

Sustainability as an Intangible Asset: A Labour-Input Measure from Job Postings*

Salma Lemkhente
ESADE Business School
`salma.lemkhente@alumni.esade.edu`

May2026

Abstract

The leading commercial ESG ratings agree only modestly on which firms are sustainable, and the disclosure-based measurement infrastructure they share rewards firms that report well rather than firms that build sustainability capability. I construct a firm-year, labour-input measure of corporate sustainability orientation, LSI, by embedding 60.7 million LinkUp job postings between 2013 and 2020 against a curated 31-topic reference corpus with the bge-m3 multilingual encoder and aggregating posting-level cosine similarities to the firm-year.

Across firms, LSI converges with the leading commercial ratings at magnitudes well below the rater-to-rater range, cross-sectionally precedes subsequent ESG-rating changes over one-to-three-year horizons, and adds significant predictive power for green patenting in the specification where commercial ratings do not. Wage-weighting the firm-year aggregator with external occupational wages produces a dollar-denominated investment flow that I cumulate into a perpetual-inventory stock, LSI-K. In the investment-q regression, the stock is incremental to both Peters-Taylor knowledge and organisation capital, establishing sustainability orientation as a distinct intangible-capital component.

1 Introduction

Sustainability has become central to corporate strategy and capital allocation, but firm-level sustainability measurement remains unusually weak. The leading commercial ESG rating providers often disagree on which firms are sustainable: across six major providers, pairwise correlations of ESG scores range from 0.38 to 0.71, far below the near-unity correlations observed for credit ratings (Berg et al., 2022). More disclosure does not necessarily resolve the problem. Christensen et al. (2022) show that rating disagreement rises with the volume of ESG-related disclosure, and Avramov et al. (2022) demonstrate that the resulting uncertainty has first-order asset-pricing implications. A parallel greenwashing literature explains why this disagreement matters. Voluntary climate disclosures by TCFD-supporting firms often contain selective and non-binding claims (Bingler et al., 2022), and high-ESG firms do not consistently behave differently

*Preliminary draft prepared for the EFMA 2026 Doctoral Consortium. Comments are welcome. I thank the consortium organisers and discussants for their feedback.

from low-ESG firms in labour, environmental, and antitrust enforcement data (Raghunandan and Rajgopal, 2024). The concern is therefore not only that current sustainability measures are noisy, but that they may partly reward firms’ ability to report sustainability rather than their investment in the capabilities needed to perform it.

The stakes are increasing because these measures now sit inside both regulation and capital allocation. Mandatory sustainability disclosure regimes are expanding across jurisdictions, including the EU’s Corporate Sustainability Reporting Directive, climate-disclosure initiatives in the United States, and the IFRS Foundation’s ISSB S1 and S2 standards. At the same time, institutional investors managing large pools of ESG-mandated capital allocate across firms using sustainability-related signals (Krueger et al., 2020; Hartzmark and Sussman, 2019). If firm-level sustainability measures are noisy, disclosure-sensitive, or vendor-specific, they constrain both the active empirical agenda in sustainability finance and the capital-allocation decisions that agenda informs.

What is missing is a firm-year measure of corporate sustainability orientation that is anchored in observable economic activity, comparable across firms and over time, and independent of any particular rating vendor’s aggregation model. This paper proposes such a measure using firms’ labour demand for sustainability-related capabilities.

Job postings provide a revealed-preference signal of labour demand. A firm advertising a sustainability-related role is not merely describing itself; it is allocating resources to acquire a capability. A growing postings-based literature uses this signal to study green-skill demand under environmental regulation (Vona et al., 2018), the green-occupation wage premium (Consoli et al., 2016), low-carbon transition skills (Saussay et al., 2022), and firm-level green hiring that predicts subsequent green patenting and profitability (Darendeli et al., 2022). This literature is the natural starting point for my construct, but it is limited in two ways for the broader sustainability-finance agenda. First, its topical scope is largely environmental, focusing on climate, energy, green skills, and green occupations. Second, its measurement approach typically relies on fixed keyword, occupation, or skill labels, which have difficulty capturing the paraphrase and synonymy that characterize unstructured job-posting text.

I extend this approach in two directions. First, I expand topical coverage from environmental-only constructs to a broader ESG and sectoral taxonomy. Second, I replace keyword counting with dense-embedding similarity to an action-language reference corpus. This approach preserves the conceptual content of multi-paragraph sustainability descriptions while allowing job postings to match those descriptions even when they use different terminology.

The resulting measure, Labour-input Sustainability Intangible (LSI), is constructed at the firm-year level from 60.7 million LinkUp job postings between 2013 and 2020, covering 10,232 ever-listed firms. Each posting is embedded with the bge-m3 multilingual dense encoder (Chen et al., 2024) and scored by cosine similarity against a curated corpus of 31 sustainability-topic anchors covering environmental, social, governance, and sectoral pillars. The anchors are synthesised from 26 sustainability-standard source families, including GRI, SASB, ISO, TCFD, ILO, OECD, EU CSRD, and the GHG Protocol, under a reproducible editorial procedure that converts disclosure-register prose into operational descriptions of substantive sustainability work. I validate this editorial step against a mechanical baseline that skips the editing procedure, as

described in Section 3.9. I then aggregate posting-level scores into a firm-year intensity measure and, in a second step, wage-weight the measure using BLS Occupational Employment and Wage Statistics to construct a dollar-denominated labour-input flow. Cumulating this flow through a perpetual-inventory method yields LSI-K, a sustainability-intangible stock comparable in units to Peters-Taylor organisation capital.

I evaluate LSI along four dimensions. First, I ask whether it converges with leading firm-level sustainability benchmarks at magnitudes consistent with capturing related-but-distinct content. Second, I test whether it tracks firm-level climate exposure in the expected direction. Third, I examine whether it predicts subsequent sustainability outcomes that commercial ESG ratings systematically miss. Fourth, I ask whether the wage-weighted and cumulated version of the construct behaves like an intangible asset in the sense formalised by Peters and Taylor (2017).

The headline results are as follows. LSI is positively but modestly correlated with commercial ESG ratings. Across firms, it converges with LSEG and S&P Global ESG scores at Spearman correlations of 0.05–0.10, magnitudes consistent with the rating-divergence literature’s expectation for independent sustainability signals. The environmental pillar of LSI also correlates with the firm-level climate-exposure measure of Sautner et al. (2023) at Spearman 0.24. These patterns suggest that LSI captures content related to existing sustainability benchmarks, but is far from a mechanical proxy for them.

