

Greenwashing Without Lying: Claim-Type Selection in Corporate Climate Disclosure*

Hyun Hak Kim[†]

September 4, 2026

Abstract

Enforcement against corporate greenwashing presumes false claims that can be detected. Matching 565 sustainability reports from 335 Korean firms, sentence by sentence, against the government-certified emissions registry of the Korea Emissions Trading Scheme, we find almost nothing to detect—and argue that this is the problem. Where statements are falsifiable, they are accurate: disclosed domestic emission totals track the registry with a median deviation of 0.01%, and every whole-firm reduction claim whose year and scope can be resolved is consistent with the record. But falsifiable statements are scarce by design: 93.5% of substantive climate statements carry no quantitative content, and among whole-firm reduction claims only 7% assert a change in total emissions, while 63% report activity effects that remain literally true regardless of the aggregate—and that firms produce at the same rate whether their certified emissions rose or fell. Adoption of verifiable formats is predicted by firm size, not emission intensity, and is not detectably priced around disclosure-regulation events. We characterize this equilibrium as greenwashing by *claim-type selection*—deception without lying—and argue that the remedy is not policing statements but mandating falsifiable formats reconciled to registries that, for large emitters, already exist.

Keywords: greenwashing; climate disclosure; cheap talk; emissions trading scheme; textual analysis; Korea

*The research was conducted as a component of the Ministry of Climate, Energy and Environment (MCEE)-funded Graduate School Specialized in Climate Change.

[†]Professor, Department of Economics, Kookmin University, Seoul, Korea.

1 Introduction

Did the companies that claimed to have cut their greenhouse gas emissions actually cut them? For a sample of large Korean firms, the answer turns out to be yes—almost without exception. Yet this reassuring answer conceals a less reassuring fact: such claims are rare. Matching 565 sustainability reports from 335 Korean firms against the government-certified emissions registry of the Korea Emissions Trading Scheme (K-ETS), we find that corporate climate communication is arranged, with remarkable consistency, so that the question of truth barely arises. Firms do not misstate verifiable facts; they select claim types for which no verification is possible.

This paper documents that pattern and proposes a simple way to think about it, which we call *claim-type selection*. Climate-related statements differ in their falsifiability. A *stock* claim—“our total emissions declined by 10% year on year”—can be checked against an external record and is false if the record moved the other way. A *flow* claim—“our initiatives reduced emissions by 7,768 tons”—reports the estimated effect of particular activities; it remains literally true even if the firm’s total emissions rose. A target (“we will cut 40% by 2030”) and a project-level achievement (“the LED retrofit at one plant saved 2,591 tCO₂eq”) are likewise shielded from contradiction by any realized aggregate outcome. If misleading communication is costly only when it can be falsified, a firm that wishes to appear green without bearing that cost should not lie; it should choose the grammar of its claims. Our evidence indicates that this is precisely what large Korean firms do.

Four findings support this interpretation. First, verifiable statements are accurate. Where reports disclose domestic total emissions in a form comparable to the K-ETS registry, the median absolute deviation from the government-certified figure is 0.01%, and every whole-firm reduction claim that could be dated and scoped against the registry (six firm-years) is consistent with it. Second, verifiable statements are rare. Of 48,219 substantive climate statements extracted from the reports, 93.5% contain no quantitative content at all; among 256 sentences in which firms claim emission reductions at the whole-firm level, only 7% are stock claims, while 63% are flow claims. Third, unverifiable claims are produced independently of performance: firms whose registry emissions rose are statistically indistinguishable, in their propensity to advertise reduction accomplishments, from firms whose emissions fell. In one illustrative case, a firm publicized a reduction accomplishment of 7,768 tons in a year when its registry emissions increased by 15.7%. Fourth, neither disclosure quality nor its absence appears to be priced: across regulatory disclosure events, differences in cumulative abnormal returns between high- and low-quality disclosers do not survive multiple-testing correction in a sample of 170 listed firms.

We further ask what determines the quality of disclosure, understood as the adoption of verifiable formats (quantified interim targets, Scope 3 accounts, domestic-segment breakdowns). Environmental materiality does not: emission intensity is unrelated to every disclosure attribute we measure (p -values between 0.62 and 0.996). Firm size is: larger firms adopt verifiable formats significantly more often. Disclosure quality thus behaves like an organizational capability rather than a response to climate risk exposure—a pattern difficult to reconcile with the view that voluntary disclosure lets exposed firms reveal their exposure.

Our study extends two strands of research. The first compares the same firm’s carbon information across outlets. Depoers et al. (2016) contrast the emission figures firms report to the Carbon Disclosure Project with those in their own corporate reports and document systematic gaps between two *voluntary* channels. We move the comparison from figures to sentences, and from a voluntary benchmark to an administrative one: the K-ETS registry is measured under a statutory monitoring protocol, audited by accredited third-party verifiers, and certified by the Ministry of Environment, which makes it the closest available analogue to an answer key for corporate climate claims. The second strand measures the vagueness of climate talk. Bingler et al. (2022) show, using ClimateBERT, that TCFD-aligned disclosures are dominated by non-committal language, and Bingler et al. (2024) link such cheap talk to higher emissions growth. Our 93.5% non-quantification rate replicates their central finding in a new institutional setting; our contribution is to explain the *mechanism* that keeps talk cheap. Cheap talk persists not because firms fabricate, but because they populate their reports with claim types that cannot be brought into contact with the record. This also speaks to the selective-disclosure literature (Marquis et al., 2016): selection operates not only over which facts to reveal but over the logical form in which anything is asserted at all.

The distinction matters for policy. Enforcement against greenwashing is typically framed as the detection and punishment of false claims. In our data there are, for practical purposes, no false claims to punish. The binding constraint is not honesty but falsifiability: 60% of firms do not disclose emissions in any form that the national registry can corroborate, and the dominant reduction rhetoric is immune to contradiction by construction. The corresponding remedy is not stricter penalties for misstatement but the mandating of verifiable formats—domestic legal-entity totals, location-based accounting, and year-on-year figures—that convert unfalsifiable prose into checkable statements.

In sum, the paper makes three contributions: an identification design that benchmarks corporate claims against an administrative registry rather than against other language or self-reports; a falsifiability typology of climate claims, and with it the claim-type-selection account of how greenwashing operates without falsehood; and a methodological template for claim-level verification. We fine-tune a Korean-language transformer classifier on human-

labeled sentences to isolate whole-firm reduction claims (out-of-fold $\kappa = 0.70$), develop a typology separating stock from flow claims using grammatical case marking, and construct a four-step procedure for aligning the year and scope of a textual claim with registry observations. We document, candidly, that a naive alignment produced spuriously significant “contradictions” that dissolved under case-by-case verification; the protocol we report is the one that survived.

