

ESG Beyond Disclosure: The Signalling Boundary and the Cumulative Development of Sustainable Finance Theory

¹CS Monika Agarwal*, ²Dr. Arpan Parashar, ³Dr. Manisha Semwal, ⁴CMA Ganesh Narwade, ⁵Dr. Shruthi Joshi,

.Department of commerce and management, Banasthali vidhyapeeth ,Rajasthan, India,

*Corresponding author. Email: csmonikagarwal@gmail.com

<https://orcid.org/0009-0001-8654-7407>

Department of commerce and management, Banasthali vidhyapeeth ,Rajasthan, India,

Email: arpanparashar@banasthali.in

Department of commerce and management, Graphic Era Deemed to be University, Dehradun, Uttarakhand, India,

Email: manishasemwal03@gmail.com

Department of commerce and management , ICBM School of Business Excellence, Hyderabad, Telangana India,

Email: narwade.ganesh@gmail.com

Department of Professional Studies, School of Commerce, Finance and Accountancy,

CHRIST (Deemed to be University) Bangalore, India.

Email: shruthijoshi.13@gmail.com

Abstract

ESG disclosure has expanded from voluntary narratives into mandatory, assurance-bearing regimes, but disclosure volume does not constitute investor-relevant information — a tension this paper identifies as the ESG Disclosure Paradox. We introduce the ESG Signalling Boundary, a diagnostic framework grounded in signalling theory and cheap-talk models of strategic communication, built on three principles — signal cost and verifiability, receiver decodability, and feedback and accountability combined by a weakest-link rule: informativeness is set by a disclosure's least-satisfied principle, not their sum. The rule yields falsifiable predictions, including an inverted-U relationship between verification cost and informativeness. An illustration from India's mandatory BRSR Core assurance cycle shows the three principles diverging within one disclosure population; illustrations from generative AI and retail-investor markets show identical disclosures on opposite sides of the boundary depending on the receiver. The framework shifts evaluation from disclosure volume to more credibility, providing a basis for cumulative sustainable finance theory.

Keywords: ESG disclosure, cheap talk, signalling boundary, sustainable finance theory, investor value relevance, BRSR Core, cumulative theory development

Introduction: From Disclosure Expansion to the ESG Disclosure Paradox Corporate accountability was once owed to shareholders alone. Freeman (1984) shifted that premise toward a stakeholder-inclusive view, arguing firms answer to a much wider set of constituents, and that shift did real work: it helped establish environmental, social, and governance conduct as information capital markets and society hold a legitimate claim to. Four decades of disclosure infrastructure followed: sustainability guidelines, integrated reporting, climate-related frameworks, and now mandatory regimes including the EU's Corporate Sustainability Reporting Directive (CSRD), the IFRS Foundation's International Sustainability Standards Board (ISSB) standards, and, in India, SEBI's Business Responsibility and Sustainability Report with its assurance-bearing BRSR Core (Securities and Exchange Board of India [SEBI], 2023).

That expansion's success has, in turn, posed a problem nobody quite intended. Disclosure has moved well past audited financial statements into metrics, narratives, targets, and third-party ratings covering emissions, labour practices, board composition, and supply-chain conduct (Christensen et al., 2021), and each new item is justified by the same logic that motivated the original shift: more disclosure should let capital allocators price sustainability-related risk and opportunity more accurately.

Volume itself is not the problem. Reporting systems have always increased in number through the introduction of indicators, and an extra indicator is not necessarily a bad thing (Verrecchia, 2001). The problem begins when an ESG metric, target, or rating, due to its nature, is considered to be informative even if it isn't priced. A disclosure may be a material disclosure without being reliable or cheap to make up; it may originate from an effective performing company and not arrive to investors who have no means to decipher it in opposition to a materiality framework; it can meet with a regulator's checklist without resulting in any consequences for being incorrect. None of all that means disclosure volume is not unimportant, it just means that volume is not the only part of disclosure that is important.

Solving that requires a change of focus: from what they disclose, to how informative that disclosure is. But this paper puts the ESG Signalling Boundary into practice by assessing disclosure based on the conditions in which it can serve as a credible signal, rather than on the question of whether a disclosure of an ESG-related claim, metric, or rating is present.

In this article three contributions are made. First, it develops the ESG Signalling Boundary, a diagnostic framework grounded in signalling theory (Spence, 1973; Connelly et al., 2011) and in the formal game-theoretic treatment of costless communication (Crawford & Sobel, 1982), situated against Verrecchia's (2001) disclosure-theory typology and a broader systematic review of the theoretical lenses used in ESG disclosure research (Del Gesso & Lodhi, 2025); the three principles combine through an explicit weakest-link rule, yielding falsifiable propositions rather than a descriptive checklist. Second, it positions the framework against a fast-growing 2025–26 empirical literature that independently tests boundary conditions on the ESG disclosure–value relationship, including signalling decoupling (Zatini et al., 2026), AI-enabled greenwashing risk (Yadav & Asongu, 2026), and sin-stock-sector conditions on ESG value creation (Kim & Oh, 2026). Third, drawing on the first full assurance cycle under India's mandatory BRSR Core regime, it shows that the three diagnostic principles can genuinely diverge rather than simply move together, which answers a natural objection to any multi-principle framework: that the principles might just be one thing wearing three names.

The ESG Disclosure Paradox: An Unresolved Challenge in Sustainable Finance Theory

Empirical evidence keeps arriving at the same uncomfortable finding: disclosure expansion does not reliably buy proportional gains in investor-relevant information. Khan et al. (2016) find that firm performance on industry-specific, financially material sustainability issues predicts future stock and accounting performance, while performance on immaterial issues does not. Christensen et al. (2022) show that as disclosure expands, disagreement among ESG rating agencies assessing the same firm rises rather than falls, and Berg et al. (2022) document only moderate correlation among major providers rating the same companies — hardly what one would expect if disclosure resolved into a single shared signal. Thewissen et al. (2026) push this point down to the level of content itself: analysing over 7,500 ESG reports, they find informativeness (proxied by analyst-forecast agreement) depends not just on whether firms disclose ESG information but on what they disclose, with topic diversity showing diminishing, not increasing, returns past a moderate point. The same parallel greenwashing literature has a similar story to tell. What is commonly observed, however, is that firms subjected to the highest level of reputational risk tend to do more symbolic disclosure instead (Delmas & Burbano, 2011), and Yu et al. (2020) show that particular governance devices across countries (independent directors and institutional-investor scrutiny) are effective at slowing down greenwashing, while disclosure rules are not.

The same tension shows up directly in studies asking whether ESG disclosure raises firm value at all. A meta-analysis of 52 Indonesian studies finds no significant aggregate relationship between ESG disclosure and firm value (Prayogo et al., 2026). Zhang et al. (2026), looking at Chinese A-share firms, find that disclosure quantity's marginal effect on firm value is not fixed but dynamic: negative when underlying ESG performance is low, and

turning positive only once performance crosses a threshold. Kim et al. (2026) report a related mechanism, an inverted U-shaped ESG–value relationship driven by interpretive uncertainty that firms can only partly counteract through sensegiving activity. Three markets, three research designs, one recurring pattern: ESG disclosure's relationship to investor value is conditional rather than automatic. That is where this paper starts.