LSI also predicts future changes in commercial ESG ratings. Today’s LSI forecasts subsequent LSEG and S&P Global rating revisions over horizons of one to three years, even controlling for current ratings and standard firm-level financial characteristics. The relationship is robust across both rating providers and multiple horizons, but is not driven by within-firm year-to-year variation. I interpret this pattern as evidence that ratings under-weight persistent firm-level sustainability orientation that LSI captures and gradually incorporate it over time. In this sense, the measure formalises the qualitative claim in Darendeli et al. (2022) that green hiring contains forward-looking information about firms’ sustainability trajectories.

The measure carries information beyond ratings in predicting green innovation. Replicating Cohen et al. (2026)’s green-patent regression with their controls and fixed-effect structure, LSI adds significant predictive power for next-year green patent counts in the specification where LSEG ESG is only marginally significant. The result simultaneously replicates the ESG-innovation disconnect on the rating side and shows that a labour-input measure helps resolve it. Firms that allocate labour toward sustainability-related capability are more likely to generate subsequent green innovation, even when commercial ESG ratings contain limited information about that outcome.

The valuation and profitability results are consistent with an intangible-investment interpretation. In regressions following the spirit of Edmans (2011) and Eisfeldt and Papanikolaou (2013), higher LSI is associated with lower near-term ROA but higher future Tobin’s q. This pattern is consistent with a costly investment that depresses contemporaneous accounting profitability while being capitalised by markets over longer horizons. Building on these flow-level results, I integrate the construct into the intangible-capital framework by wage-weighting LSI into a dollar-denominated investment flow and cumulating it into the stock measure LSI-K. In augmented investment-q regressions, LSI-K remains informative after controlling for Peters-

Taylor knowledge capital and organisation capital, and cross-sectional quintile sorts recover the elevated-q, low-ROA signature that [Eisfeldt and Papanikolaou \(2013\)](#) document for organisation capital. I interpret this evidence as a framework proposal rather than a final capitalisation model: the central depreciation rate is set at $\delta = 0.20$, with $\delta \in \{0.15, 0.30\}$ as a bracket, and pinning down the empirically correct depreciation rate is outside the scope of the present paper.

The evidence also imposes an important boundary condition. LSI does not predict fewer subsequent Reprisk environmental, social, or governance incidents in any horizon, severity-weighting scheme, or fixed-effects specification I test. A Sun-Abraham staggered event study around firms' first incidents ([Sun and Abraham, 2021](#)) likewise rules out reactive responses of economically meaningful magnitude across a ± 21 -quarter window. This null result is informative. LSI should not be interpreted as a direct measure of realised sustainability performance or near-term risk reduction. It is better understood as a measure of firms' revealed labour-market investment in sustainability capability.

The paper contributes to four literatures. First, it extends the empirical intangible-capital literature ([Corrado et al., 2009](#); [Peters and Taylor, 2017](#); [Crouzet et al., 2022](#); [Sichel, 2025](#)) by introducing a firm-year, dollar-denominated measure of sustainability labour investment and showing that it is not spanned by standard knowledge- and organisation-capital measures. Second, it contributes to ESG measurement ([Berg et al., 2022](#); [Christensen et al., 2022](#); [Avramov et al., 2022](#); [Raghunandan and Rajgopal, 2024](#); [Bingler et al., 2022](#)) by providing an action-oriented benchmark that is independent of voluntary disclosure choices and rating-vendor aggregation rules. Third, it contributes to empirical sustainability finance ([Cohen et al., 2026](#); [Darendeli et al., 2022](#); [Pástor et al., 2021](#)) by showing that sustainability labour input predicts green innovation and future rating revisions beyond commercial ESG scores. Fourth, it contributes to the job-postings-as-sustainability literature ([Vona et al., 2018](#); [Consoli et al., 2016](#); [Saussay et al., 2022](#); [Darendeli et al., 2022](#)) by moving beyond environmental keyword and occupation measures toward a broader ESG taxonomy scored with dense text embeddings and made interoperable with the intangible-capital framework through wage weighting.

The rest of the paper proceeds as follows. Section 2 situates the contribution within the four related literatures: intangible capital, ESG measurement, job postings as labour-input infrastructure, and substantive sustainability outcomes. Section 3 describes the construct, the curated corpus, the embedding-and-scoring pipeline, the wage-weighting procedure, and the validation strategy. Section 4 documents the data and sample construction. Section 5 reports the empirical results. Section 6 integrates the evidence, clarifies the boundary conditions and lays out future extensions.

2 Related Literature

This paper connects four research streams: intangible-capital measurement, firm-level sustainability measurement, job postings as labour-input data, and the empirical literature on sustainability outcomes. I use this section to identify the measurement gap that motivates LSI and to clarify the type of object the measure is intended to capture: an input-side measure of sustainability-related capability investment, not a direct measure of realised sustainability performance.

2.1 The intangible-capital framework and the missing sustainability category

The own-account intangible-investment framework developed by [Corrado et al. \(2005, 2009\)](#) shows that a large share of business investment takes the form of assets produced inside the firm for the firm’s own future use. These assets include software, brand equity, training, organisational design, and the labour devoted to scientific and engineering R&D. The framework is deliberately open-ended: when a new category of durable, firm-specific capability becomes economically material, it should enter the intangible-capital catalog with its own investment measure ([Sichel, 2025](#)). This makes the boundary of the catalog an empirical and measurement question rather than a fixed accounting category.

At the firm level, [Peters and Taylor \(2017\)](#) operationalise this logic by constructing a firm-year stock of intangible capital as the sum of knowledge capital, cumulated from R&D, and organisation capital, cumulated from a capitalised share of SG&A. Adding this stock to the denominator of the investment-q regression materially improves the investment-q relation. [Eisfeldt and Papanikolaou \(2013\)](#) show that organisation capital is priced in the cross-section of equity returns, and [Eisfeldt et al. \(2022\)](#) extend the asset-pricing logic to a broader intangible-value factor. [Crouzet et al. \(2022\)](#) emphasise that, despite this progress, several conceptually relevant intangible categories remain outside the accepted empirical catalog because they lack firm-level measures.

Sustainability capability is a natural candidate for such a missing category. Firms devote labour, training, monitoring systems, audit procedures, supplier-engagement protocols, and managerial routines to sustainability-related work. These inputs are produced inside the firm, are at least partly durable and firm-specific, and may generate future benefits through regulatory preparedness, customer access, reputational resilience, and eligibility for sustainability-mandated capital. Yet the existing intangibles taxonomy does not contain a firm-year, dollar-denominated measure of own-account sustainability investment. The construct I introduce is designed to fill that gap by translating sustainability-related labour demand into a wage-weighted investment flow and, subsequently, into a perpetual-inventory stock comparable to existing intangible-capital measures.