The remainder of the paper proceeds as follows. Section 2 develops the claim-type selection argument and derives our hypotheses. Section 3 describes the Korean institutional setting and the registry. Section 4 presents the data and measurement pipeline. Section 5 reports the results, Section 6 discusses implications, and Section 7 concludes.

2 Theory and Hypotheses

2.1 From misrepresentation to claim-type selection

Greenwashing is conventionally defined as the conjunction of poor environmental performance and positive environmental communication (Delmas and Burbano, 2011), though later work notes that in practice the label is applied through accusation rather than adjudication, precisely because outside observers can rarely establish whether a green message is misleading (Seele and Gatti, 2017). The definition is silent, however, about *how* the positive communication is produced, and the literature has explored two broad answers. The first is misrepresentation: firms overstate what they do or achieve, and enforcement is a matter of detecting falsehoods, whether by auditors, activists, or regulators (Lyon and Maxwell, 2011). The second is selection: firms disclose truthfully but partially, revealing favorable facts and withholding unfavorable ones (Marquis et al., 2016; Kim and Lyon, 2015). In the selective-disclosure view the unit of choice is the *fact*—which metrics, plants, or programs to mention.

We propose that an important third margin lies between these two: the choice of *claim type*. Holding fixed both the underlying facts and the decision to communicate positively about them, statements differ in whether they can, even in principle, be brought into contact with an external record. A firm that wishes to appear green at minimal expected cost need not misrepresent, and need not even suppress unfavorable facts; it can compose its communication out of assertion forms that no record can contradict. We call this margin *claim-type selection*. It is conceptually closest to the “cherry-picking” documented by Bingler et al. (2022), who show that TCFD-aligned disclosures concentrate on non-material categories; our argument is that the selection operates not only across topics but across the logical form

of individual sentences, and we test it at that level.

2.2 Sustainability reporting between impression management and verification

A substantial literature on sustainability reporting anticipates our starting intuition—that reports are instruments of impression rather than information—without being able to complete the argument. Content analyses of GHG reporting document the strategies: energy firms rationalize their climate impacts through selective justification (Talbot and Boiral, 2018), firms legitimize the negative aspects they cannot avoid mentioning through abstraction and corrective framing (Hahn and Lülfs, 2014), and the quality of GHG information—a “credence” good the reader cannot independently assess—fails to improve with reporting experience (Comyns and Figge, 2015). The umbrella reviews of greenwash catalogue these devices at every level from product labels to selective disclosure (Lyon and Montgomery, 2015). What this largely qualitative tradition cannot do, however, is adjudicate: lacking an external record, it can classify a statement as an impression-management device but cannot say whether the impression is false. Our design closes exactly this gap. The registry converts a subset of impression-management devices into testable propositions, and the empirical surprise is that the tested subset passes: impression management here operates not by distorting checkable content but by rationing it.

The claim-type lens also reframes the decoupling literature. Organizational scholarship distinguishes policy–practice decoupling, where adopted policies go unimplemented, from means–ends decoupling, where policies are genuinely implemented yet fail to move the outcomes they nominally serve (Bromley and Powell, 2012); sustainability reporting has been read as the façade that manages the resulting gap (Cho et al., 2015). Flow claims are the linguistic signature of means–ends decoupling: the firm truthfully reports its means—the retrofit, the fuel switch, the tons attributed to each—while the end, the aggregate trajectory, goes unstated. Our contribution to this literature is to show that the façade requires no false statements at all, and to measure its dominant grammatical form directly. The closest analogue is recent work on the *temporal* axis of disclosure: GHG targets are routinely set in hindsight, shifted, or disbanded without explanation (Avidan et al., 2026). Claim-type selection is the cross-sectional counterpart—where temporal inconsistency evades accountability by moving the goalposts across reports, claim-type selection evades it within a single report by choosing assertion forms no record can reach.

Two adjacent strands sharpen what is and is not new in our approach. The first is the rapidly growing measurement literature that applies language models to climate disclosure.

Since domain-adapted transformers became available (Webersinke et al., 2021), this work has scored disclosures for specificity, commitment, and cheap-talk content, and a recent survey catalogues more than sixty such studies (Calamai et al., 2025). Its common limitation, from our vantage point, is the benchmark: language is scored against other language, against ESG ratings, or against self-reported emissions, so that even sophisticated “greenwashing detectors” cannot say whether any particular claim is *true*. Where hard numbers have been compared across outlets, both sides of the comparison were voluntary self-reports (Depoers et al., 2016). Our contribution to this strand is the benchmark itself: a statutory, third-party-verified, government-certified registry, matched to claims at the sentence level.

The second strand is the Korean evidence, which has examined the K-ETS and carbon disclosure largely through the lens of firm value and disclosure volume. Mandatory emissions disclosure is associated with lower valuations and a smaller foreign investor base (Han et al., 2023); carbon emissions relate to value differently for business-group affiliates than for independent firms (Lee and Cho, 2021); and content analyses find that firms with favorable carbon performance disclose more (Kim and Kim, 2022). Most tellingly for our purposes, firms covered by the K-ETS did not reduce emissions in its early phases (Oh et al., 2023), yet their environmental *ratings* improved (Lee et al., 2024)—a wedge between measured performance and its public representation that this literature records but does not explain. Our claim-level evidence supplies the missing mechanism: the representation improves because it is built from claim types that performance cannot contradict. No prior Korean study, to our knowledge, matches report language to the certified registry at the level of individual statements.

2.3 A falsifiability-cost framework

Consider a firm whose realized change in total emissions over a reporting period is Δ , observed by the firm but communicated to outsiders through a report. The firm chooses a set of climate-related claims; each claim yields an impression benefit and carries an expected cost of being shown false. Readers reward green impressions, and—critically—the impression conveyed by “reduction” language is largely insensitive to its logical form: “we reduced our emissions by 10%” and “our initiatives reduced emissions by 45,000 tons” leave similar traces on a non-expert reader, although only the first is an assertion about Δ .

The expected cost side is where claim types separate. A claim can be falsified only if three alignment conditions hold between the claim and some external record: the claim’s *scope* coincides with the record’s (same legal entities, same gas accounting boundary), its *period* can be dated, and its *metric* is the one the record measures. Table 1 classifies the

Table 1: A typology of emission-reduction claims by falsifiability

Claim type	Example	Checkable against registry?	Cost if aggregate emissions rose ($\Delta > 0$)
Stock (aggregate level)	“Total emissions fell 10% year on year.”	Yes	Falsified: claim asserts $\Delta < 0$
Flow (activity effect)	“Our initiatives cut emissions by 45,000 tons.”	No	None: literally true even if $\Delta > 0$
Target (future)	“We will cut 40% by 2030.”	Not yet	None within the period
Project (partial scope)	“The LED retrofit saved 2,591 tCO ₂ eq at one plant.”	No (scope mismatch)	None
Intensity	“Emissions per revenue fell 13%.”	No (different metric)	None

reduction-claim types we observe in sustainability reports by these conditions, taking as the external record a mandatory emissions registry that certifies firm-level total emissions per calendar year.