We refer to this as the ESG Disclosure Paradox: the more broadly a disclosure regime is cast as covering which area of a company a disclosure must cover, the more it appears it has the power to inform capital markets, and the more difficult it is to say which disclosures provide the information. Needless to say, it's not a reason not to expand disclosure. The process of materiality mapping, assurance standards and rating methodologies are some empirical steps toward theoretical advances in defining what sustainability information is decision useful (Khan et al., 2016; Christensen et al., 2021).

What remains unresolved is a harder, more specific question: when does an expanding regime produce genuinely new investor-relevant information, and when does it merely add volume the market has no way to price?

This distinction matters increasingly as disclosure engages generative AI, mandatory assurance regimes, and market contexts with very different investor populations. A disclosure could be produced by an AI system that has been trained to produce fluent sustainability narratives, is compliant with the minimum disclosure requirement of a jurisdiction, and is produced using a materiality framework that is different from that of the retail investor, but for which the disclosure appears credible for none of the reasons that make disclosure informative. Addressing this means shifting attention from where ESG disclosure is required to what makes it function as a signal.

This paradox is not solely a conceptual concern. As the sections that follow develop the Signalling Boundary's three diagnostic principles, we return repeatedly to a wider set of 2025–26 empirical papers that independently document facets of the same paradox: decoupling between rhetoric and routine, uneven receiver processing, boundary conditions on when disclosure creates value, and AI-enabled erosion of credibility. Taken together, they offer convergent, if fragmented, evidence for a problem this paper treats as one diagnostic logic rather than four separate ones.

The ESG Signalling Boundary: An Adaptive Principle of Disclosure Evaluation

The ESG Signalling Boundary is aimed squarely at the evaluative problem disclosure expansion has created: whether a given disclosure functions as a credible signal of underlying firm conduct, or merely satisfies a reporting obligation without informing capital allocation at all. This distinction matters for cumulative theory-building in sustainable finance, because progress in this field depends on more than novelty in new metrics or new taxonomies; it depends on whether disclosure carries information a receiver can actually act on (Whetten, 1989; Connelly et al., 2011).

The ESG Signalling Boundary is a diagnostic tool to examine the effectiveness of an ESG disclosure as an independent signal versus noise, regardless of the underlying performance or disclosure objectives of the company making the disclosure. It does not specify what topics companies must report, but does not set fixed reporting requirements. Rather, it focuses on disclosure based on the conditions under which it could communicate with the capital markets.

One clarification is worth making explicit here. Receiver decodability is, by construction, a property of a specific receiver population rather than of the disclosure itself, so the boundary is more precisely a property of a disclosure–receiver pair: two different receivers can place the identical disclosure on opposite sides of it, as the retail-versus-institutional illustration below shows directly. For brevity we still speak, as convention allows, of a disclosure being "inside" or "outside" the boundary; read this throughout as shorthand for inside or outside the boundary for a specified receiver population. Cost and accountability do not carry this same relativity: they are properties of the disclosure and its regulatory environment rather than of any one receiver, which is why only the decodability principle below is framed in explicitly receiver-relative terms.

Choosing signalling theory as the organising lens here is deliberate, not a default choice. Del Gesso and Lodhi's (2025) systematic review of 142 accounting studies on ESG and CSR disclosure finds that signalling theory is only the fifth most frequently applied lens in this literature, behind stakeholder, legitimacy, institutional, and

agency theory. Fu et al. (2022) show what a more developed signalling account can add even within a narrower question — a curvilinear, cost-driven relationship between ESG signalling breadth and equity risk — but their focus stays empirical rather than evaluative. Uyar et al. (2020) ask directly whether CSR reporting signals or greenwashes, but resolve it only within one sector, the worldwide logistics industry; the Signalling Boundary offers ex-ante criteria that could answer the same question for any ESG disclosure, in any sector. Asif and Searcy (2026) pursue a related deterrence goal but draw on several foundational management theories rather than committing to one organising lens, trading mechanism-precision for broader institutional coverage.

It helps to place the Signalling Boundary against Verrecchia's (2001) influential typology of the disclosure literature. Verrecchia distinguishes discretionary-based disclosure (why managers choose to disclose, holding information endowment fixed), association-based disclosure (how disclosure is associated with investor behaviour, without modelling the disclosure decision itself), and efficiency-based disclosure (which arrangements are unconditionally best for a class of information problems, before any firm's information is even realised). ESG disclosure research in sustainable finance has drawn heavily on the first tradition and on association-based work linking disclosure to valuation and rating outcomes (Khan et al., 2016; Christensen et al., 2022). What the Signalling Boundary adds is a diagnostic layer neither tradition supplies: it does not model why a firm discloses or measure the statistical association between disclosure and outcome, but specifies the conditions, cost, decodability, accountability, under which a disclosure of any origin becomes capable of generating that association at all. In Verrecchia's own terms, the question is logically prior to association-based measurement: not whether disclosure and price move together, but under what conditions they should be expected to.

Conventional materiality treats importance as a fixed classification of topics. The Signalling Boundary does not work that way: rather than drawing a permanent line between material and immaterial disclosure, it locates an adaptive boundary that shifts with the cost of producing disclosure, the capacity of the receiving market, and the mechanisms available to discipline misrepresentation. It is not a checklist or a list of topics, but an interface for assessing and evaluating, and the cost, decodability and accountability all play a role in where any particular disclosure falls.

The result is a change in the focus of the evaluation from disclosure presence to signalling mechanism. But the question is, does a disclosure have to be expensive enough to be believable, decipherable enough to be understood, and verifiable enough to be trusted if it covers a topic that is labelled as ESG? While identification of a reported measure is good theory, a theory must also articulate the mechanism by which the reported measure can or cannot tell a capital-market receiver (Whetten, 1989). This commitment is based on the three diagnostic principles discussed next: signal cost and verifiability, receiver decodability, and feedback and accountability.

Operationalizing the ESG Signalling Boundary: Three Diagnostic Principles

A framework is not a really meaningful one until it becomes a set of criteria by which one could determine whether a disclosure is a credible signal. The Signalling Boundary is three interconnected principles, each of which answers a different question, in turn: why a disclosure is believed, whether the receiver of that disclosure can understand it, and whether there are any consequences for having it misrepresented.

This is not a random division, it comes from the five elements Connelly et al. (2011) have established as being necessary to any signalling episode: signaler, signal itself (cost and fit), receiver (attention and interpretation), feedback, and signalling environment (countersignals). The three principles below arise naturally within the context of a classification of informational failure where it can actually happen. signaler-side conditions motivate cost and verifiability, receiver-side conditions motivate decodability, and environmental and dynamic conditions motivate feedback and accountability, each corresponding to a distinct locus of failure in the foundational review.

