Mirrors to Walking the Talk: Decoupling in the Contemporary World. *The Academy of Management Annals*, 6(1), 483–530. <https://doi.org/10.1080/19416520>
- Cai, W., Lee, E., Xu, A. L., & Zeng, C. (2019). Does corporate social responsibility disclosure reduce the information disadvantage of foreign investors? *Journal of International Accounting, Auditing and Taxation*, 34, 12–29.
- Cho, C. H., Michelon, G., & Patten, D. M. (2012). Impression management in sustainability reports: An empirical investigation of the use of graphs. *Accounting and the Public Interest*, 12(1), 16–37.
- Cho, C. H., & Patten, D. M. (2007). The role of environmental disclosures as tools of legitimacy: A research note. *Accounting, Organizations and Society*, 32(7–8), 639–647.
- Christensen, D. M., Serafeim, G., & Sikochi, A. (2022). Why is corporate virtue in the eye of the beholder? The case of ESG ratings. *The Accounting Review*, 97(1), 147–175.
- Clarkson, P. M., Li, Y., Richardson, G. D., & Vasvari, F. P. (2008). Revisiting the relation between environmental performance and environmental disclosure. *Accounting, Organizations and Society*, 33(4–5), 303–327.
- Dai, N. T., Du, F., Young, S. M., & Tang, G. (2018). Seeking legitimacy through CSR reporting: Evidence from China. *Journal of Management Accounting Research*, 30(1), 1–29.
- Delmas, M. A., & Burbano, V. C. (2011). The drivers of greenwashing. *California Management Review*, 54(1), 64–87.
- Dhaliwal, D. S., Li, O. Z., Tsang, A., & Yang, Y. G. (2011). Voluntary nonfinancial disclosure and the cost of equity capital. *The Accounting Review*, 86(1), 59–100.
- Dorfleitner, G., & Utz, S. (2024). Green, green, it's green they say: A conceptual framework for measuring greenwashing on firm level. *Review of Managerial Science*, 18(12), 3463–3486.
- Du, X., Jian, W., Zeng, Q., Du, Y., & Chang, Y. (2017). Do lenders applaud corporate environmental performance? Evidence from Chinese private-owned firms. *Journal of Business Ethics*, 143(1), 179–207.
- Eccles, R. G., Ioannou, I., & Serafeim, G. (2014). The impact of corporate sustainability on organizational processes and performance. *Management Science*, 60(11), 2835–2857.
- Gentzkow, M., Kelly, B., & Taddy, M. (2019). Text as data. *Journal of Economic Literature*, 57(3), 535–574.
- Hahn, R., & Lülfs, R. (2014). Legitimizing negative aspects in GRI-oriented sustainability reporting. *Journal of Business Ethics*, 123(3), 401–420.
- Hawn, O., & Ioannou, I. (2016). Mind the gap: The interplay between external and internal CSR actions. *Strategic Management Journal*, 37(13), 2569–2588.
- He, X., Si, Z., & Yu, X. (2025). The influential factors of CSR reporting in China. *Accounting, Finance & Governance Review*, 34.
- Hrasky, S. (2012). Carbon footprints and legitimization strategies: Symbolism or action? *Accounting, Auditing & Accountability Journal*, 25(1), 174–198.
- Kirtac, K., & Germano, G. (2024). Sentiment trading with large language models. *Finance Research Letters*, 62, 105227.
- Kostka, G., & Nahm, J. (2017). Central-Local Relations: Recentralization and Environmental Governance in China. *The China Quarterly*, 231, 567–582.
- Lan, Y., Qin, M., & Hu, F. (2024). Emerging trends in AI-driven financial NLP: A multilingual perspective. *Finance Research Letters*, 65, 105508.
- Li, W., & Qi, G. (2021). Bureaucratic agency, corporate political connections and CSR. *Management Decision*, 59(8), 1875–1895.