Only the stock claim satisfies all three alignment conditions; every other type fails at least one by construction. A flow claim reports the estimated effect of chosen activities against an implicit business-as-usual counterfactual that no registry records; a project claim’s scope is narrower than any registry entity; an intensity claim substitutes a ratio the registry does not certify; a target postpones the period beyond any existing record.

Two behavioral predictions follow for a firm that values green impressions and bears costs only upon demonstrated falsehood. First, whatever stock claims it does make should be true: the expected-cost asymmetry bites exactly there, and a false stock claim is the one message an outside party can convict. Second, the *composition* of its claims should tilt away from the stock form, and the production of unverifiable claims should be decoupled from Δ : because flow, project, and intensity claims cost nothing regardless of the aggregate outcome, there is no mechanism linking their incidence to realized performance. In the resulting communication equilibrium, reduction rhetoric is abundant, truthful in its falsifiable component, and uninformative about aggregate performance—a micro-foundation for the “cheap talk” regularity documented at the corpus level by Bingler et al. (2022). The logic parallels strategic-communication models in which message credibility derives from the cost of lying (Crawford and Sobel, 1982): where the institutional environment prices only falsifiable lies, talk migrates to where it is free.

2.4 Hypotheses

The framework organizes our empirical analysis into one descriptive question and four hypotheses.

RQ1. What is the claim-type composition of corporate emission-reduction communication?

Claim-type selection predicts that the stock form is rare relative to unverifiable forms. This is descriptive because the framework does not pin down magnitudes, only the direction of the tilt.

Hypothesis 1 (H1): *Verifiable (stock) reduction claims, where they occur, are consistent with the registry record.*

Hypothesis 2 (H2): *The incidence of unverifiable (flow) reduction claims is unrelated to the firm’s realized change in registry emissions.*

H1 and H2 jointly discriminate claim-type selection from the misrepresentation view of greenwashing, which predicts falsified stock claims, and from a naive honesty view, under which reduction rhetoric of any form should track realized reductions.

The third hypothesis concerns who adopts verifiable formats—quantified interim targets, Scope 3 accounts, and domestic-segment breakdowns—given that adopting them is costly and self-binding. Economics-based disclosure theory predicts that firms with genuinely favorable positions separate themselves by disclosing verifiable information (Verrecchia, 1983; Clarkson et al., 2008), and legitimacy theory predicts that environmentally exposed firms disclose more in response to social pressure (Patten, 1992). Both views tie disclosure quality to environmental *materiality*—in our setting, emission intensity. An alternative capability view ties it instead to organizational resources: verifiable formats require measurement systems, consolidated carbon accounting, and staff, which scale with firm size irrespective of exposure. The two views yield competing predictions that our data can separate:

Hypothesis 3a (H3a, materiality): *Disclosure quality increases with emission intensity.*

Hypothesis 3b (H3b, capability): *Disclosure quality increases with firm size, conditional on intensity.*

Finally, if verifiable disclosure separates types, investors should price it when disclosure regulation becomes more or less likely: firms already publishing verifiable formats stand to

bear lower compliance costs and enjoy lower information risk. Conversely, in the claim-type-selection equilibrium in which published talk carries little information about performance, there is little for the market to price.

Hypothesis 4 (H4): *Around regulatory disclosure events, cumulative abnormal returns do not differ systematically between firms with high and low disclosure quality.*

We state H4 in its null-consistent form deliberately: under our framework the absence of differential pricing is not a failed test but an equilibrium property. We nonetheless treat it as the weakest link in the chain—absence of evidence bounds, rather than proves, the claim—and report power along with the estimates.

3 Institutional Setting

3.1 The K-ETS registry as an answer key

Launched in 2015 under the Act on the Allocation and Trading of Greenhouse Gas Emission Permits (2012), the Korea Emissions Trading Scheme (K-ETS) was the first nationwide mandatory cap-and-trade program in East Asia. Inclusion is not elective: any company whose annual emissions exceed 125,000 tCO₂eq, or that operates an installation exceeding 25,000 tCO₂eq, is a covered entity. In its third phase (2021–2025) the scheme covered 685 companies across 69 subsectors, accounting for 73.5% of national greenhouse gas emissions (International Carbon Action Partnership, 2025). Covered emissions comprise direct emissions and indirect emissions from externally supplied electricity and heat—that is, Scope 1 plus location-based Scope 2 at the level of the domestic legal entity.

What makes the K-ETS useful for our purposes is not the trading system but its measurement infrastructure. Each covered company operates under an approved monitoring plan; its annual emissions statement must be audited by an independent, government-accredited third-party verifier before submission; and the verified statement is then reviewed and certified by the Certification Committee of the Ministry of Environment. Because allowance obligations—and therefore money—ride on these figures, both the firm and the verifier face sanctions for misstatement, and firms have, if anything, an incentive to avoid *understating* reductions. The certified firm-year emissions are published through the national greenhouse gas management system, yielding a public, government-certified panel of firm-level emissions.

The registry therefore satisfies, by construction, the three alignment conditions of Section 2.3 for exactly one claim type. Its scope is fixed (domestic legal entity), its period is

fixed (calendar year), and its metric is fixed (total Scope 1 plus location-based Scope 2, in tCO₂eq). A corporate statement asserting a change in total emissions can be checked against it; a statement about activity effects, projects, intensity ratios, or future targets cannot. In this sense the registry functions as an answer key—not for everything a firm says, but for the one falsifiable thing it might say. To our knowledge, no comparably certified, statutory, firm-level emissions record exists alongside a large voluntary reporting corpus in most jurisdictions; where prior work has compared carbon figures across outlets, both sides of the comparison were voluntary self-reports (Depoers et al., 2016).

Two features of the setting bear on interpretation. First, the registry covers the domestic legal entity, whereas sustainability reports commonly consolidate global operations. This wedge is not merely a nuisance: a firm choosing its reporting boundary chooses, at the same time, whether its statements can be checked. We therefore treat the disclosure of a domestic-segment breakdown as a measurable attribute of disclosure quality—the adoption of a *verifiable format*—rather than assuming away the mismatch, and we restrict claim verification to statements whose scope demonstrably matches the registry. Second, because covered firms already report verified emissions to the state, any vagueness in their public reports is not a matter of capacity to measure: the certified number exists. What varies is the decision to expose it to comparison.