2.2 The measurement limits of existing firm-level sustainability metrics

The dominant firm-level sustainability data infrastructure consists of commercial ESG ratings, disclosure-based and text-based measures, and event databases of corporate sustainability incidents. Each modality is useful, but each captures a different object from the input-side capability investment I study.

The first limitation is rating divergence. [Berg et al. \(2022\)](#) document that pairwise correlations between leading ESG rating providers range from 0.38 to 0.71 and decompose the disagreement into scope, measurement, and weighting differences. [Christensen et al. \(2022\)](#) show that disagreement among raters increases with the volume of ESG-related disclosure, suggesting that additional disclosure is not necessarily transformed into a common sustainability signal. [Avramov et al. \(2022\)](#) show that rating uncertainty matters for the equilibrium pricing of sustainable assets in the framework of [Pástor et al. \(2021\)](#). These findings motivate measures that are independent of any single rating provider’s scope choices, measurement rules, and aggregation

weights.

A second limitation is the distinction between disclosure and underlying capability. [Bingler et al. \(2022\)](#) use ClimateBERT to show that voluntary climate disclosures by TCFD-supporting firms often contain cheap talk and cherry-picking. [Bingler et al. \(2024\)](#) link cheap talk in climate disclosure to subsequent emissions growth and negative news coverage. [Marquis et al. \(2016\)](#) show that civil-society scrutiny and global transparency norms discipline selective disclosure, while [Lyon and Maxwell \(2011\)](#) provides the formal logic for greenwashing under audit threat. [Raghunandan and Rajgopal \(2024\)](#) directly compare responsible-firm labels with enforcement actions in labour, environmental, and antitrust domains and find limited evidence that high-ESG firms behave differently. At the same time, [Khan et al. \(2016\)](#) show that performance on material sustainability issues predicts future financial performance. Taken together, this literature does not imply that sustainability is economically irrelevant; it implies that disclosure-based measures can mix substantive capability with reporting quality and impression management.

A third strand develops text-based measures of climate exposure, attention, and sentiment. [Sautner et al. \(2023\)](#) construct firm-quarter climate-exposure measures from earnings-call transcripts, decomposing exposure into opportunity, physical, regulatory, and sentiment components. [Engle et al. \(2020\)](#) use climate news to build hedging portfolios, while [Pankratz et al. \(2023\)](#) link physical climate exposure to firm operating performance and investor surprises. [Webersinke et al. \(2022\)](#) documents the technical infrastructure behind domain-adapted climate language models. These measures are valuable, but they primarily capture what firms and market participants discuss or are exposed to. They do not directly measure the labour input firms commit to building sustainability-related capability.

The gap, therefore, is not the absence of sustainability data. It is the absence of a firm-year measure that is independent of rating aggregation, less exposed to voluntary self-presentation, and anchored in observable input-side economic activity.

2.3 Job postings as a labour-input measurement infrastructure

Online job postings provide high-frequency, granular information on firms' labour demand. [Hershbein and Kahn \(2018\)](#) and [Deming and Kahn \(2018\)](#) use Burning Glass vacancy data to study changes in skill requirements across occupations, firms, and labour markets. [Acemoglu et al. \(2022\)](#) use vacancy text to identify AI-exposed occupations and changes in AI-related demand. In finance, [Babina et al. \(2024\)](#) combine employee data and job postings to measure firm-level AI investment and show that it predicts product innovation and firm growth. [Abis and Veldkamp \(2024\)](#) use postings to construct firm-level data and AI stocks for investment-management firms, applying a perpetual-inventory logic to labour inputs in knowledge production.

This literature establishes the template I use: postings can be treated as revealed-preference evidence of investment in a specific intangible capability. A firm advertising a role is willing to pay the market wage for the corresponding skill bundle. When aggregated over firms and years, the postings signal can therefore be interpreted as a labour-input investment flow rather than merely as a textual descriptor of the firm.

A related environmental literature applies postings data to green skills and green occupations. [Vona et al. \(2018\)](#) show that environmental regulation increases demand for green engineering

and managerial skills. [Consoli et al. \(2016\)](#) document how green jobs differ from non-green jobs in skill and human-capital requirements, and [Saussay et al. \(2022\)](#) study low-carbon skill demand and wage gaps in vacancy data. These papers show that postings contain economically meaningful information about environmental labour demand.

The closest paper to mine is [Darendeli et al. \(2022\)](#), who construct a firm-year green-hiring score from green vocabulary in job postings and show that green hiring predicts subsequent green patenting and profitability in settings where ESG scores have limited explanatory power. I use this result first as a validation benchmark. If LSI is a credible labour-based sustainability measure, it should reproduce the central postings-to-green-innovation result: labour input should predict green patenting where commercial ESG ratings do not. I then extend their contribution in three directions. First, I broaden the construct from environmental green hiring to a more holistic sustainability measure covering environmental, social, governance, and sectoral topics. Second, I move from keyword-based classification to dense-embedding similarity against an action-language reference corpus, allowing the measure to capture paraphrase and synonymy in unstructured job text. Third, I use the lead-lag relation between LSI and future ESG rating revisions to interpret why the labour-input signal matters: it captures persistent sustainability capability that ratings appear to recognise only gradually. The contribution is therefore not only to confirm the green-hiring result with a broader labour-based measure, but to place it inside a sustainability-intangible framework.

2.4 Substantive sustainability outcomes

The outcome literature clarifies which empirical tests a useful sustainability-capability measure should pass, and which tests instead delimit its scope.

The first relevant outcome is green innovation. [Aghion et al. \(2016\)](#) show how environmental policy directs technical change toward green innovation. [Cohen et al. \(2026\)](#) document that firms with lower ESG ratings, especially traditional energy producers, generate a disproportionate share of high-impact green patents, and that ESG ratings fail to predict green innovation output. [Darendeli et al. \(2022\)](#) provide one resolution by showing that green hiring predicts subsequent green patenting. This literature motivates green patents as a central test of whether LSI captures sustainability-related capability beyond commercial ESG ratings. Patent-firm linkages and patent-value estimates are drawn from the KPSS patent data ([Kogan et al., 2017](#)).