Signal Cost and Verifiability

The first diagnostic principle is one of the grounds on which it can be accepted as true. Spence (1973) has answered that question in the following way. The signalling theory suggests that if an informed party is going to be able to credibly convey private information to an uninformed party, then this must be brought about at an irrationally high cost if done dishonestly. If not, it is sent by both high and low performing firms and contains no information

(Connelly, 2011). Exactly this asymmetric cost is imposed by assuring and audited emissions data, but not by unaudited narrative text, which can be reproduced at near zero marginal cost. Fu et al. (2022) find that signal cost accelerates as ESG scope broadens — a curvilinear, U-shaped relationship that scales with the breadth and specificity of what is being signalled. Reber et al. (2022) find something similar from a different angle: voluntary ESG disclosure and higher ratings predict lower idiosyncratic volatility and downside risk in a firm's first year after IPO, precisely when investors have the least independent basis for evaluation. Debt markets price the same logic even more directly. Taddeo et al. (2026), across 411 S&P 500 firms, show that the gap between a firm's ESG disclosure and its actual performance greenwashing, in their terms gets priced into a persistently higher after-tax cost of debt. That's contrary to what Zatini et al. (2026) find, which is described below, that equity-oriented signalling-decoupling measures are not penalized but instead are rewarded. Whether to value low verifiability as risk or reward is a choice that may be made on the basis of the type of receiver; the latter is directly addressed by the second principle.

This cost logic, however, can't go on like this forever. The original separating equilibrium of Spence (1973) assumes that the cost of the signal cannot be too low or the low signaler will be able to mimic and the cost cannot be too high or the high signaler will be prevented from signalling truthfully. Once it exceeds that level, push verification costs and even highly successful companies can find themselves on the verge of irrational compliance, settling into a "pooling" equilibrium at the mandatory minimum. The curvilinear relationship between signal cost and equity risk, which Fu et al. (2022) find, is a good fit for this bound. Our proposal, accordingly, is that net informativeness is inverted-U in verification cost: it rises with cost up to an optimum, and past that point, further increases in mandated verification cost push compliant firms toward disclosing only what is strictly required, so marginal informativeness falls even while formal compliance stays universal. This is directly testable using the same cohort-phasing structure behind the BRSR Core illustration below. SEBI's staged rollout by market capitalisation created cohorts facing different compliance timelines and cost burdens; the inverted-U prediction says informativeness gains from assurance should not scale uniformly with cohort-level cost.

This kind of account has its formal counterpart that is worth considering: cheap-talk models of strategic communication, which pose the question of whether a message can be informative with no cost, and be credible even if it is unverifiable (Crawford & Sobel, 1982). Crawford and Sobel demonstrate that the communicative value of a message is only partial and only when sender and receiver interests coincide, even when the message is free of cost. As preferences change, the best informative balance tips over to pure noise — a babbling balance where the sender tells the receiver no more than what he or she was expecting to hear. Here, cost is not valuable per se; rather, it serves as a proxy for the deeper condition for informativeness, the sender-receiver alignment, as cheap-talk identifies. There is generally no perfect alignment between a firm and its investors regarding ESG disclosure: while firms have incentives to babble with regard to ESG performance in order to reap a reputational benefit, investors do not have the same incentives, as costly separation would eliminate the benefits of babbling in that scenario. In short, verifiable cost takes the place of voluntary unaudited narrative disclosure which is unable to guarantee an alignment in its own right.

The question this first principle puts to any disclosure, then, is simple to state: would its cost structure deter a low-performing firm from producing it dishonestly? Many disclosures satisfy a reporting requirement without imposing any such cost. A regulation may compel a stated target. A rating methodology may reward a policy document's mere presence. A voluntary framework may accept a narrative claim at face value. These disclosures can carry real market consequences, but they differ, in their informational foundations, from one that would actually be expensive for a dishonest signaler to fabricate. Hoai's (2026) survey evidence on the ESG Assurance Expectation Gap is useful here: it distinguishes a reasonable gap in what assurance can ever guarantee from a deficient-standard gap, where the criteria themselves are too weak, and a deficient-performance gap, where the process exists but is poorly executed, and finds it is specifically the deficient-performance dimension that erodes stakeholder trust most severely. The diagnostic question is not whether a disclosure exists, or carries the label of assurance, but whether its cost and verification rigour jointly separate genuine claims from fabricated ones.

Diagnostic Question: *Does the disclosure impose a cost of production that a firm without the underlying performance would find irrational to pay?*

Receiver Decodability

The second principle concerns the interpretive structure through which a signal is actually received. Even a costly, honest signal fails to establish investor-relevant information if the receiver cannot decode it (Connelly et al., 2011). Materiality scholarship shows value relevance depends on whether disclosed information maps onto industry-specific financial mechanisms a receiver recognises (Khan et al., 2016). Institutional-investor survey evidence points to the same constraint from the practitioner side: the absence of standardised, comparable ESG reporting is a primary impediment to using such information at all (Amel-Zadeh & Serafeim, 2018), suggesting this receiver-side limitation persists even when disclosure is costly and perfectly honest.

This kind of account has its formal counterpart that is worth considering: cheap-talk models of strategic communication, which pose the question of whether a message can be informative with no cost, and be credible even if it is unverifiable (Crawford & Sobel, 1982). Crawford and Sobel demonstrate that the communicative value of a message is only partial and only when sender and receiver interests coincide, even when the message is free of cost. As preferences change, the best informative balance tips over to pure noise — a babbling balance where the sender tells the receiver no more than what he or she was expecting to hear. Here, cost is not valuable per se; rather, it serves as a proxy for the deeper condition for informativeness, the sender-receiver alignment, as cheap-talk identifies. There is generally no perfect alignment between a firm and its investors regarding ESG disclosure: while firms have incentives to babble with regard to ESG performance in order to reap a reputational benefit, investors do not have the same incentives, as costly separation would eliminate the benefits of babbling in that scenario. In short, verifiable cost takes the place of voluntary unaudited narrative disclosure which is unable to guarantee an alignment in its own right.

So the question this second principle asks is not whether a disclosure exists, but whether it explains an outcome through a receiver's actual capacity to interpret it. The same disclosure can sit on different sides of the boundary depending purely on who receives it. An institutional analyst with sector-specific materiality maps may decode a disclosure that a retail investor, lacking any comparable frame of reference, cannot use at all, a divergence that becomes acute in retail-dominated markets. Whether a disclosure is material in principle is beside the point; what matters is whether its receiver population is actually positioned to treat it as such.

Diagnostic Question: *Does the disclosure explain investor outcomes through a receiver's demonstrated capacity to attend to and interpret it against a materiality framework?*

Feedback and Accountability

The third is consequence. Informativeness is not solely determined by cost and decodability, as both can be subtly diminished over time without detection and punishment if they are not used properly (Connelly et al., 2011). What accountability adds is a feedback loop of scrutiny — from regulatory action, court action, analyst re-thinking, or continuous media and civil society monitoring — in the form of subsequent disclosure and pricing behaviour (Marquis et al., 2016), which creates an endogenous mechanism of market discipline that extends beyond the immediate disclosure act. This principle was brought to life in a way for a cross-country study on greenwashing by Yu et al. (2020). Independent-director monitoring and institutional-investor monitoring have a pronounced effect on reducing greenwashing without the effect of firm-level governance, and the effect of country-level anti-corruption enforcement and cross-listing status is only significant in combination with firm-level governance. Accountability, that is, is seldom a single lever.