3.2 Voluntary reporting and the disclosure-regulation timeline

Against this mandatory registry stands an almost entirely voluntary public-reporting regime. Korean sustainability reports are published at the firm’s discretion, typically annually, without a prescribed format or assurance requirement; the Korea Chamber of Commerce and Industry maintains a repository to which large firms submit their reports, which we use for sampling (Section 4). Mandatory sustainability disclosure has been repeatedly announced and repeatedly deferred. The Financial Services Commission’s original plan to phase in ESG disclosure from 2025 was postponed in October 2023 to “2026 or later”; in February 2026 the Commission released a draft roadmap under which disclosure becomes mandatory in stages beginning in 2028, starting with listed groups whose consolidated assets exceed KRW 30 trillion.¹

The combination—certified private measurement, discretionary public disclosure—is what gives the setting its analytical leverage. Firms know their audited totals; outsiders can look those totals up; and yet the public narrative is composed almost entirely of statements the registry cannot touch. The next section describes how we measure that narrative and bring

¹Financial Services Commission press releases of October 16, 2023 (deferral) and February 25, 2026 (draft roadmap). We exploit both announcements as events in Section 5.

it into contact with the record.

4 Data and Methods

4.1 Sample and corpus construction

Our report corpus originates from the sustainability-report repository maintained by the Korea Chamber of Commerce and Industry, which we scraped in full and supplemented with a small number of manually collected reports, yielding 570 report PDFs. After screening out English-language editions and quarantining eight reports that the repository had attributed to the wrong member of a business group,² the estimation corpus comprises 565 firm-report pairs from 335 firms, published between 2022 and 2026 (with 88, 166, 148, 130, and 33 reports per publication year, respectively). Of these firms, 171 can be matched to financial statements filed with the Korean regulatory filing system (DART), 141 belong to a large business group designated by the Korea Fair Trade Commission, and 94 appear in our K-ETS registry panel.

Text extraction from Korean sustainability reports is nontrivial because the documents interleave dense data tables with prose in multi-column layouts. We extract text with a layout-aware parser that removes table regions using detected table bounding boxes, splits lines at large horizontal gaps, and flags residual table fragments (rows with three or more data values and no Korean predicate ending) for exclusion; tables are retained separately for the structured extraction described below. The resulting prose corpus contains 81,030 sentences. This separation matters: in an earlier iteration that mixed table fragments into prose, our vagueness measure correlated at $r = -0.96$ with the share of digits per sentence—measuring parsing yield rather than rhetoric—and the correlation became statistically indistinguishable from zero only after tables and table fragments were removed.

4.2 Measuring the narrative and report-level attributes

Each prose sentence with climate-related content is classified into four categories: concrete targets (quantified future commitments), past performance (quantified realized outcomes),

²Partial-name matching in the repository assigned, for example, a battery subsidiary’s report to its similarly named parent. We detected these cases through implausible emission comparisons, verified each by inspection, and excluded the affected firms when no valid report remained. The eight quarantined files (folder → actual reporting entity) are: SK → SK Eternix (2025); Sebang → Sebang Global Battery (2023, 2024); SeAH Steel → SeAH Steel Holdings (2023); Shinsegae → Shinsegae Property (2024); Hanwha → Hanwha Vision (2026); Hyosung → Hyosung Advanced Materials (2023) and Hyosung Chemical (2025). Collection subsequently switched to exact matching on the company-name core extracted from report titles.

non-quantitative climate statements, and noise (boilerplate, tables of contents, assurance statements). Rule-based classification proved inadequate—against human labels drawn from a random representative sample its agreement was $\kappa = 0.14$, largely because vague statements lack reliable surface markers—so we fine-tune a Korean RoBERTa transformer on human-labeled sentences. The labeling proceeded in three rounds totaling 962 sentences; crucially, evaluation uses only a 200-sentence simple random sample of the corpus, with enriched strata used for training alone, so that reported performance is not inflated by over-sampling rare classes. Guideline revisions between rounds—classifying tables of contents, assurance statements, and GRI index lines as noise, and requiring an explicit figure for the two quantified classes—are documented in the replication materials. The final classifier attains out-of-fold $\kappa = 0.57$ (accuracy 0.76) on the representative sample after a logit adjustment that corrects the prior shift between training and corpus class frequencies. For our headline quantity—the share of substantive climate statements that contain no quantitative content—the model is nearly unbiased: on the representative sample it estimates 89.6% against a direct human estimate of 89.9% (95% CI 83.6–94.4), and applied to the full corpus it yields 93.5% of 48,219 substantive statements. We report the model figure with the human-sample confidence interval alongside throughout.

From the same corpus we code six report-level attributes of disclosure quality: whether the firm quantifies an interim (pre-2050) reduction target, discloses Scope 3 emissions, discloses a domestic-segment emissions breakdown, mentions the Science Based Targets initiative, names a net-zero target year, and the non-quantification share defined above. In addition, we determine *verifiable-format adoption* from the tables themselves: a structured scan of every report’s emission tables records whether a domestic legal-entity total appears at all, using merged-cell-preserving table extraction. Where a domestic total is disclosed, we compare it directly with the registry figure for the same firm-year; this number-to-number comparison, in the spirit of Depoers et al. (2016), underpins the accuracy results in Section 5.

4.3 Identifying and verifying reduction claims

The core of the paper requires isolating sentences in which a firm claims to have reduced emissions at the whole-firm level. We proceed in two stages, favoring recall first and precision second. A deliberately broad lexical filter (past-tense reduction verbs co-occurring with greenhouse-gas terms) retrieves 1,241 candidate sentences from 265 firms. A stratified sample of 342 candidates was then hand-labeled by the author into four classes: whole-firm reduction achievement (94), partial-scope achievement (59), target or plan (79), and not applicable (110)—the last including statements about other actors (“the EU will cut. . .”) and sentence

fragments. The exercise confirmed that rules alone cannot do this work: the rule-based prefilter’s precision for whole-firm claims was 53%, and a third of true whole-firm claims were hiding in its residual category. Labeling followed three ordered questions—future tense $\rightarrow$ target or plan; partial scope $\rightarrow$ partial achievement; subject not the reporting firm, or an unintelligible fragment, $\rightarrow$ not applicable; otherwise whole-firm achievement—and the labeler consulted neither the source PDFs nor any emissions data, so labels cannot encode knowledge of subsequent verification outcomes.

We fine-tune and compare two Korean transformer models on these gold labels under identical five-fold, two-seed protocols, selecting by the F1 of the whole-firm class ($F1 = 0.85$, $\kappa = 0.70$ for the larger model, which we adopt). Sentences are presented to the model with their neighboring sentences as context, mirroring the information the human labeler used. A five-seed ensemble then classifies all 1,241 candidates, yielding 256 whole-firm reduction claims from 131 firms.

Verifying a textual claim against a registry observation requires aligning three things—metric, scope, and period—and each alignment step embodies a substantive decision that we make explicit.