- Lin, X., Yu, J., Wang, S., & Zhang, N. (2023). ESG greenwashing and equity mispricing: Evidence from China. *Finance Research Letters*, 58, 104470.
- Liu, X., & Anbumozhi, V. (2009). Determinant factors of corporate environmental information disclosure. *Journal of Cleaner Production*, 17(6), 593–600.
- Loughran, T., & McDonald, B. (2011). When is a liability not a liability? *Journal of Finance*, 66(1), 35–65.
- Loughran, T., & McDonald, B. (2016). Textual analysis in accounting and finance: A survey. *Journal of Accounting Research*, 54(4), 1187–1230.
- Lublóy, Á., Keresztúri, J. L., & Berlinger, E. (2025). Quantifying firm-level greenwashing. *Journal of Environmental Management*, 373, 123399.
- Lyon, T. P., & Maxwell, J. W. (2011). Greenwash: Corporate environmental disclosure under threat of audit. *Journal of Economics & Management Strategy*, 20(1), 3–41.
- Lyon, T. P., & Montgomery, A. W. (2015). The means and end of greenwash. *Organization & Environment*, 28(2), 223–249.
- Marquis, C., & Qian, C. (2014). Corporate social responsibility reporting in China: Symbol or substance? *Organization Science*, 25(1), 127–148.
- Marquis, C., Toffel, M. W., & Zhou, Y. (2016). Scrutiny, norms, and selective disclosure. *Organization Science*, 27(2), 483–504.
- Mathebula, T., Adebayo, J., & Thompson, K. (2024). Multilingual large language models and financial NLP. *Finance Research Letters*, 66, 105644.
- Meyer, J., & Rowan, B. (1977). Institutionalized Organizations: Formal Structure as Myth and Ceremony. *American Journal of Sociology*, 83, 340–363.
- Michelon, G., Pilonato, S., & Ricceri, F. (2015). CSR reporting practices and disclosure quality. *Critical Perspectives on Accounting*, 33, 59–78.
- Parsa, S., Dai, N., Belal, A. et al. (2021) Corporate social responsibility reporting in China : political, social and corporate influences. *Accounting and Business Research*, 51(1), 36–64.
- Shen, H., Loh, L., Zhou, Y., Zhang, X., & Ormazabal, G. (2023). ESG in China: A review of practice and research. *China Journal of Accounting Research*, 16(4), 100325.
- Song, S., Jun, A., Luo, T., & Ma, S. (2024). Political legitimacy and CSR reporting. *Global Finance Journal*, 60, 100942.
- Suchman, M. C. (1995). Managing Legitimacy: Strategic and Institutional Approaches. *The Academy of Management Review*, 20(3), 571–610. <https://doi.org/10.2307/258788>
- SynTao. (2022). A report on the state of CSR reporting in China 2022. SynTao Green Finance.
- Wang, Q. (2025). Generative AI-assisted evaluation of ESG practices and information delays. *Finance Research Letters*, 74, 106757. <https://doi.org/10.1016/j.frl.2025.106757>
- Wang, Y., Hu, F., & Wang, Y. (2024). Analyst coverage and greenwashing: Evidence from Chinese A-share firms. *International Review of Economics & Finance*.
- Walker, K., & Wan, F. (2012). The harm of symbolic actions and green-washing. *Journal of Business Ethics*, 109(2), 227–242.
- Webersinke, N., Kraus, M., Bingler, J. A., & Leippold, M. (2022). ClimateBERT: A pretrained language model for climate-related text. *arXiv preprint*.
- Wood, K., Pyun, C., & Pham, H. (2025). Beyond green labels: Assessing mutual funds' ESG commitments. *Finance Research Letters*, 74, 106713.
- Yuan, X., Li, C., Li, Z., & Wang, J. (2022). ESG disclosure and corporate financial irregularities. *Journal of Cleaner Production*, 332, 130011.