What this third principle asks, then, is whether a disclosure regime builds in feedback as a genuine accountability mechanism, rather than simply trusting the signaler's good faith. This is a different question from generic regulatory compliance: firms may well expand disclosure out of reputational pressure or efficiency objectives without any feedback loop capable of catching a false claim (Delmas & Burbano, 2011). Accountability only becomes a signalling mechanism once scrutiny can actually change the cost calculus of the next disclosure. As the illustrative application below shows, this feedback condition can lag well behind cost and decodability, even within a single regulatory cycle; it does not resolve automatically just because a mandate exists.

Diagnostic Question: *Does a mechanism exist that can detect and penalise misrepresentation, changing the cost of the next disclosure?*

The Signalling Boundary Amid a Fast-Growing 2025–26 Literature on ESG Signalling Boundary Conditions

The three diagnostic principles above address a conceptual gap that recent empirical scholarship has started approaching from several directions at once, without quite naming it as such. This manuscript was drafted alongside an identifiable cluster of studies, several of them in Business Strategy and the Environment, each testing a specific boundary condition on the ESG disclosure–value relationship without proposing the diagnostic logic that would tie them together. Organising this cluster by the mechanism each paper isolates shows how timely the problem is, and what an ex-ante framework can add alongside ex-post measurement.

A first set of studies examines directly whether ESG disclosure generates value, instead of taking that outcome for granted. Agrawal (2026), using CRISIL's SEBI-aligned ratings for 380 Indian firms, finds that environmental and governance practices produce value only when backed by firm size, age, profitability, and recognition, whereas CSR itself remains largely symbolic. Kim and Oh (2026) draw a comparable conclusion through a different channel: the connection between ESG disclosure and performance declines steeply in sin-stock sectors. Surya et al. (2026), and separately Tran et al. (2026), verify the same overall pattern for Indonesian and Southeast Asian firms, this time through political connections and financial distress. Four studies, four mechanisms, and one shared takeaway: the ESG–value relationship is contingent rather than consistent. Yet none of them explains why.

A second line of research considers what can substitute for formal, third-party verification when credibility still must be established. Khan et al. (2026) show, for Turkish firms, that business-group affiliation and institutional blockholder monitoring enhance ESG valuation by standing in for underdeveloped assurance markets. Zatini et al. (2026) address the same issue from the other side by distinguishing costly routinisation from environmental rhetoric that is not matched across 593 large-cap firms, and they find that markets often reward the rhetorical pattern rather than penalise it. Taken together, the two studies complement each other: one identifies a substitute that strengthens a weak regime, while the other documents what occurs when it is absent. Table 2 (row 7) rates this arrangement as only partially within the boundary under the weakest-link rule, even in settings where an informal substitute is at work—an interpretation offered here without fully resolving it.

A third cluster looks past institutional analysts to how other kinds of receivers process ESG signals. Tan et al. (2026) survey 415 Chinese consumers and find that episodic ESG signal quality is a more influential purchase-intention cue than a firm's enduring reputation — an echo of Kim et al.'s (2026) interpretive-uncertainty effect discussed above. Neither paper addresses this paper's investor-receiver question directly, but both confirm something the second principle already assumes: timing, framing, and interpretation shape a signal's effect independently of the sender's track record.

One terminology note is worth adding here. Lee et al. (2026) use the same phrase, "ESG paradox" — also in Business Strategy and the Environment — but for something different: a firm-level choice between innovation and washing under stakeholder-legitimacy pressure. That sits one step upstream of the receiver-side problem this paper addresses, and the two accounts are compatible without restating each other — their explanation of why firms produce washing-type disclosure plausibly describes exactly the disclosures the Signalling Boundary would classify as outside it.

Generative AI supplies a fifth thread. Yadav and Asongu (2026), across 1,260 firm-years in six Asian economies, find that AI adoption in sustainability reporting increases greenwashing intensity unless constrained by verification and regulatory stringency. Qiu et al. (2026) find close to the opposite: an AI-based readability metric applied to Chinese CSR reports is negatively associated with ESG rating divergence. The difference comes down to what the technology is doing — production versus verification — and that distinction maps directly onto the Signalling Boundary's own separation of signal cost from receiver decodability.

Put these five clusters together and a pattern emerges: signalling-based accounts of ESG disclosure are converging, from independent directions, on cost, receiver characteristics, and verification, exactly the loci this paper's three principles derive from. What each of these studies contributes is measurement. They isolate one mechanism and test its empirical footprint, while the Signalling Boundary tries to integrate all three into a single framework usable *ex ante*. The relationship between the two is cumulative, not competitive: a framework becomes more useful when it anticipates the patterns such studies document, and a finding becomes more interpretable when read against an explicit diagnostic logic. This literature is moving fast enough that this section can only reflect the evidence available as the manuscript was drafted; the framework's job is to organise a fast-moving evidence base, not to have the last word on it.

Read this way, these moderators start to look less like independent discoveries and more like differently named proxies for the same three underlying primitives. Table 3 recodes each moderator discussed above onto the diagnostic principle it actually proxies, and states what the Signalling Boundary predicts that the original study — framed, understandably, around its own moderator — never tested.

The Signalling Boundary as an Integrated Diagnostic Framework

The Signalling Boundary does not treat its three diagnostic principles as independent conditions, and it is not a checklist in which satisfying one criterion is enough to establish investor relevance. What we propose instead is a weakest-link rule: a disclosure's net informativeness is set by whichever of the three principles is least satisfied, not by their sum or their average. Score cost and verifiability, decodability, and accountability on the same ordinal scale, and informativeness equals the minimum of the three scores — never a weighted or additive combination. Signal cost and verifiability address whether a disclosure can be believed. Receiver decodability addresses whether it can be understood. Feedback and accountability address whether misrepresentation carries any consequence. Fail any one of these three questions, and a disclosure is uninformative on that dimension, no matter how strongly it satisfies the other two.

This choice of aggregation rule is not a minor technical detail; it is what lets the framework claim to be a theory rather than a classification scheme. Consider the alternative. A compensatory or additive rule, in which exceptional cost could offset weak accountability, would imply that a maximally expensive, fully audited disclosure with no feedback mechanism could still register as informative overall. There is no theory in signalling to justify that trade-off: a receiver who can't hold a costly claim responsible will have no way to distinguish a one-time costly signal from a costly claim that goes on to be unpoliced. The weakest-link rule has a precise testable corollary: marginal investments in the principle that is strongest are worth no marginal gain in informativeness. Those wishing to strengthen a disclosure within the boundary should focus on the binding constraint, rather than on any particular principle that might be the most economical to strengthen—a principle that, in fact, can be shown to be wrong not only on this or that example of it but in its essence: if the empirical evidence indicates that a stronger version of an already-strong principle leads to a higher rating of informativeness net of the weaker principle's rating, then the weakest-link principle is mistaken, not only that particular example of it, but the principle itself. Table 2 mechanically applies the rule to each phenomenon that it is applied to, in accordance with the scoring of the principles themselves; a three value verdict scale is directly derived from the principles' own scoring: "inside the boundary" when the weakest principle is fully satisfied; "partially inside the boundary" when it is only partially satisfied; "outside the boundary" when it is unsatisfied. Each verdict names the binding constraint the principle, or principles, actually doing the work including cases where the verdict is not the one a purely regulatory or reputational account would predict.