Metric and scope screens. We first exclude claim sentences whose metric the registry does not certify: intensity claims (per-revenue ratios), financed-emissions claims by financial firms, statements about other actors, and sentences that themselves concede an aggregate increase. We then separate *stock* from *flow* claims (Table 1) using Korean case marking, which turns out to be nearly diagnostic: in a stock claim, “emissions” is the subject of an intransitive decline (*paechullyang-i/-eun kamsohaetta*, “emissions fell”), whereas in a flow claim it is the object of a transitive reduction with a stated amount (*paechullyang-eul X-ton kamchukhaetta*, “we reduced emissions by X tons”). Claims quantified relative to a base year more than one year prior are routed to a separate baseline check rather than the year-on-year comparison.

Period resolution. Report titles do not reliably identify the data year: in our corpus, a report labeled with year t contains data through t about as often as through $t - 1$. We therefore resolve each claim’s data year in a fixed order: (i) if the sentence states an emissions *level*, match it against the firm’s registry levels and adopt the year of a match within 1.5%—the strongest evidence, given that disclosed domestic totals track the registry to within a fraction of a percent; (ii) otherwise use an explicit year mentioned in the sentence; (iii) otherwise use the latest data year appearing in the report’s own emission tables; (iv) otherwise, compare against both candidate years and record a verdict only if the registry moved in the same direction in both. Claims whose stated level matches no registry year within tolerance are classified as scope-mismatched—the claim is about a different reporting boundary—and

marked unverifiable rather than false. Two further guards apply: a tonnage immediately followed by a reduction verb (“fell by 26,000 tons”) is a change, not a level, and is skipped in step (i); and a claimed percentage matching one candidate year’s registry change within ± 2 percentage points resolves the year directly. No claim’s verdict was assigned by hand.

We adopted this protocol after a cautionary experience that we report rather than bury: under the naive convention that a report labeled t covers year $t-1$, we initially found that 43% of claiming firm-years showed registry *increases*, a nominally significant result ($p = 0.02$). Case-by-case inspection of every apparent contradiction showed each to be an artifact of year misalignment or claim-type conflation; under the protocol above, the contradictions dissolve. The episode illustrates how easily claim-verification exercises can manufacture greenwashing findings, and we treat it as a methodological result in its own right.

4.4 Registry, financial, and market data

The registry panel covers 95 firms exactly matched to K-ETS certified emissions over 2020–2024 (418 firm-years); 94 of them publish reports in our corpus. Name matching uses full legal names with explicit exclusion lists, since partial matching in a business-group economy reliably attaches subsidiaries to parents. Financial controls (assets, leverage, profitability) come from DART filings; business-group affiliation comes from the Fair Trade Commission’s official rosters of designated-group member companies, which identify affiliates whose names do not contain the group name. For the event study we use daily closing prices for the 171 financially matched firms (168 with sufficient price history) and the KOSPI index, estimating market-model abnormal returns over trading days $[-250, -30]$ relative to each event and cumulating them over $[-5, +20]$, $[-1, +1]$, and $[0, +5]$; inference uses Welch and Mann–Whitney tests with Holm correction across the full comparison grid, specified before estimation.

5 Results

5.1 The composition of climate talk

Table 2 reports descriptive statistics and Table 3 the classifier performance underlying all text-based measures. The first result is the sheer scarcity of checkable content. Across the corpus, 93.5% of substantive climate statements contain no quantitative content (human-sample estimate 89.9%, 95% CI 83.6–94.4); the share is remarkably stable across subsamples, moving by less than a percentage point as the sample grew from 9 to 335 firms. At the level of report attributes, only 40% of firms (31 of the 77 whose emission tables we audited in

Table 2: Descriptive statistics (335 report-publishing firms)

Variable	N	Mean	SD	Median
Non-quantification share (prose)	335	0.935	0.047	0.940
Quantified interim target (0/1)	335	0.492	0.458	0.500
Scope 3 disclosure (0/1)	335	0.792	0.374	1.000
Domestic-segment disclosure (0/1)	335	0.316	0.433	0.000
SBTi mention (0/1)	335	0.248	0.410	0.000
Net-zero year stated (0/1)	335	0.729	0.416	1.000
Business-group affiliate (0/1)	335	0.421	0.494	0.000
ln(Total assets)	171	12.801	0.676	12.690
Leverage (debt/equity)	171	2.461	6.936	1.136
Emission intensity (tCO ₂ eq/KRW bn)	81	23.818	79.219	4.901

Financial variables are available for 171 firms; emission intensity requires both the K-ETS registry match and sales data (81 of the 94 registry-matched firms).

Table 3: Classifier performance, evaluated on representative samples only

Task	Method	κ	Acc.	Macro F1	F1 (key class)
Sentence stance (4 classes, $n=200$)	Rule-based	0.076	—	—	—
	RoBERTa-base + prior adj.	0.568	0.763	0.748	—
Reduction-claim type (4 classes, $n=342$)	RoBERTa-base	0.677	0.762	0.742	0.834
	RoBERTa-large (adopted)	0.704	0.782	0.765	0.847

Out-of-fold estimates (5-fold, averaged over two seeds). “Key class” is the whole-firm reduction claim. Enriched strata are used for training only; all figures are computed on random representative samples.

structured form) disclose emissions in a table that a domestic registry could corroborate at all; the corpus-wide prose-based measure gives a similar 31.6% (Table 2).

The same pattern reappears, in sharper form, within the reduction claims themselves (Table 4, Panel A). Of 256 sentences in which firms assert whole-firm emission reductions, only 18 (7.0%)—20 including base-year-relative variants—take the stock form that the registry can falsify. Flow claims outnumber them nine to one (161 sentences, 62.9%), with intensity claims (5.9%) and non-comparable metrics adding further unfalsifiable variants. Reduction rhetoric, in other words, is plentiful; reduction *assertions* are rare. This is the tilt predicted by claim-type selection, and its magnitude is the paper’s central descriptive fact (Figure 1).

5.2 Accurate stock claims, decoupled flow claims

Where firms do expose numbers to comparison, the numbers are right. In the hand-audited pilot subsample of disclosed domestic emission totals, the median absolute deviation from

Table 4: Whole-firm reduction claims: composition and verification

<i>Panel A: composition of 256 whole-firm reduction claim sentences (131 firms)</i>		
Claim type	Sentences	Share
Stock (aggregate-level, falsifiable)	18	7.0%
Stock, base-year relative	2	0.8%
Flow (activity effect, unfalsifiable)	161	62.9%
Intensity metric	15	5.9%
Non-comparable metric (financed etc.)	5	2.0%
Concedes an increase	4	1.6%
Residual (fragmentary)	51	19.9%
<i>Panel B: verification of stock claims against the registry (firm-years)</i>		
Verifiable, consistent with registry	6	
Verifiable, contradicted by registry	0	
Unverifiable: scope mismatch	2	
Unverifiable: registry year unavailable	1	
Unverifiable: data year unresolvable	5	
<i>Panel C: flow claims and realized registry direction (firm-years)</i>		
Flow-claiming firm-years matched to registry	30	
of which registry emissions <i>rose</i>	14	46.7%
Base rate of increases, all registry firm-years		58.0%
Binomial test vs. base rate (two-sided)		$p = 0.27$

Composition classified by the case-marking and screening rules of Section 4.3. Verification uses the four-step year-resolution protocol; scope mismatch means the stated emissions level matches no registry year within 1.5%, indicating a different reporting boundary.

the registry figure is 0.01% across 23 firm-year observations; automated extraction over the full corpus adds a further eleven observations within 1.3%, including several exact matches to the ton (Figure 2). Turning from figures to sentences, Panel B of Table 4 reports the verification of stock claims: of the firm-years in which the claim’s year and scope could be pinned down, all six are consistent with the registry, and none is contradicted; Table 5 lists all six with the year-resolution basis and translated excerpts. The remaining stock claims are unverifiable for identifiable reasons—two state levels that match no registry year, indicating a consolidated or global boundary, and six cannot be dated or fall outside the registry window—and we decline to score them either way.