The second relevant outcome is corporate sustainability incidents. Reprisk-based studies show that adverse ESG events are economically meaningful. [Glossner \(2021\)](#) studies ESG-incident recidivism and investor underreaction, while [Krüger \(2015\)](#) documents the shareholder-wealth effects of corporate-goodness news. For LSI, incident data provide a discriminant-validity test. If LSI is an input-side capability measure, it need not move one-for-one with realised incident counts, which are shaped by operational exposure, media visibility, compliance failures, and scale. A null relation between LSI and subsequent incidents therefore helps distinguish sustainability capability investment from realised short-horizon incident avoidance.

The third relevant outcome is firm performance and valuation. [Khan et al. \(2016\)](#) show that performance on financially material sustainability issues predicts future financial performance. [Edmans \(2011\)](#) provides the closest intangibles-and-performance analogue: employee satisfaction,

an intangible firm attribute absent from financial statements, predicts long-horizon equity returns. In the intangible-capital literature, [Eisfeldt and Papanikolaou \(2013\)](#) shows that organisation capital is priced in the cross-section. These papers motivate testing whether LSI has the classic intangible-investment signature: short-run costs reflected in accounting profitability and longer-run valuation benefits reflected in Tobin’s q .

Finally, sustainable-investing models clarify why sustainability-related capability may be relevant for capital markets. [Pástor et al. \(2021\)](#) develop an equilibrium model in which sustainable assets earn lower expected returns but may outperform after positive shocks to ESG demand. [Pástor et al. \(2022\)](#) decompose green returns into expected-return and unexpected-return components, while [Hartzmark and Sussman \(2019\)](#) document mutual-fund-flow responses to salient sustainability ratings. [Bolton and Kacperczyk \(2021\)](#) show that carbon emissions are priced in the cross-section of US equity returns, and [Krueger et al. \(2020\)](#) document institutional investors’ attention to climate risks. [Pástor et al. \(2025\)](#) reviews the broader field. I do not attempt to build a new sustainability factor in this paper. Instead, I use valuation and profitability patterns to ask whether sustainability labour input behaves like an intangible-capital component.

2.5 Where this paper sits

The literature leaves a clear measurement gap. ESG ratings are widely used but disagree and partly reflect disclosure and aggregation choices. Text-based climate measures capture exposure, attention, and sentiment rather than labour input. Event databases capture realised adverse outcomes rather than ex ante capability investment. Job postings provide a revealed-preference signal of labour demand, but existing green-hiring measures are primarily environmental, keyword-based, and not integrated into the dollar-denominated intangible-capital framework.

This paper fills that gap by constructing a firm-year measure of labour-input sustainability investment. LSI is independent of voluntary sustainability disclosure and rating-vendor aggregation rules; broadens the postings-based green-hiring literature from environmental skills to a wider ESG and sectoral taxonomy; and converts the resulting labour-demand signal into a wage-weighted investment flow and perpetual-inventory stock. In the language of [Corrado et al. \(2005, 2009\)](#), the claim is that sustainability-related capability is an own-account intangible: it is produced inside the firm, partly durable, firm-specific, and accumulated through identifiable labour inputs. The empirical contribution is to show where this input-side measure is informative—commercial rating revisions, green innovation, and intangible-capital valuation patterns—and where it is not: realised short-horizon ESG incident reduction. That boundary is central to the interpretation of the construct. LSI measures firms’ revealed investment in sustainability-related capability, not realised sustainability performance itself.

3 Construct and Measurement

This section defines LSI and describes how I construct it from job postings. The procedure has four steps. First, I define sustainability capability as an own-account intangible produced through firm-level labour input. Second, I construct an action-language reference corpus that describes

substantive sustainability work. Third, I score each posting by dense-embedding similarity to that corpus. Fourth, I aggregate posting-level scores into firm-year intensity, wage-weighted investment-flow, and perpetual-inventory stock measures.

3.1 The construct: own-account labour input for sustainability capability

I treat firm-level sustainability capability as an own-account intangible in the sense of [Corrado et al. \(2005, 2009\)](#). The asset is produced inside the firm, is at least partly durable and firm-specific, and may generate future economic benefits through regulatory preparedness, access to sustainability-sensitive customers and investors, and the accumulation of organisational routines that markets may capitalise even when they are not immediately visible in accounting earnings ([Krueger et al., 2020](#); [Hartzmark and Sussman, 2019](#); [Pástor et al., 2021, 2022](#)). The target construct is therefore the dollar value of labour demand that a firm directs in a given period toward producing or maintaining sustainability-related capability. This is the sustainability analogue of [Peters and Taylor \(2017\)](#)’s organisation-capital logic, where part of SG&A is interpreted as own-account investment in organisational capability.

I observe this labour demand through job postings rather than realised employment. Job postings are useful for three reasons. First, a posting is a revealed-preference signal that the firm is willing to pay the market wage for the corresponding skill bundle. [Acemoglu et al. \(2022\)](#) and [Babina et al. \(2024\)](#) validate the postings-as-investment interpretation for artificial-intelligence capability; I apply the same logic to sustainability capability. Second, postings are available at high frequency and at the open-position level, while firm-disclosed sustainability employment data are sparse and potentially exposed to the disclosure-versus-action bias documented in [Bingler et al. \(2022, 2024\)](#). Third, postings can be aggregated into firm-year labour-demand flows that are close in spirit to the CHS own-account investment concept.

I do not observe whether each posting results in a realised hire. I therefore interpret postings as labour-demand signals rather than realised employment flows. This limitation is common in postings-based work ([Hershbein and Kahn, 2018](#); [Deming and Kahn, 2018](#); [Babina et al., 2024](#); [Darendeli et al., 2022](#)). In the empirical analysis, firm fixed effects absorb stable differences in posting intensity and fill rates, while wage-weighting converts observed labour demand into an economically scaled investment proxy. Wage-weighting does not solve the fill-rate problem, but it prevents low-wage and high-wage vacancies from contributing equally to the dollar-denominated investment flow.

Let i index firms, t calendar years, and P_{it} the set of postings published by firm i in year t . The measurement problem is to assign each posting $p \in P_{it}$ a sustainability-capability score s_p , aggregate these scores to the firm-year level, and then convert the resulting labour-demand signal into a dollar-denominated investment flow and stock.

3.2 The reference corpus

The construct requires a reference corpus that defines what sustainability work looks like. I make four design choices explicit: the source materials, the topical taxonomy, the action-language register, and the curation procedure.