Addressing Reviewer and Supervisor Comments

- 1. Alternative explanation of the post-2022 increase in Social GWI (the observed effect may reflect post-pandemic reporting trends rather than rhetoric migration):* Additional pillar-level analysis shows that the increase is concentrated almost entirely in the Social pillar, while Environmental and Governance dimensions remain relatively stable. This asymmetric response is more consistent with a reallocation of symbolic disclosure effort than with a general reporting trend affecting all ESG dimensions equally.
- 2. Potential Anglophone bias of GPT-4o and construct validity concerns:* The study relies on multilingual Chinese ESG reports and applies identical scoring criteria across all firms and years. Robustness checks using alternative prompt formulations produced highly stable aggregate GWI values (7.50–7.67), suggesting that the main conclusions are not driven by prompt wording or language-specific artifacts.
- 3. Difficulty of measuring Social-pillar performance may inflate Social greenwashing estimates:* This limitation is acknowledged. However, the same RES–SES scoring logic was applied across all ESG pillars. The fact that only the Social dimension exhibits a structural post-2022 increase supports the interpretation that the observed shift reflects disclosure behavior rather than a purely measurement-related issue.
- 4. Absence of a formal control group for the regulatory effect:* The dissertation does not claim strict causal identification. Instead, it uses a balanced panel and within-firm pillar decomposition to provide evidence consistent with regulatory-induced disclosure reallocation. Future research may apply difference-in-differences designs when suitable control groups become available.
- 5. Limited covariates and lack of external validation:* These limitations are explicitly acknowledged in the dissertation. Future research should incorporate governance characteristics, political connections, ESG controversies, regulatory penalties, and market-based outcomes to further validate the Greenwashing Index.
- 6. Prompt sensitivity and ranking instability (Spearman 0.50–0.81):* The objective of the GWI is not firm-level ranking but group-level and policy-level inference. Despite moderate rank variation, all prompt versions reproduce the same aggregate conclusions, supporting the use of GWI as a macro-level measurement instrument.



Prompt Templates

A.1 Rhetorical ESG Score (RES) — System prompt:

You are an expert ESG analyst specializing in corporate sustainability reporting. Your task is to evaluate the rhetorical dimension of ESG disclosure in the provided corporate document excerpts. Assess the strength, ambition, and frequency of ESG commitments, strategic pledges, forward-looking targets, and aspirational language. Do NOT consider the presence or absence of numerical data. Score each of the three ESG pillars (Environmental, Social, Governance) on a scale from 0 to 100 using the following anchors: 0 — no ESG rhetoric; 25 — minimal mention without commitments; 50 — moderate ESG language with some commitments; 75 — strong narrative with clear targets; 100 — comprehensive, ambitious multi-year strategy with explicit framework references (GRI, TCFD, UN SDGs). Return your response exclusively as a JSON object with fields: e_res (0–100), s_res (0–100), g_res (0–100), rationale (2–3 sentences). Do not include any text outside the JSON object.

A.2 Substantive ESG Score (SES) — System prompt:

You are an expert ESG analyst specializing in corporate sustainability reporting. Your task is to evaluate the substantive dimension of ESG disclosure by identifying verifiable quantitative data in the provided document excerpts. Score each pillar on the presence of specific numerical metrics only. For Environmental: absolute Scope 1 GHG emissions (required), Scope 2 (preferred), specific energy consumption figures, water withdrawal volumes. For Social: numerical injury rates (LTIFR/TRIR), percentage of women in management, training hours or expenditure per employee. For Governance: percentage of independent directors, explicit inclusion of ESG metrics in executive compensation (binary with description). Anchors: 0 — no quantitative data; 25 — one or two metrics; 50 — half of the checklist; 75 — most items with trend data; 100 — full checklist with independently verifiable figures and year-on-year comparisons. Return as JSON: e_ses, s_ses, g_ses (0–100), rationale. No text outside the JSON.



LLM scoring pipeline — technical specification

The scoring pipeline operates as follows: (1) PDF ingestion via PyMuPDF with character-level text extraction and GB2312/GBK to UTF-8 transcoding; (2) segmentation into 1,500-token chunks with 150-token overlap; (3) embedding computation using OpenAI text-embedding-3-large (3,072 dimensions); (4) FAISS indexing with inner-product similarity; (5) per-pillar ESG-relevance RAG query retrieving top-k = 10 chunks; (6) dual LLM scoring with Prompt A (RES) and Prompt B (SES) at temperature = 0; (7) JSON parsing and schema validation; (8) GWI computation as the simple difference $RES - SES$ at the pillar level, with composite GWI as the equal-weighted mean across pillars. All API calls are logged with timestamps, model version strings, and token counts for reproducibility documentation.