Six firm-years is a small base, and we do not claim to have proven the absence of false statements. The proper reading combines H1 with RQ1: the population of falsifiable claims is itself so small that honesty within it is cheap to maintain. What the data rule out is the misrepresentation account of greenwashing, in which positive communication rests on statements an auditor could convict.




Figure 1: Composition of the 256 whole-firm reduction claim sentences. Falsifiable stock claims (dark) are outnumbered nine to one by unfalsifiable flow claims.

Table 5: All resolved stock claims and registry outcomes

Firm	Year	Basis	Registry Δ	Claim (translated excerpt)
Hyundai WIA	2023	report coverage	−11.0%	“emissions fell 26,000 tons year on year”
KOGAS	2024	%-match	−11.3%	“cut greenhouse gases 11% from the prior year”
Korea Zinc	2023	two-sided	−5.0%	“achieved an 11% year-on-year reduction”
Samsung C&T	2024	explicit year	−0.8%	“2024 emissions down 10% from 2023”
Lotte Chemical	2022	report coverage	−12.0%	“emissions declined [with lower utilization]”
Sungshin Cement	2023	report coverage	−2.4%	“total emissions decreased in the current period”

Registry Δ is the firm’s certified year-on-year change. All six claims assert a decline and the registry declines in all six cases. Verification is of direction; two firms’ claimed magnitudes (10–11%) exceed the registry change (−0.8%, −5.0%), consistent with consolidated reporting boundaries wider than the domestic registry entity.

If flow claims were an honest summary of aggregate progress, firms whose registry emissions rose should produce fewer of them. They do not. Panel C of Table 4 matches flow-claiming firm-years to the registry (Figure 3): 46.7% of them fall in years when the firm’s certified emissions *increased*, statistically indistinguishable from the 58.0% base rate of increases across all registry firm-years ($p = 0.27$). Advertising reduction accomplishments is thus unrelated to the direction of realized aggregate change—precisely the decoupling that the falsifiability-cost framework predicts, since flow claims cost nothing either way.

Concrete instances convey the mechanism. A casino operator publicized reduction accomplishments of 7,768 tons in a year when its registry emissions rose 15.7%; a defense manufacturer reported cutting 1,209 tCO₂eq while its certified total rose 11.8%; a semiconductor firm described initiatives saving 12,029 tCO₂eq against a certified increase of 10.1%.




Figure 2: Disclosed domestic emission totals against K-ETS certified figures (log–log scale). Circles are hand-audited pilot observations; triangles are automated, scope-matched extractions. The line is the 45° locus on which report and registry agree.

None of these statements is false. Each reports a genuine activity whose estimated effect is real; each leaves the aggregate trajectory unstated; and in each case the certified aggregate moved the other way.

5.3 Size, not exposure, predicts verifiable disclosure

Table 6 confronts the materiality and capability views. Emission intensity—the materiality measure—predicts nothing: across all six disclosure attributes its coefficient is essentially zero (p -values 0.33–0.996 in the registry subsample). Firm size, by contrast, is a strong and consistent predictor of exactly the attributes that expose the firm to verification: conditional on intensity, log assets carries coefficients of +0.235 ($p = 0.003$) for quantified interim targets, +0.182 ($p = 0.002$) for Scope 3 disclosure, and +0.218 ($p = 0.004$) for domestic-segment disclosure, with the same pattern (smaller coefficients) in the larger financial sample. Business-group affiliation adds little once size is controlled: its only conventionally significant association is with naming a net-zero year (+0.241, $p = 0.044$)—a declarative attribute—with a marginal association for Scope 3 ($p = 0.089$); it does not predict the self-binding




Figure 3: Flow claims are decoupled from realized performance. Bars show the distribution of certified year-on-year emission changes across all registry firm-years; markers locate the firm-years in which firms advertised reduction accomplishments.

attributes, and we read the group results as suggestive at most. Notably, no regressor explains the non-quantification share, consistent with its near-universality: vagueness is the baseline, and what varies with capability is the adoption of verifiable formats on top of it.

These results are difficult to reconcile with either economics-of-disclosure separation (H3a via signaling) or legitimacy pressure (H3a via exposure), both of which tie disclosure to materiality. They are what the capability view predicts: verifiable disclosure is produced by firms that can afford the accounting infrastructure, irrespective of how much carbon is at stake.

Nor are these conclusions an artifact of using each firm's latest report. Re-estimating on the report-year panel (329 report-years, 171 firms, publication years 2022–2026) with year fixed effects and standard errors clustered by firm leaves them unchanged: log assets predicts all three verifiable formats (+0.118 to +0.158, all $p \leq 0.02$), and in the intensity subsample the intensity coefficient remains null throughout ($p = 0.69$ – 0.77) while the size coefficients are, if anything, larger. Group affiliation is significant for two attributes in the panel but not in the cross section—an instability across specifications that is itself the reason we treat group effects as suggestive.

Table 6: Determinants of disclosure quality: size versus materiality

	Interim target	Scope 3	Dom. segment	SBTi	Net-zero yr	Non-quant. share
<i>Panel A: financially matched sample ($n = 171$)</i>						
ln(Assets)	+0.130** (0.012)	+0.160*** (0.000)	+0.082* (0.088)	+0.090** (0.035)	+0.024 (0.607)	+0.002 (0.587)
Group affiliate	+0.039 (0.615)	+0.150** (0.026)	+0.097 (0.152)	+0.006 (0.916)	+0.087 (0.229)	+0.002 (0.767)
<i>Panel B: registry subsample with intensity ($n = 81$)</i>						
ln(Assets)	+0.235*** (0.003)	+0.182*** (0.002)	+0.218*** (0.004)	+0.047 (0.504)	+0.010 (0.883)	+0.006 (0.330)
Group affiliate	-0.156 (0.205)	+0.199* (0.089)	+0.017 (0.895)	+0.102 (0.290)	+0.241** (0.044)	+0.003 (0.853)
ln(Intensity)	-0.000 (0.996)	-0.020 (0.699)	-0.034 (0.640)	-0.005 (0.919)	+0.051 (0.364)	+0.002 (0.751)

OLS coefficients with HC1 robust p -values in parentheses; all regressions control for leverage. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.4 The market does not price disclosure quality

Table 7 reports the event study. Across two regulatory events, three windows, and three disclosure-quality splits—eighteen pre-specified comparisons over 170 firms—two differences are nominally significant and none survives Holm correction. The eighteen comparisons would be expected to produce about one nominal hit by chance. An earlier version of this test on a 73-firm subsample produced a single nominally significant comparison ($p = 0.032$); in the enlarged sample the same comparison returns $p = 0.689$, a clean non-replication that reinforces the null reading.