Take a fully assured, audited emissions figure that a retail investor simply has no materiality framework to interpret. It is costly and honest, and yet under the weakest-link rule it is uninformative to that particular receiver, regardless of how much it cost to produce. The same logic runs the other way: a decodable disclosure with no accountability mechanism registers as uninformative on the accountability dimension even when cost and decodability are both strong, and an accountable regime built on cheap, unverifiable input has nothing substantive left to hold firms accountable for in the first place. Table 1, placed at the end of the manuscript per journal

formatting requirements, summarises the three diagnostic principles and their logic; Table 2 puts them to work. Figure 1, also placed at the end per journal formatting, diagrams this aggregation logic end to end.

A natural objection to any multi-principle framework is that the principles might simply co-vary with one underlying driver, most obviously regulatory stringency, rather than being separable. This has real force where a single regulatory shock moves an entire population at once: a jurisdiction-wide assurance mandate raises verification cost and, by creating an auditable record, plausibly strengthens accountability for every covered firm simultaneously. However as demonstrated in the following illustrative application this co-variance does not apply generally. For firms that choose to voluntarily provide assurance beyond the core, within the same population under the same SEBI requirement, there is no corresponding increase in the level of decodability, as the level of decodability depends on the receiver's ability to distinguish these types and scopes of assurance, which raising the cost does not ensure. When one of the three is activated, then the three are activated as well, not the same, and exactly where they diverge provides a true indicator of the regime's remaining faults. None of this excludes regulatory, institutional, or technological mechanisms from the picture interdisciplinary integration can genuinely advance theory when mechanisms from other fields connect meaningfully to signalling outcomes. What the framework asks instead is whether the dominant explanatory mechanism actually advances our understanding of cost, interpretation, and accountability, which is ultimately what separates disclosure that merely carries an ESG label from disclosure shown, through these mechanisms, to inform capital allocation.




Figure 1. The Signalling Boundary's Integrated Diagnostic Logic

Illustrative Application of the Signalling Boundary

The Signalling Boundary is illustrated in three ways below which show how it evaluates various disclosure treatments. No one gives a definite judgment on a category. Each instead shows that the same phenomenon can carry a different signalling status depending on the mechanism through which cost, decodability, and accountability are established. The sequence goes from a technology development that is conceptually assessed (generative AI), to a technology instrument that is assessed based on actual outcomes (mandatory assurance) to a structural market condition (retail-investor-dominated markets).

Generative Artificial Intelligence. The history of disclosure has been the one area where fluent reporting on sustainability is both the most difficult and the least cost-effective to verify and copy. AI-assisted narrative becomes theoretically informative only when its claims remain tied to verifiable, costly commitments, such as audited metrics or capital expenditure, rather than to textual fluency itself. That fits signalling theory's emphasis on cost asymmetry over presentation quality (Spence, 1973; Connelly et al., 2011), and matches cross-country evidence that AI adoption in sustainability reporting increases greenwashing intensity unless constrained by verification and regulatory stringency (Yadav & Asongu, 2026). The same technology can also sit on the other side of the boundary when applied to verification rather than production. Qiu et al. (2026) find that an AI-based readability metric applied to existing disclosure is associated with lower ESG rating divergence, unless weakened by greenwashing itself. AI is therefore relevant not because it produces more disclosure, but because it can erode or reinforce signal cost depending on whether it is deployed on generation or verification. Table 2 (row 4) formalises the risk: AI-generated narrative can score higher on apparent decodability than human-written narrative without being any more costly or accountable. Fluency and informativeness are not the same thing.

Mandatory Assurance Regulation. Regulatory assurance mandates offer a concrete setting to evaluate the three diagnostic principles against real disclosure outcomes. India's SEBI has phased in mandatory reasonable assurance on defined BRSR Core indicators, introduced by circular in July 2023 and extended in stages to progressively larger cohorts by market capitalisation (SEBI, 2023). An analysis of the first full assurance cycle, covering 94 NIFTY100 companies' FY2024–25 filings, illustrates how the three principles apply, and diverge, within a single regulatory population (KPMG Assurance and Consulting Services, 2026). These figures come from a single, non-peer-reviewed source that does not state why these 94 constituents were included, nor whether counts were independently verified against filings, so the patterns below are best read as illustrative rather than independently audited findings. On signal cost and verifiability, the picture is unambiguous: all 94 companies obtained reasonable assurance — the costlier of the two options SEBI's 2024 amendment permits — on their mandatory Core indicators rather than the lower-cost "assessment" alternative, moving the covered indicators inside the boundary on cost by construction. Voluntary regimes look different. Surya et al. (2026) find that external ESG assurance does not significantly moderate the ESG–value relationship among Indonesian firms, attributing this to weak assurance differentiation in an immature market. SEBI's mandatory instrument, by contrast, appears to be doing signalling work a voluntary one may not.

The same population, however, shows that cost and decodability do not move together. Fifty-seven of the 94 companies voluntarily extended assurance beyond the mandatory Core scope, yet 36 of those 57 chose only limited assurance, a lower-cost, lower-verifiability standard. The frameworks covered (GRI, 41; supplementary BRSR, 41; climate frameworks such as SASB, TCFD, and GHG, 15) are not consistently distinguished from Core assurance in public-facing disclosure. A receiver must read the scope paragraph closely to learn that a headline claim blends reasonable assurance on nine indicators with limited or no assurance elsewhere, precisely the decodability failure principle two identifies; Hoai's (2026) deficient-standard/deficient-performance distinction offers a vocabulary for this blended claim. A portfolio-level reading is also worth noting here: firms choosing the cheapest voluntary tier just beyond a binding mandate may be substituting a minimum credibility label for genuine transparency at exactly the margin regulation does not reach. This waterbed effect is not something the Signalling Boundary can detect phenomenon-by-phenomenon, since signalling equilibria are properties of a sender's whole portfolio, not one disclosure; confirming it would require observing whether firms reduce effort on adjacent, still-voluntary margins as mandates tighten, an extension we note but do not test here. Accountability, over this same cycle, is similarly mixed. Only one company, in construction, received a qualified opinion, citing gaps in data, evidence, and internal controls a low failure rate that could reflect either genuine compliance or weak detection.