Source materials. A common shortcut in empirical ESG measurement is to anchor a construct on a single disclosure framework, such as GRI or SASB. That choice is convenient, but it exposes the measure to the same disclosure bias that motivates this paper. Disclosure standards are written to tell firms what to report, not necessarily to describe the operational work required to build sustainability capability. A measure anchored only on disclosure prose may therefore reward postings that describe sustainability reporting rather than postings that describe sustainability work.

I therefore draw source material from a broad set of standard-setting, regulatory, management-system, labour, human-rights, climate-accounting, sustainable-finance, and sectoral benchmark documents. The final corpus uses primary documents from 26 source families, including GRI, SASB, ISO, TCFD, ILO, OECD, the GHG Protocol, UN Guiding Principles, IFC Performance Standards, CDP, EU CSRD/ESRS, the EU Taxonomy, PCAF, TNFD, SBTi, PRI, SA8000, LEED, and other sustainability frameworks.¹ This multi-source design reduces dependence on any single framework’s scope choices, a concern made salient by the rating-divergence evidence in [Berg et al. \(2022\)](#).

Topical taxonomy. The corpus is organised into 31 topic anchors grouped into four categories: environmental, social, governance, and sectoral. Table 20 reports the full taxonomy with topic IDs and titles. I do not score postings against pillar-level anchors directly. Instead, pillar scores are constructed as the maximum similarity across the topics within each pillar, as described in Section 3.5. This avoids duplicating source material across multiple levels of aggregation and lets narrow topic matches contribute cleanly to pillar-level scores.

Action-language register. The most important editorial choice is the register of the reference corpus. The greenwashing and cheap-talk literatures show that ESG disclosure often contains aspirational verbs, reporting scaffolding, and standard-setting language that can be weakly connected to substantive action ([Bingler et al., 2022, 2024](#); [Marquis et al., 2016](#); [Lyon and Maxwell, 2011](#); [Raghunandan and Rajgopal, 2024](#)). A corpus built directly from such text risks scoring postings as sustainability-intensive because they resemble disclosure language rather than because they describe operational sustainability work.

I address this problem at the corpus level. Each topic anchor is written in an action-language register that removes disclosure scaffolding, aspirational language, and standard-setting preambles. The anchors instead foreground operational verbs—such as audit, monitor, measure, train, certify, remediate, enforce, inspect, model, mitigate, source, screen, investigate, and escalate—and the substantive work objects those verbs act on: controls, training content, supplier-screening procedures, remediation processes, audit routines, emissions inventories, safety systems, and governance protocols. This mirrors the logic of [Loughran and McDonald \(2011\)](#): when a generic reference embeds the wrong conceptual register, the solution is to redesign the reference around the specific construct one wants to measure.

¹The full source-family list, document versions, extraction dates, and file paths are reported in Appendix A. ESCO and O*NET are excluded from the anchor corpus and reserved for external validation, so that the validation tests in Section 3.9 are not mechanically circular. ISSB S1 and S2 are excluded from the main corpus for time-period coherence, because their 2023 disclosure language postdates most of the 2013–2020 empirical sample.

Prompted curation and reproducibility. I instantiate each topic anchor through a prompted-synthesis procedure. A locked prompt template takes as inputs the source documents for a topic and a concept-focus description, and instructs the model to synthesise an anchor in the action-language register described above. The output is capped at 4,096 tokens, matching the truncation cap used on the posting side.²

Prompted curation does not eliminate researcher discretion. The source selection, topic taxonomy, prompt design, and model choice are all researcher decisions. Its advantage is that these decisions are observable and repeatable: the source files, prompt template, generated anchors, and code can be archived and rerun. To benchmark the value of the editorial step, I also construct a mechanical alternative corpus. This alt-corpus uses the same source documents, extracted with a standard PDF-to-text pipeline and chunked with the bge-m3 tokenizer, but skips the prompted action-language synthesis. The curated and mechanical corpora therefore differ in editorial register, not in source coverage. Section 3.9 uses this comparison to evaluate whether the action-language filter changes the empirical behaviour of the measure.

3.3 From keyword counting to dense embeddings

The traditional text-as-data approach in economics and finance is dictionary-based: the researcher specifies a list of words or phrases associated with the construct, and the document-level measure is a function of their frequency (Gentzkow et al., 2019). This approach has produced influential and replicable measures, including financial sentiment (Loughran and McDonald, 2011), firm-level political risk (Hassan et al., 2019), and news-implied volatility (Manela and Moreira, 2017). It is transparent and inexpensive, but it is not well matched to the measurement problem here.

I use dense embeddings for three reasons. First, sustainability vocabulary is wide and evolving. A keyword list may miss semantically equivalent descriptions that use different terminology. For example, a posting that mentions “Scope 3 emissions accounting” and one that mentions “value-chain greenhouse-gas inventory development” describe closely related work even if they do not share the same keywords. Dense sentence embeddings map semantically similar text to nearby points in a vector space, allowing cosine similarity to generalise across paraphrase and synonymy (Reimers and Gurevych, 2019; Karpukhin et al., 2020). Cosine similarity on firm-level text has an established methodological precedent in finance: Hoberg and Phillips (2010, 2016) use cosine similarity on 10-K product descriptions to construct text-based network industries (TNIC) classifications, replacing fixed SIC/NAICS categories with continuous firm-pair similarity. The methodological move I make is analogous in spirit but operates on the posting-versus-anchor axis rather than the firm-versus-firm axis: the unit of comparison is each individual posting against each topic anchor, and the firm-year measure aggregates these posting-level similarities.

Second, the reference corpus consists of multi-paragraph action-language descriptions rather than term lists. A keyword extraction would discard much of the structure encoded in the anchors: the relevant work objects, responsibilities, monitoring routines, and role-specific verbs. Embedding similarity preserves the prose-level content of the reference and compares it to the

²The full prompt template is reproduced in Appendix A.3. The current version uses Claude Opus 4.7 with a long-context window because the procedure requires reading all source documents for a topic in one pass. The exact model choice is an implementation detail rather than a conceptual requirement.

prose-level content of the posting.

Third, the measure requires a continuous score. Domain-adapted models such as Climate-BERT, FinBERT, and ESG-BERT are often deployed as classifiers that assign text to discrete categories (Webersinke et al., 2022; Bingler et al., 2022; Araci, 2019; Schimanski et al., 2024). A classifier would require a labelled training set and would return discrete posting-level decisions. By contrast, embedding similarity produces a continuous score for every posting-topic pair, generalises naturally across the 31 topic anchors, and lets the substantive content of the curated reference define the construct.