One pattern deserves mention without overclaiming. Both nominal hits attach to the February 2026 roadmap event, and the more coherent of them is a +2.65 percentage-point three-day CAR for firms disclosing a domestic-segment breakdown (t - $p = 0.017$; the same variable is the only one to approach significance in a cross-sectional regression with size and group controls). The direction—verifiable-format firms gaining when mandatory disclosure becomes concrete—is what a compliance-cost or information-risk story would predict. With a Holm-adjusted p of 0.30 we treat it as suggestive only; we return to it as a design for future work in Section 6. Taken together, H4’s null is consistent with the equilibrium of Section 2.3: talk that carries no information about performance gives the market little to price.

Table 7: Event study: CAR differences by disclosure quality (170 firms, 18 pre-specified comparisons)

Event	Split	Window	Diff.	t - p	Holm- p	n
Oct 2023 deferral	Scope 3	[-1,1]	-0.64%	0.355	1.000	138+25
Oct 2023 deferral	Scope 3	[-5,20]	-1.39%	0.586	1.000	138+25
Oct 2023 deferral	Scope 3	[0,5]	-1.91%	0.314	1.000	138+25
Oct 2023 deferral	Dom. segment	[-1,1]	-0.22%	0.715	1.000	62+101
Oct 2023 deferral	Dom. segment	[-5,20]	+2.52%	0.216	1.000	62+101
Oct 2023 deferral	Dom. segment	[0,5]	+0.22%	0.817	1.000	62+101
Oct 2023 deferral	Interim target	[-1,1]	-0.23%	0.689	1.000	89+74
Oct 2023 deferral	Interim target	[-5,20]	+0.20%	0.914	1.000	89+74
Oct 2023 deferral	Interim target	[0,5]	+0.76%	0.399	1.000	89+74
Feb 2026 roadmap	Scope 3	[-1,1]	+1.26%	0.318	1.000	136+24
Feb 2026 roadmap	Scope 3	[-5,20]	-7.29%	0.099	1.000	136+24
Feb 2026 roadmap	Scope 3	[0,5]	+0.89%	0.523	1.000	136+24
Feb 2026 roadmap	Dom. segment	[-1,1]	+2.65%	0.017	0.303	61+99
Feb 2026 roadmap	Dom. segment	[-5,20]	+3.55%	0.257	1.000	61+99
Feb 2026 roadmap	Dom. segment	[0,5]	+1.62%	0.223	1.000	61+99
Feb 2026 roadmap	Interim target	[-1,1]	+1.15%	0.258	1.000	86+74
Feb 2026 roadmap	Interim target	[-5,20]	-6.71%	0.023	0.395	86+74
Feb 2026 roadmap	Interim target	[0,5]	-0.26%	0.840	1.000	86+74

Market-model CARs; difference is high- minus low-quality group mean. Holm correction over all 18 comparisons. No comparison survives correction.

6 Discussion

Matching 565 sustainability reports from 335 Korean firms against the government-certified emissions registry, we find a communication regime that is at once scrupulously accurate and systematically uninformative. Verifiable statements are correct to a fraction of a percent, and every whole-firm reduction claim that could be checked survives the check. But such claims are seven percent of reduction rhetoric; the rest consists of activity effects, intensity ratios, projects, and targets that no registry can contradict, produced at the same rate whether certified emissions rose or fell. Disclosure quality tracks organizational size, not carbon exposure, and the market prices none of it detectably. This section draws out what these findings imply for theory, for the ethics of corporate communication, and for regulatory design.

6.1 Theoretical contributions

These results carry three implications for theory. For the cheap-talk literature, they supply a mechanism: vagueness is not a measurement problem—firms hold audited numbers and cite them accurately when they cite them at all—but a portfolio choice over falsifiability, the

equilibrium of an enforcement environment that prices only falsifiable falsehoods (Bingler et al., 2022; Crawford and Sobel, 1982; Lyon and Maxwell, 2011). For the selective-disclosure literature, they extend the locus of selection one level deeper than the choice of facts (Marquis et al., 2016; Kim and Lyon, 2015): selection over the logical form of assertions suppresses nothing and is therefore harder to police. For research on disclosure determinants, the null on emission intensity resists both the signaling prediction that superior performers separate through verifiable disclosure (Verrecchia, 1983; Clarkson et al., 2008) and the legitimacy prediction that exposed firms disclose under pressure (Patten, 1992); with no market demand for the information (H4), disclosure quality defaults to a by-product of administrative capacity. Finally, the evidence gives the impression-management and decoupling literatures a quantitative anchor: the organizational façade of Cho et al. (2015) is built, in our data, not from false statements but from grammatically unfalsifiable ones, and its dominant form can be counted.

6.2 Ethical implications: deception without lying

The ethical character of the conduct we document is distinctive: deception without lying. A sentence of the form “we reduced emissions by X tons” does not state, but reliably conveys, that emissions fell (Grice, 1975). In Habermasian terms—whose validity claims of truth, sincerity, appropriateness, and understandability underpin established measures of CSR report credibility (Lock and Seele, 2017)—such a sentence honors the truth claim while defeating sincerity: nothing asserted is false, yet the assertion is selected to convey what the speaker knows the record would contradict. Because the conveyed proposition is never asserted, it is never false, and no truth-in-advertising regime reaches it. Ethical and regulatory assessment of corporate communication should therefore attend to the distribution of claim types, not only to the truth of individual claims—the same move financial reporting made when it stopped litigating narratives and mandated formats.

6.3 Implications for policy and practice

The policy implication follows directly from the empirical structure. Enforcement premised on detecting falsehoods will find little to act on here; the binding constraint is the supply of falsifiable statements. The remedy is to mandate the formats that make claims checkable—domestic legal-entity totals, location-based Scope 2 accounting, year-on-year figures, and an explicit reconciliation between the public report and the firm’s certified registry entry. The last item is nearly free where, as in Korea, every covered firm already holds a certified number. The suggestive market response to the 2026 roadmap among verifiable-format firms

hints that such mandates may even be rewarded rather than merely borne.