A timing pattern fits this picture: 30 of the 94 companies signed their assurance report more than fifty days after their audit opinion, with a maximum gap of 98 days. That could indicate weaker accountability cadence, though it is just as compatible with benign process integration (KPMG Assurance and Consulting Services, 2026). An assurance mandate is therefore not one boundary-crossing event: it can secure cost while leaving decodability and accountability only partially resolved. Table 4 formalises three of these observables as candidate proxies, alongside their main construct-validity threats: assurance type, scope specificity, and the opinion/signing-lag pair.

Retail-Investor-Dominated Markets. Retail-investor-heavy ownership structures illustrate why the Signalling Boundary evaluates receiver capacity rather than disclosure content alone. Institutional investors report that inconsistent, incomparable ESG data is a primary barrier to using it in investment decisions (Amel-Zadeh & Serafeim, 2018), a barrier considerably more binding for retail investors, who typically lack dedicated analyst coverage or a shared materiality framework. A cost-verified, materially aligned disclosure may therefore sit inside the boundary for an institutional receiver and outside it for a retail one, not because the disclosure changes, but because decodability does. This is particularly consequential in emerging markets such as India, where retail participation has grown alongside disclosure mandates without a matching expansion in the interpretive infrastructure decodability presumes: analyst coverage, financial-literacy programming, and standardised materiality guidance. This matches evidence that non-investor ESG receivers also process signals unevenly by framing and timing rather than content (Tan et al., 2026), suggesting uneven decodability is structural to ESG signalling environments generally. Chen et al. (2026) show, using Chinese investor-firm interactive platforms, that retail investors' ESG-related inquiries measurably reduce ESG-washing where information asymmetry is high and managerial short-termism is pronounced. This extends the framework rather than contradicting it: decodability and accountability capacity are not fixed by investor type, but depend on the interpretive and interactive infrastructure made available to investors, here a structured firm-investor dialogue platform. Table 2 (row 6) accordingly frames that verdict as contingent on infrastructure access, not retail status per se.

Advancing Sustainable Finance Theory Through the Signalling Boundary

The prospect of the Signalling Boundary is to use disclosure evaluation as a theory-building exercise. It does not accept the existence of an ESG labelled disclosure as a basis for it providing informativeness, but rather requires research and standard-setters to show how a disclosure provides an insight for investors. The focus of disclosure theory is not on the quantity of disclosure, but on how disclosure is done.

Illustrative propositions that would not follow from a disclosure-volume account are obtained by combining the three diagnostic principles, using the weakest link rule. But investor-value relevance is inverted-U in verification cost for a given disclosure volume and context, increasing to an optimum as verification cost increases until it must be reduced: if the required level of disclosure is at or below that which is considered "disclosure", then firms must signal compliance, and the higher the cost of doing so, the more disclosure they are likely to provide, but only up to a point; after that point, investor-value relevance is likely to decline as disclosure is pushed to the minimum required to comply with the disclosure rule, rather than the optimal level for revealing information (Proposition 1). Following a new mandatory assurance regime, the three principles should satisfy in sequence: cost first, at mandated compliance; decodability second, as interpretive infrastructure catches up; accountability last, as detection and enforcement mature, which matches the BRSR signing-date evidence above (Proposition 2). Cross-provider rating disagreement should concentrate in disclosures the framework scores as outside or only partially inside the boundary rather than spread uniformly, and disagreement on disclosures scored fully inside should be no higher than on comparable audited financial-statement items (Proposition 3, sharpening Christensen et al.'s (2022) aggregate finding into a boundary-conditional one). Decodability is a property of a disclosure–receiver pair, not a disclosure alone, so the identical disclosure should be scored on opposite sides of the boundary for high- and low-decodability receivers even when cost and accountability are fixed (Proposition 4). Given that the cheap-talk foundations of the cost principle suggest that value relevance should rise as alignment of interest between the sender and the receiver increases, towards a zero value as the fully divergent case of the babbling equilibria suggested by the cheap-talk theory approaches (Proposition 5, following Crawford & Sobel, 1982), the existence of the babbling equilibria can be tested by comparing firms where executive compensation is directly

tied to externally audited ESG metrics with other firms in the same industry that have the same accounting regime but are not linked to externally audited ESG metrics.

This matters particularly for emerging domains such as generative AI, assurance-based regulation, and retail-investor-dominated markets, where scholars routinely engage mechanisms from computer science, public administration, and behavioural finance. The Signalling Boundary does not require rejecting these mechanisms. It requires showing how their integration changes whether a disclosure informs capital allocation. Cross-disciplinary engagement, on this view, is a resource for theory-building, not a substitute for signalling-specific explanation.

Scope Conditions and Limitations

Five scope conditions bound the diagnostic claims made here. First, the three principles are more likely to move together, and so less independently diagnostic, under a uniform, market-wide regulatory shock than in cross-sectional comparisons within a single population. The BRSR Core illustration demonstrates divergence within one regulatory cohort, but cross-regime designs would still need to account for this co-variance risk. Second, the framework remains conceptual and diagnostic rather than fully measurement-ready. Table 4 lists a first-pass candidate observable for each principle, with its main construct-validity threat, but translating these into validated indices of the kind Zatini et al. (2026) construct for routinisation and greenwashing risk remains a necessary next step. Third, the illustrations here draw substantially on India's regulatory experience and on Indian and Asian emerging-market evidence. Whether the same boundary conditions hold in regimes with different enforcement intensity, analyst density, or ownership structure remains an open empirical question the framework aims to make tractable, not resolve. Fourth, the empirical literature testing these boundary conditions is itself expanding rapidly, and this paper's positioning reflects the literature as it stood when the manuscript was prepared. Fifth, the framework as applied here evaluates one disclosure phenomenon at a time; it does not model whether firms substitute effort across a portfolio of disclosures in response to a binding mandate at one margin i.e. the waterbed effect noted in the Illustrative Application.

Conclusion

The amount of information companies report on their ESG practices has dramatically grown. The resulting issue is a less obvious one: as a metric, narration or rating gets introduced into ESG reporting, it does not necessarily mean that it is used to inform capital allocation. So, the issue before us is not if disclosure is too little or too much, but how disclosure can be increased while still being investor relevant.

The Signalling Boundary is one solution. It is not a line drawn around topics that count as "content," but rather a question as to whether a disclosure has its own cost, as well as reaches a receiver who has the capacity to comprehend it, and is accountable to a feedback mechanism that can detect misrepresentation. These three conditions, rather than being additive, are connected to each other through their worst performer, that is, the most demanding principle is the limiting condition. They work together to link the production, interpretation, and consequence: the type of time a claim gets when it is produced, when it is understood, and when it is held accountable.

The point of view also impacts the potential course of sustainable finance theory going forward. The AI world-building and assurance requirements, and the inclusion of retail investors should not be seen as theories of investor value, merely because they are also ESG-labelled. Their contribution is instead based on whether they alter the costs, decodability or accountability in a way that enhances understanding of what information disclosure is and isn't informative about. As such, the Signalling Boundary may enhance disclosure, rather than limit it. The framework took up the idea of expanding the scope of what firms can report and clarifies what it needs to do to become cumulative, investor-relevant knowledge. It is not the task of sustainable finance research to limit its empirical scope, rather it is a task to bring each expansion of this scope into theoretical significance.