H1a: Existence of systematic greenwashing

Tests of H1a (Mean GWI > 0)

Variable	N	Mean	Median	t-statistic	p-value (one-sided > 0)
GWI_E raw	924	19.04	25	34.74	< 0.001
GWI_S raw	924	-1.01	0	-1.68	0.953
GWI_G raw	924	7.14	0	18.02	< 0.001
GWI composite raw	924	8.39	8.33	24.15	< 0.001

✓ Systematic greenwashing confirmed at the composite level

H2: Determinants of greenwashing

Panel OLS with year fixed effects

Model	Variable	Coefficient	Std. error	p-value
M1: extractive only	Extractive	+3.607	1.426	0.011
	Intercept	6.011	0.993	< 0.001
M2: extractive + SOE	Extractive	+2.951	1.473	0.045
	SOE	+2.038	0.985	0.038
	Intercept	5.081	1.059	< 0.001
M3: with interaction	Extractive	+5.695	2.380	0.017
	SOE	+2.403	1.040	0.021
	Extractive × SOE	-3.678	2.945	0.212
	Intercept	4.915	1.077	< 0.001

✓ Extractive and SOE effects supported
✗ Interaction effect not supported

H1b: Time dynamics of greenwashing

Time-trend regressions for composite GWI

Specification	Variable	Coefficient	Std. error	p-value
M1: linear year trend	Year (centred)	+0.639	0.201	0.001
	Intercept	6.476	0.834	< 0.001
M2: Post2022 dummy	Post2022	+2.547	0.768	0.001
	Intercept	7.301	0.620	< 0.001

✓ Greenwashing increased over time

H3: Post-2022 regulatory effect

Within-firm FE estimates

Specification	Variable	Coefficient	Std. error	p-value
Composite GWI, firm FE	Post2022	+2.547	0.829	0.002
E pillar GWI, firm FE	Post2022	-0.909	1.290	0.481
S pillar GWI, firm FE	Post2022	+7.617	1.284	< 0.001
G pillar GWI, firm FE	Post2022	+0.934	0.916	0.308

✓ Increase concentrated in Social pillar



Prompt-sensitivity analysis

Comparison	N	Spearman ρ	Pearson r	Mean diff	Within 25 pts
e_res original vs Alt1	100	0.527	0.628	6.25	100%
s_res original vs Alt1	100	0.635	0.665	9.25	98%
g_res original vs Alt1	100	0.681	0.719	3.75	100%
e_ses original vs Alt1	100	0.745	0.756	10.00	98%
s_ses original vs Alt1	100	0.738	0.739	9.25	97%
g_ses original vs Alt1	100	0.619	0.602	10.00	97%
RES mean original vs Alt1	100	0.682	0.757	5.75	95%
SES mean original vs Alt1	100	0.809	0.795	7.58	97%
GWl composite original vs Alt1	100	0.643	0.603	6.17	98%
e_res original vs Alt2	100	0.714	0.762	5.00	100%
s_res original vs Alt2	100	0.333	0.321	19.29	89%
g_res original vs Alt2	100	0.614	0.711	5.79	100%
e_ses original vs Alt2	100	0.726	0.732	11.75	94%
s_ses original vs Alt2	100	0.658	0.671	12.25	96%
g_ses original vs Alt2	100	0.556	0.582	11.50	95%
RES mean original vs Alt2	100	0.669	0.728	7.50	98%
SES mean original vs Alt2	100	0.774	0.759	10.17	96%
GWl composite original vs Alt2	100	0.502	0.521	7.83	99%

The dual-scoring pipeline was re-executed on a random 100-firm-year subsample (seed = 42) under two alternative prompt formulations. Alt1 substitutes the rubric vocabulary (e.g., “ambitious” → “extensive”, “verifiable quantitative” → “specific numerical”) while preserving anchor descriptions. Alt2 reorders anchors and restructures the rubric while preserving conceptual content.

Aggregate distributional moments are highly stable: the one-sample t-test of H1a yields $t = 6.76, 6.84, 6.93$ in the original, Alt1, and Alt2 versions ($p < 10^{-9}$ in all three), with mean composite GWI of 7.67, 7.67, and 7.50. Mean absolute differences between paired scores range from 3.75 to 19.29 points on the 0–100 scale, with 95–100% of paired observations falling within 25 points (one anchor spacing). Spearman rank correlations span 0.50–0.81. This pattern of stable distributional moments combined with imperfectly preserved firm-year rank orderings motivates the group-level interpretation of the GWI adopted throughout the dissertation.