I use a general-purpose multilingual encoder rather than a domain-adapted ESG encoder in the main specification. This is a conservative choice. A domain-adapted encoder trained on ESG disclosure may encode some of the disclosure-versus-action bias that the action-language corpus is designed to remove. A general-purpose encoder lets the curated reference and the posting text do the construct work.³

3.4 The embedding model

I use bge-m3 (Chen et al., 2024) as the embedding model. The choice is driven by four features of the setting.

First, bge-m3 supports long inputs. Job postings often contain detailed descriptions of responsibilities, qualifications, reporting lines, and compliance duties; sustainability-relevant content can appear far below the opening lines. I cap inputs at 4,096 tokens to bound compute cost while preserving nearly all substantive content in the validation sample.⁴

Second, bge-m3 is multilingual. Most LinkUp postings are in English, but non-English postings are not random: they are concentrated among firms headquartered in Continental Europe and Asia-Pacific. Dropping them or translating them would introduce either sample-selection or translation noise. A multilingual encoder allows postings to be scored in their original language.

Third, bge-m3 is designed for retrieval. Because the scoring task is to retrieve the closest sustainability-topic anchor for each posting, retrieval quality is more directly relevant than classification accuracy on a fixed ESG label set.

Fourth, the model has open weights. Closed-API encoders are vulnerable to silent vendor updates that can change embeddings over time and complicate exact replication. Versioned open weights allow the embedding step to be reproduced.

I prepend the posting title to the posting description before tokenisation, because some postings carry the sustainability-specific content in the title while using generic managerial language in the body. I use the model’s default dense representation and L2-normalise all posting and anchor embeddings before scoring.⁵

³A robustness exercise using climate- and ESG-adapted encoders as alternative embedding backbones is planned for the appendix version of the paper. The purpose is not to select the best encoder ex post, but to quantify how much of the LSI variation is robust to encoder choice.

⁴The token-length distribution and over-cap share are reported in Appendix A. In the stratified validation sample, fewer than 0.01% of postings exceed the 4,096-token cap.

⁵As a robustness check, I compare the default pooling rule with mean pooling on the validation sample. The resulting posting-level scores correlate at Spearman 0.97, so the pooling choice does not materially affect the measure.

3.5 Posting-level scoring

Let x_p denote the L2-normalised embedding of posting p , and let a_k denote the L2-normalised embedding of topic anchor k , for $k = 1, \dots, 31$. The topic-level similarity score is

$$c_{pk} = x_p' a_k. \quad (1)$$

Because both vectors are normalised, c_{pk} is the cosine similarity between posting p and topic anchor k . The primary posting-level LSI score is the maximum similarity across all topic anchors:

$$s_p = \max_{k \in \{1, \dots, 31\}} c_{pk}. \quad (2)$$

The maximum is appropriate because a posting can be substantively sustainability-related through a strong match to one topic even if it does not match the full taxonomy. A climate-mitigation engineering role, for example, should score as sustainability-related because it matches the climate anchor, even if it does not also match biodiversity, labour rights, or governance anchors.

For pillar-level analyses, let $\mathcal{K}_E$, $\mathcal{K}_S$, $\mathcal{K}_G$, and $\mathcal{K}_{Sec}$ denote the environmental, social, governance, and sectoral topic sets. I define the posting-level pillar scores as

$$s_p^m = \max_{k \in \mathcal{K}_m} c_{pk}, \quad m \in \{E, S, G, Sec\}. \quad (3)$$

The maximum rather than the mean avoids mechanically diluting scores for pillars with more topics and keeps the interpretation topic-specific: a posting belongs to a pillar when it strongly matches at least one topic in that pillar.

I also compute alternative posting-level aggregators as robustness measures: the mean of the four pillar maxes, the mean of the top three topic similarities, and the mean of the top five topic similarities.⁶

3.6 Firm-year aggregation

The primary firm-year intensity measure is the mean posting-level score across all postings published by firm i in year t :

$$LSI_{it}^{mean} = \frac{1}{N_{it}} \sum_{p \in P_{it}} s_p, \quad (4)$$

where $N_{it} = |P_{it}|$. This measure captures the average sustainability orientation of the firm's labour demand in a given year. The mean separates sustainability intensity from posting scale: larger firms post more vacancies mechanically, but they do not receive a higher intensity score unless their typical posting is more sustainability-related.

For pillar-level analyses, I analogously define

⁶Across the validation sample, these alternative aggregators correlate with the primary `total_max` score at Spearman correlations above 0.95. I therefore use `total_max` as the primary posting-level score and report the alternatives as sensitivity columns.

$$LSI_{it}^m = \frac{1}{N_{it}} \sum_{p \in P_{it}} s_p^m, \quad m \in \{E, S, G, Sec\}. \quad (5)$$

I also report tail-intensity measures, defined as the share of firm-year postings above selected cross-sectional percentiles of the posting-level score.⁷ The wage-weighted aggregator, developed next, is the version used to construct the dollar-denominated investment flow.

3.7 Wage-weighting and the dollar-denominated investment flow

The intensity measures above are dimensionless. For the intangible-capital interpretation, I require a dollar-denominated investment flow that can be compared with SG&A-based organisation-capital investment in [Peters and Taylor \(2017\)](#). I construct this flow by assigning each posting an occupation-level wage.

For each posting, I assign an external occupation-level annual wage drawn from a national wage-statistics series, deflated to constant 2017 dollars using a consumer-price index. The base year is the midpoint of the 2013–2020 sample window. The wage series, occupation-code mapping, and deflator are described in [Appendix B.3](#).

Let w_p denote the real annual wage assigned to posting p . The wage-weighted mean score is

$$LSI_{it}^{wmean} = \frac{\sum_{p \in P_{it}} s_p w_p}{\sum_{p \in P_{it}} w_p}. \quad (6)$$

This measure preserves the interpretation of sustainability intensity but gives greater weight to higher-wage roles. The dollar-denominated investment flow is

$$I_{it}^{LSI} = \sum_{p \in P_{it}} s_p w_p. \quad (7)$$

This flow is the cosine-weighted wage bill associated with sustainability-oriented labour demand in firm i and year t , expressed in 2017 dollars. It is not an accounting expense observed on the income statement; it is an economically scaled proxy for own-account sustainability labour input. Its advantage is that it is expressed in the same monetary units as SG&A-based intangible investment, allowing comparison with existing intangible-capital stocks.

The wage match covers approximately 95% of postings and 98% of posting-firm-years; coverage details and the residual unmatched cases are documented in [Appendix B.3](#).