References

1. Agrawal, A. (2026). Value or veneer? Testing the financial credibility of ESG signaling in emerging markets. *Business Strategy and the Environment*. Advance online publication. <https://doi.org/10.1002/bse.71166>
2. Amel-Zadeh, A., & Serafeim, G. (2018). Why and how investors use ESG information: Evidence from a global survey. *Financial Analysts Journal*, 74(3), 87–103. <https://doi.org/10.2469/faj.v74.n3.2>
3. Asif, M., & Searcy, C. (2026). Beyond the green façade: Building a framework to deter ESG greenwashing. *Quality & Quantity*, 60(1), 2585–2605. <https://doi.org/10.1007/s11135-025-02368-4>
4. 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>
5. Bromley, P., & 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. <https://doi.org/10.1080/19416520.2012.684462>
6. Chan, Y.-F., Ngoc, L. M., Oladapo, M. O., Osemeke, G., & Akhtar, I. (2025). When authenticity doesn't pay: Validating an ESG communication authenticity framework and explaining stakeholder–investor decoupling. *Sustainability*, 17(19), Article 8922. <https://doi.org/10.3390/su17198922>
7. Chen, R., Wang, M., Yi, Y., & Gao, Y. (2026). Democratizing corporate governance: How retail investors' ESG attention reduces ESG-washing. *International Review of Economics & Finance*, 105, Article 104897. <https://doi.org/10.1016/j.iref.2026.104897>
8. 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. <https://doi.org/10.2308/TAR-2019-0506>
9. Christensen, H. B., Hail, L., & Leuz, C. (2021). Mandatory CSR and sustainability reporting: Economic analysis and literature review. *Review of Accounting Studies*, 26(3), 1176–1248. <https://doi.org/10.1007/s11142-021-09609-5>
10. Connelly, B. L., Certo, S. T., Ireland, R. D., & Reutzel, C. R. (2011). Signaling theory: A review and assessment. *Journal of Management*, 37(1), 39–67. <https://doi.org/10.1177/0149206310388419>
11. Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6), 1431–1451. <https://doi.org/10.2307/1913390>
12. Del Gesso, C., & Lodhi, R. N. (2025). Theories underlying environmental, social and governance (ESG) disclosure: A systematic review of accounting studies. *Journal of Accounting Literature*, 47(2), 433–461. <https://doi.org/10.1108/JAL-08-2023-0143>
13. Delmas, M. A., & Burbano, V. C. (2011). The drivers of greenwashing. *California Management Review*, 54(1), 64–87. <https://doi.org/10.1525/cmr.2011.54.1.64>
14. Freeman, R. E. (1984). *Strategic management: A stakeholder approach*. Pitman.
15. Fu, L., Boehe, D. M., & Orlitzky, M. O. (2022). Broad or narrow stakeholder management? A signaling theory perspective. *Business & Society*, 61(7), 1838–1880. <https://doi.org/10.1177/00076503211053018>
16. Hoai, N. T. (2026). The mediating role of trust in bridging the ESG assurance credibility chasm. *Discover Sustainability*, 7(1), Article 369. <https://doi.org/10.1007/s43621-026-02747-3>
17. Khan, A., Lasyoud, A. A., & Abbas, M. (2026). Validated signalling: How ownership structure enhances ESG disclosure credibility in an emerging market. *Business Strategy and the Environment*. Advance online publication. <https://doi.org/10.1002/bse.71260>
18. Khan, M., Serafeim, G., & Yoon, A. (2016). Corporate sustainability: First evidence on materiality. *The Accounting Review*, 91(6), 1697–1724. <https://doi.org/10.2308/accr-51383>

19. Kim, Y., & Oh, T. (2026). How ESG disclosure creates value: The mediating effect of ESG performance and the moderating role of sin stock sectors. *Corporate Social Responsibility and Environmental Management*. Advance online publication. <https://doi.org/10.1002/csr.70966>
20. Kim, Y., Park, C., Seo, Y., & Roh, T. (2026). Sensegiving, ESG, and firm value: Mitigating interpretive uncertainty in South Korea. *Business Strategy and the Environment*, 35(3), 4240–4255. <https://doi.org/10.1002/bse.70394>
21. KPMG Assurance and Consulting Services LLP. (2026, February 28). *Accounting and auditing update, Issue No. 115: Chapter 1—Emerging trends in BRSR reporting by listed companies*. KPMG in India. <https://kpmg.com/in/en/insights/2026/02/aau-accounting-and-auditing-update-february-2026.html>
22. Lee, M.-J., Pak, A., Seo, D., & Roh, T. (2026). Navigating the ESG paradox: Strategic pathways between innovation and washing under stakeholder scrutiny. *Business Strategy and the Environment*, 35(5), 6890–6913. <https://doi.org/10.1002/bse.70525>
23. Marquis, C., Toffel, M. W., & Zhou, Y. (2016). Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, 27(2), 483–504. <https://doi.org/10.1287/orsc.2015.1039>
24. Pinto, I., & Gaio, C. (2025). ESG performance and information asymmetry: The moderating role of ownership concentration. *Corporate Social Responsibility and Environmental Management*, 32(5), 5747–5759. <https://doi.org/10.1002/csr.3260>
25. Prayogo, E., Sudarma, M., Purwanti, L., & Adib, N. (2026). Environmental, social, and governance (ESG) disclosure and firm value in Indonesia: A meta-analysis. *International Journal of Business and Society*, 27(1), 349–374. <https://doi.org/10.33736/ijbs.12828.2026>
26. Qiu, X., Gan, X., & Zhang, C. (2026). Generative AI-based readability as a measure of disclosure quality: Evidence from ESG rating divergence. *International Review of Economics & Finance*, 109, Article 105364. <https://doi.org/10.1016/j.iref.2026.105364>
27. Reber, B., Gold, A., & Gold, S. (2022). ESG disclosure and idiosyncratic risk in initial public offerings. *Journal of Business Ethics*, 179(3), 867–886. <https://doi.org/10.1007/s10551-021-04847-8>
28. Securities and Exchange Board of India. (2023, July 12). *BRSR core: Framework for assurance and ESG disclosures for value chain* (Circular No. SEBI/HO/CFD/CFD-SEC-2/P/CIR/2023/122). SEBI.
29. Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3), 355–374.
30. Surya, R. A. S., Andreas, Halim, E. H., & Zarefar, A. (2026). ESG performance and firm value in Indonesia: Do political connections and external assurance matter? *Journal of Risk and Financial Management*, 19(2), Article 131. <https://doi.org/10.3390/jrfm19020131>
31. Taddeo, S., Regoli, A., Weber, O., & Carè, R. (2026). Green talk, costly walk: The financial cost of greenwashing. *Business Strategy and the Environment*. Advance online publication. <https://doi.org/10.1002/bse.70595>
32. Tan, Q. L., Hashim, S., Abdullah, N. L., Zaki, H. O., & Zheng, Z. (2026). Enduring versus episodic ESG signals: How corporate reputation and ESG signal quality shape perceived ESG and purchase intention. *Corporate Social Responsibility and Environmental Management*. Advance online publication. <https://doi.org/10.1002/csr.70745>
33. Thewissen, J., Shrestha, P., Arslan-Ayaydin, Ö., & Yan, B. (2026). Topics that matter: ESG reports and information asymmetry. *Journal of Financial Research*. Advance online publication. <https://doi.org/10.1111/jfir.70053>
34. Tran, T. T. N., Hoang, T. C., & Dau, T. T. K. (2026). Does financial distress shape the effect of ESG disclosure on firm value? *Management Research Review*. Advance online publication. <https://doi.org/10.1108/MRR-05-2025-0441>