6.4 Limitations and future research

Five limitations bound these claims. The registry covers domestic legal entities while reports often consolidate globally, so “no contradictions” means none among scope-matched statements—though the boundary choice is itself part of what we measure. The verified stock-claim base is small (six firm-years); our strong claim is the scarcity of falsifiable statements, not a proven absence of false ones. One fifth of claim sentences remain fragmentary extraction residuals, and classifier agreement, while far above rule-based baselines, is imperfect. The setting is one country and a self-selected population of report publishers. And the determinants and event-study evidence is associational; we avoid causal language throughout. The most direct extension is causal: the February 2026 roadmap, with its staged asset-threshold coverage, offers a difference-in-differences design in which verifiable-format adoption is the treatment and the full listed universe the control pool. A second is comparative: wherever mandatory registries coexist with voluntary reporting—the EU ETS being the obvious case—the claim-type composition and its verification can be replicated, turning our protocol into a portable audit instrument.

7 Conclusion

Corporate climate disclosure in Korea is accurate wherever it can be checked and constructed, almost everywhere else, so that it cannot be. That combination—not falsehood, but the systematic selection of unfalsifiable claim types—is what greenwashing looks like in a regime of voluntary reporting beside a mandatory registry. Where certified numbers already exist, the distance between what firms know and what they assert is not a measurement gap. It is a choice—and a regulable one.

References

- Avidan, M., Walls, J. L., Dowell, G. W., and Cauderay, V. (2026). Temporal (in)consistency in sustainability disclosure. *Journal of Business Ethics*. Forthcoming, doi:10.1007/s10551-025-06236-x.
- Bingler, J. A., Kraus, M., Leippold, M., and Webersinke, N. (2022). Cheap talk and cherry-picking: What ClimateBert has to say on corporate climate risk disclosures. *Finance Research Letters*, 47:102776.

- Bingler, J. A., Kraus, M., Leippold, M., and Webersinke, N. (2024). How cheap talk in climate disclosures relates to climate initiatives, corporate emissions, and reputation risk. *Journal of Banking and Finance*, 164:107191.
- Bromley, P. and Powell, W. W. (2012). From smoke and mirrors to walking the talk: Decoupling in the contemporary world. *Academy of Management Annals*, 6(1):483–530.
- Calamai, T., Balalau, O., Le Guenedal, T., and Suchanek, F. M. (2025). Detecting greenwashing: A natural language processing literature survey. arXiv preprint arXiv:2502.07541.
- Cho, C. H., Laine, M., Roberts, R. W., and Rodrigue, M. (2015). Organized hypocrisy, organizational façades, and sustainability reporting. *Accounting, Organizations and Society*, 40:78–94.
- Clarkson, P. M., Li, Y., Richardson, G. D., and Vasvari, F. P. (2008). Revisiting the relation between environmental performance and environmental disclosure: An empirical analysis. *Accounting, Organizations and Society*, 33(4–5):303–327.
- Comyns, B. and Figge, F. (2015). Greenhouse gas reporting quality in the oil and gas industry: A longitudinal study using the typology of “search”, “experience” and “credence” information. *Accounting, Auditing & Accountability Journal*, 28(3):403–433.
- Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.
- Delmas, M. A. and Burbano, V. C. (2011). The drivers of greenwashing. *California Management Review*, 54(1):64–87.
- Depoers, F., Jeanjean, T., and Jérôme, T. (2016). Voluntary disclosure of greenhouse gas emissions: Contrasting the Carbon Disclosure Project and corporate reports. *Journal of Business Ethics*, 134(3):445–461.
- Grice, H. P. (1975). Logic and conversation. In Cole, P. and Morgan, J. L., editors, *Syntax and Semantics, Vol. 3: Speech Acts*, pages 41–58. Academic Press, New York.
- Hahn, R. and Lülfs, R. (2014). Legitimizing negative aspects in GRI-oriented sustainability reporting: A qualitative analysis of corporate disclosure strategies. *Journal of Business Ethics*, 123(3):401–420.
- Han, H. H., Lee, J., and Wang, B. (2023). Greenhouse gas emissions, firm value, and the investor base: Evidence from Korea. *Emerging Markets Review*, 56:101048.

- International Carbon Action Partnership (2025). Korea emissions trading system (K-ETS). ETS Factsheet, ICAP. <https://icapcarbonaction.com/en/ets/korea-emissions-trading-system-k-ets>.
- Kim, E.-H. and Lyon, T. P. (2015). Greenwash vs. brownwash: Exaggeration and undue modesty in corporate sustainability disclosure. *Organization Science*, 26(3):705–723.
- Kim, S. A. and Kim, J. D. (2022). Voluntary carbon disclosure (VCD) strategy under the Korean ETS: With the interaction among carbon performance, foreign sales ratio and media visibility. *International Journal of Environmental Research and Public Health*, 19.
- Lee, J., Park, H., and Pham, M. K. (2024). The effects of carbon emissions trading on profitability and value: Evidence from Korean listed firms. *Journal of International Financial Management & Accounting*, 35(3):760–799.
- Lee, J.-H. and Cho, J.-H. (2021). Firm-value effects of carbon emissions and carbon disclosures—evidence from Korea. *International Journal of Environmental Research and Public Health*, 18(22):12166.
- Lock, I. and Seele, P. (2017). Measuring credibility perceptions in CSR communication: A scale development to test readers’ perceived credibility of CSR reports. *Management Communication Quarterly*, 31(4):584–613.
- Lyon, T. P. and Maxwell, J. W. (2011). Greenwash: Corporate environmental disclosure under threat of audit. *Journal of Economics & Management Strategy*, 20(1):3–41.
- Lyon, T. P. and Montgomery, A. W. (2015). The means and end of greenwash. *Organization & Environment*, 28(2):223–249.
- Marquis, C., Toffel, M. W., and Zhou, Y. (2016). Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, 27(2):483–504.
- Oh, N., Miteva, D. A., and Lee, Y. (2023). Impact of Korea’s emissions trading scheme on publicly traded firms. *PLOS ONE*, 18(5):e0285863.
- Patten, D. M. (1992). Intra-industry environmental disclosures in response to the Alaskan oil spill: A note on legitimacy theory. *Accounting, Organizations and Society*, 17(5):471–475.
- Seele, P. and Gatti, L. (2017). Greenwashing revisited: In search of a typology and accusation-based definition incorporating legitimacy strategies. *Business Strategy and the Environment*, 26(2):239–252.

- Talbot, D. and Boiral, O. (2018). GHG reporting and impression management: An assessment of sustainability reports from the energy sector. *Journal of Business Ethics*, 147(2):367–383.
- Verrecchia, R. E. (1983). Discretionary disclosure. *Journal of Accounting and Economics*, 5:179–194.
- Webersinke, N., Kraus, M., Bingler, J. A., and Leippold, M. (2021). ClimateBert: A pre-trained language model for climate-related text. arXiv preprint arXiv:2110.12010.