3.8 The intangible stock LSI-K

I cumulate the wage-weighted dollar flow into a stock measure, K_{it}^{LSI} , using the perpetual-inventory method. This is the standard convention for intangible-capital stock construction ([Hall, 1990](#); [Corrado et al., 2009](#); [Peters and Taylor, 2017](#)). The law of motion is

$$K_{it}^{LSI} = (1 - \delta)K_{i,t-1}^{LSI} + I_{it}^{LSI}, \quad (8)$$

⁷The tail measures use the 75th, 90th, and 95th percentiles of `total_max`. They are reported as sensitivity measures because they emphasise concentrated sustainability hiring rather than the average orientation of all postings.

where δ is the depreciation rate and I_{it}^{LSI} is the wage-weighted investment flow from Equation 7. The initial condition follows the convention in Hall (1990), as adopted by Peters and Taylor (2017):

$$K_{i,t_0}^{LSI} = \frac{I_{i,t_0}^{LSI}}{g + \delta}, \quad (9)$$

where g is the steady-state growth rate of the investment-flow series. I estimate g as the pooled mean log growth of the wage-weighted dollar flow across firms, which gives approximately $g = 0.014$ per year.

I report results for depreciation rates $\delta \in \{0.15, 0.20, 0.30\}$, matching the range used for organisation-capital and R&D-capital stocks in Peters and Taylor (2017). The headline specification uses $\delta = 0.20$. Results in Section 5 are not materially sensitive to this choice.

3.9 Validation strategy

The validation exercise asks whether LSI behaves like a credible proxy for labour-input sustainability investment. I run five tests on a stratified posting-level sample of 250,000 postings drawn evenly across calendar years. Detailed results are reported in Appendix A.7; this section summarises the design and the main interpretation.

First, I conduct a face-validity check using the top- and bottom-ranked postings by `total_max`. The top-ranked postings are substantive sustainability roles, including climate-risk strategy, EHS management, DEI programmes, sustainability reporting, and corporate ESG analysis. The bottom-ranked postings are routine retail, food-service, and operational roles with no sustainability content. Table 21 reports illustrative examples.

Second, I validate LSI against the O*NET green-occupation flag developed by Dierdorff et al. (2009), which identifies occupations as green-enhanced, green-new-and-emerging, or green-increased-demand. This benchmark is useful for coarse external validation because it is widely used in the green-skills literature (Vona et al., 2018; Consoli et al., 2016). It is not, however, designed to distinguish operational sustainability work from disclosure-oriented sustainability work.

Third, I run a Walk-versus-Talk discrimination test using Chen (2024)’s classification of O*NET green occupations into action-oriented and disclosure-oriented roles. Because LSI is designed to capture sustainability capability rather than generic ESG vocabulary, this is the more construct-relevant benchmark. A successful action-language measure should score Walk roles above Talk roles, conditional on occupation.

Fourth, I compare the curated corpus to the mechanical alt-corpus described in Section 3.2. The alt-corpus performs slightly better on the coarse O*NET green-occupation benchmark, while the curated corpus performs better on the Walk-versus-Talk discrimination task. I interpret this as the expected trade-off between broad green-vocabulary detection and action-language discrimination. The curated corpus sacrifices a small amount of performance on a noisy occupation-level green benchmark in exchange for sharper separation between operational sustainability work and disclosure-oriented sustainability work.

Fifth, I assess internal coherence. I report cross-pillar correlations, rank consistency across

alternative posting-level aggregators, year-by-year trends in the full estimation sample, truncation diagnostics, and robustness to pooling choice. The alternative firm-year aggregators are highly rank-correlated with the primary measure, and the time-series evolution of LSI is smooth across pillars.

Taken together, the validation tests support the interpretation of LSI as a labour-input sustainability measure. The measure passes coarse external validation, performs better on the construct-relevant Walk-versus-Talk task than the mechanical corpus, is robust to alternative aggregators and pooling rules, and produces stable firm-year rankings. The firm-year convergent-validity tests against LSEG ESG, S&P Global ESG, and the [Sautner et al. \(2023\)](#) climate-exposure measure are reported in Section 5.1.

4 Data and sample

The empirical analysis uses one primary data source (LinkUp job postings) and six external firm-year panels for matching, controls, benchmarks, and outcomes (Compustat for firm fundamentals, LSEG ESG and S&P Global ESG for commercial sustainability ratings, the firm-level climate-exposure measure of [Sautner et al. \(2023\)](#), the Reprisk database of corporate sustainability incidents, and the [Kogan et al. \(2017\)](#) extended patent file linked to a green-patent classification). I describe each at the conceptual level here. Vendor-specific implementation detail (database identifiers, field codes, pull procedures, identifier crosswalk pipeline, and variable-construction recipes) is in Appendix B.

4.1 LinkUp postings and text cleaning

The primary data source is the LinkUp posting database, a commercial collection of job postings scraped daily from the careers pages of public and private employers worldwide. Two features motivate the choice over alternative postings vendors such as Lightcast (Burning Glass) and Revelio Labs. First, LinkUp scrapes postings directly from each employer’s company website rather than aggregating from job boards. The direct-from-source scraping captures a single canonical posting per role without cross-board deduplication ambiguity and captures the employer’s own description rather than an aggregator’s standardised summary. Second, the LinkUp company identifier is an employer-entity identifier that is stable over time, which avoids the firm-identifier fragmentation common in aggregator databases. Each posting carries a creation date and a removal date, which lets me attach an unambiguous calendar-time stamp to every posting and aggregate to the firm-quarter frequency for the event-study analysis in Section 5 without dating imputation.

Each posting is lightly cleaned before embedding. The cleaning concatenates the posting title and description (titles often carry the sustainability-specific content while bodies are generic) and applies three normalisation tiers: encoding repair (line-ending normalisation, escaped-unicode expansion, HTML-entity unescaping, mojibake correction, zero-width and control-character removal); structural-noise removal (HTML tags, URLs, emails, phone numbers, postal codes); and whitespace tidying (multi-space collapse, decorative-symbol-line removal, orphan-navigation-line removal such as “Apply Now” and “Share this Job”, blank-line normalisation). The cleaning

is deliberately light. Heavier passes that would strip HRIS labels and boilerplate language are skipped because they risk removing substantively informative text (diversity and environmental-commitment statements often live inside “About the Company” boilerplate, for example). The length filter is minimal: keep any posting with at least one non-whitespace character after cleaning. Approximately 99 percent of raw postings are retained across all years. The full Tier-by-Tier specification is in Appendix B.1.