35. Uyar, A., Karaman, A. S., & Kilic, M. (2020). Is corporate social responsibility reporting a tool of signaling or greenwashing? Evidence from the worldwide logistics sector. *Journal of Cleaner Production*, 253, Article 119997. <https://doi.org/10.1016/j.jclepro.2020.119997>
36. Verrecchia, R. E. (2001). Essays on disclosure. *Journal of Accounting and Economics*, 32(1–3), 97–180. [https://doi.org/10.1016/S0165-4101\(01\)00025-8](https://doi.org/10.1016/S0165-4101(01)00025-8)
37. Whetten, D. A. (1989). What constitutes a theoretical contribution? *Academy of Management Review*, 14(4), 490–495. <https://doi.org/10.5465/amr.1989.4308371>
38. Yadav, A., & Asongu, S. A. (2026). The AI sustainability paradox: How verification and regulation synergize to curb greenwashing in emerging markets. *Corporate Social Responsibility and Environmental Management*. Advance online publication. <https://doi.org/10.1002/csr.70737>
39. Yu, E. P., Luu, B. V., & Chen, C. H. (2020). Greenwashing in environmental, social and governance disclosures. *Research in International Business and Finance*, 52, Article 101192. <https://doi.org/10.1016/j.ribaf.2020.101192>
40. Zatini, G., Torricella, N., & della Porta, A. (2026). Costly signals and cheap talk: Measuring the decoupling of ESG routinisation and greenwashing risk. *Business Strategy and the Environment*. Advance online publication. <https://doi.org/10.1002/bse.71079>
41. Zhang, Z., Zhang, B., You, J., Wang, N., & Xu, X. (2026). Is more ESG disclosure always better? New evidence from ESG disclosure quantity and firm value. *International Review of Economics & Finance*, 108, Article 105318. <https://doi.org/10.1016/j.iref.2026.105318>

Table 1. The Signalling Boundary's Three Diagnostic Principles

Diagnostic Principle	Core Explanatory Logic	Diagnostic Question	Indicative Alternative Logic If Absent
Signal Cost & Verifiability	Costly, asymmetric production separates honest from fabricated claims	Does the disclosure impose a cost of production that a firm without the underlying performance would find irrational to pay?	Cheap talk, compliance formality, narrative impression management
Receiver Decodability	Value relevance requires a receiver capable of interpreting disclosure against a materiality framework	Does the disclosure explain investor outcomes through a receiver's demonstrated capacity to attend to and interpret it?	Information overload, receiver indifference, unusable complexity
Feedback Accountability	Scrutiny capable of detecting and penalising misrepresentation disciplines future signalling	Does a mechanism exist that can detect and penalise misrepresentation, changing the cost of the next disclosure?	Unmonitored self-report, symbolic compliance, decoupled practice

Table 2. Illustrative Application Across Disclosure Phenomena

Disclosure Phenomenon (Receiver, Where Relevant)	Signal Cost & Verifiability	Receiver Decodability	Feedback Accountability	Boundary Verdict (Binding Constraint)
Mandatory BRSR Core KPI under reasonable assurance (NIFTY100, FY2024–25)	✓	Partial	Partial	Partially inside — binding: decodability, accountability
Voluntary assurance beyond Core scope, assurance type undistinguished in disclosure	Partial	X	X	Outside — binding: decodability, accountability

Disclosure Phenomenon (Receiver, Where Relevant)	Signal Cost & Verifiability	Receiver Decodability	Feedback Accountability	Boundary Verdict (Binding Constraint)
Unaudited narrative sustainability report	X	Partial	X	Outside — binding: cost, accountability (cheap talk)
AI-generated sustainability narrative, no assurance	X	✓	X	Outside — binding: cost, accountability (fluent cheap talk)
Materially aligned metric, read by institutional analyst	✓	✓	✓	Inside the boundary
Same metric, read by retail investor lacking platform-based interpretive infrastructure	✓	X	✓	Outside — binding: decodability (infrastructure-contingent)
Ownership-monitored disclosure, formal assurance underdeveloped (concentrated-ownership emerging market)	Partial	Partial	Partial	Partially inside — binding: all three (tied); informal substitute possible
Aggregated third-party ESG rating, opaque methodology	Partial	Partial	Partial	Partially inside — binding: all three (tied); rating divergence

Table 3. Recoding Published ESG–Value Moderators onto the Signalling Boundary's Three Principles

Published Moderator	Principle It Proxies	Untested Prediction
Firm size, age, profitability (Agrawal, 2026)	Cost & Verifiability	Tracks verification cost, not size; equal-cost smaller firms should match.
Market recognition (Agrawal, 2026)	Receiver Decodability	Tracks analyst coverage, not recognition; equal-coverage firms should match.
Sin-stock sector (Kim & Oh, 2026)	Feedback Accountability	& Tracks scrutiny intensity, not sector label; comparably scrutinised sectors should match.
Political connections (Surya et al., 2026)	Feedback Accountability	& Tracks enforcement faced, not connection status; unconnected firms in low-enforcement settings should match.
Financial distress (Tran et al., 2026)	Cost & Verifiability	Tracks sustained verification spending, not distress; firms maintaining assurance should not reverse.
Ownership concentration (Pinto & Gaio, 2025)	Receiver Decodability	Tracks private-channel access, not concentration; equal-access dispersed owners should match.
Blockholder monitoring, business-group affiliation (Khan et al., 2026)	Feedback Accountability	& Tracks monitoring exercised, not group membership; equally monitored standalone firms should match.
Topic diversity (Thewissen et al., 2026)	Receiver Decodability	Tracks receiver capacity, not topic count; higher-capacity receivers should turn later.

Published Moderator	Principle It Proxies	Untested Prediction

Table 4. Candidate Observables for the Three Diagnostic Principles

Principle	Candidate Observable	Data Source	Main Threat
Cost & Verifiability	Assurance type/scope (none, limited, reasonable; Core-only vs. extended)	Assurance report; BRSR filing	Ordinal proxy only; rigor varies within a tier.
Receiver Decodability	Specificity of scope language (mandatory vs. voluntary distinguished)	Assurance scope paragraph	Necessary, not sufficient; receiver uptake unconfirmed.
Feedback & Accountability	Qualified-opinion rate; audit-to-assurance signing lag	Assurance and audit reports	Short lag also consistent with weak detection (see above).