4.2 Universe funnel and estimation panel

The 2007–2023 LinkUp universe contains 257.9 million posting records across 61,519 unique employer entities. Approximately 22 percent are scrape-failure records for which LinkUp did not capture the full text before the original posting page was removed; the remaining 200.4 million description-bearing postings across 50,291 entities are the population from which I draw the embedding sample. Restricting to the 2013–2020 estimation window (pre-2013 has insufficient coverage to support firm-year aggregation; post-2020 was deferred to bound consortium-draft compute cost) yields 60.7 million postings for 10,232 ever-listed firms, which collectively contribute 42,734 firm-years with at least one description-bearing posting. The ever-listed restriction uses LinkUp’s point-in-time reference file and is what makes the downstream Compustat, CRSP, and ESG-panel merges feasible. No exclusions are applied on country, financial-sector membership, or within-window delisting, to avoid survivorship bias in the forward-looking outcome regressions. Table 1 reports the full funnel from raw LinkUp records to the estimation panel.

Table 1: Sample construction: from the LinkUp universe to the estimation panel

Stage	Postings	Firms	Firm-years
1. LinkUp records (2007-2023, all firms)	257,908,254	61,519	—
2. LinkUp postings with scraped description text	200,410,559	50,291	—
3. Ever-listed firms (2007-2023; estimation universe)	—	10,232	—
4. Postings in estimation universe (2013-2020)	60,664,465	10,232	—
5. Firm-years with at least one posting (LSI base)	—	10,232	42,734
6. Firm-years with wage-matched postings (BLS OEWS)	—	10,082	41,977
7. Firm-years with Compustat financial coverage	—	6,671	26,865
8. Firm-years with LSEG ESG coverage	—	4,839	18,825


Notes. Stage 1 reports the universe of LinkUp posting records over 2007-2023. Stage 2 restricts to records for which LinkUp scraped the underlying job-posting text directly from the company website (77.7 percent of records); these are the postings I can embed.

Figure 1 plots the year-by-year mean of LSI across the full estimation sample. The headline aggregator rises by roughly 0.026 between 2015 and 2020 with the four pillar means tracking each other closely, consistent with the broader rise in sustainability hiring documented in the postings literature (Darendeli et al., 2022).

4.3 External data sources and matching coverage

I match the LinkUp estimation panel to six external firm-year panels through an identifier-crosswalk infrastructure that maps the LinkUp company identifier to gvkey (Compustat), permno (CRSP), primary RIC (LSEG), and the supporting identifier set (ISIN, CUSIP, SEDOL, LEI). The crosswalk covers 7,622 of the 10,232 firms (74.5 percent) with at least one validated gvkey;

Figure 1: Year-by-year mean LSI on the full estimation sample (60.7 million postings, 10,232 firms)



Notes. Each line plots the cross-posting mean of the indicated score for each calendar year. The solid black line is the headline `total_max` aggregator (the maximum cosine similarity across all 31 topic anchors).

the unmatched 25.5 percent are firms either absent from WRDS Global coverage, listed on minor exchanges, or recorded as listed in LinkUp’s reference file without an active record in any standard market-data database during 2013–2025. The crosswalk methodology is documented in Appendix B.4.

Table 2 summarises the coverage of each external panel as a fraction of the LSI firm-year base.

Table 2: Panel coverage by external source

Source	Firm-years	Unique firms	Coverage %
LSI (intensity, primary measure)	42,734	10,232	100.0
LSI wage-weighted (BLS OEWS + CPI-U)	41,977	10,082	98.2
NAICS-2 industry classification	42,342	10,056	99.1
Compustat financial fundamentals	26,865	6,671	62.9
Tobin’s q (mkvalt / ceq, Compustat)	25,761	6,432	60.3
Peters-Taylor intangible-capital stock k_{int}	24,138	5,762	56.5
LSEG (Refinitiv) ESG combined score	18,825	4,839	44.1
S&P Global ESG total score	14,171	4,301	33.2
Sautner et al. climate change exposure	22,924	5,653	53.6
Reprisk ESG incidents (any)	10,292	3,298	24.1
KPSS patents (any)	6,596	1,658	15.4
KPSS green patents (CPC Y02, > 0)	2,164	620	5.1

Notes. Coverage percentages are relative to the LSI base of 42,734 firm-years (10,232 firms over 2013–2020). Each external panel is merged onto the LSI base using the identifier crosswalks summarised in Section B.4: Compustat via gvkey, LSEG ESG via RIC, S&P Global ESG via gvkey, Reprisk via ISIN, KPSS patents via permno (US-listed firms only), and Sautner et al. via gvkey.

The six external panels are as follows. **Compustat** (financial fundamentals: assets, sales, employment, leverage, profitability, capital expenditure, R&D, and the Peters-Taylor knowledge-capital and organisation-capital stocks) covers 62.9 percent of LSI firm-years and is the binding constraint for any regression combining LSI with accounting controls. The **LSEG combined ESG score** (the headline commercial ESG benchmark) covers 44.1 percent. The **S&P Global total ESG score and Environmental, Social, and Economic Governance pillar scores** from the Corporate Sustainability Assessment cover 33.2 percent. The **climate-change exposure measure** of Sautner et al. (2023), decomposed into total, opportunity, regulatory, and physical components, covers 53.6 percent. The **Reprisk** database of corporate ESG incidents, with severity, reach, novelty, and pillar tags per incident, covers 24.1 percent. The **KPSS extended patent file** (Kogan et al., 2017) matched to two green-patent classifications (the CPC Y02 climate-mitigation family used by Cohen et al. (2026) and Darendeli et al. (2022), and the broader OECD ENV-TECH 2023 classification) covers 15.4 percent, necessarily restricted to US-listed patenting firms. Wage-weighting for the dollar-denominated LSI flow uses the U.S. Bureau of Labor Statistics OEWS national mean annual wage series at the 6-digit Standard Occupational Classification level; approximately 98 percent of posting-firm-years receive a wage match.

4.4 Variable definitions and descriptive statistics

The unit of analysis is the firm-year. The headline measure is $LSI_{i,t}$, the firm-year mean of posting-level cosine similarities defined in Section 3.6, with the pillar means $E_{i,t}$, $S_{i,t}$, $G_{i,t}$, $Sec_{i,t}$ as the four pillar-specific aggregators. The wage-weighted dollar-denominated flow $I_{i,t}$ from Section 3.7 feeds the perpetual-inventory stock $LSI-K_{i,t}$ used in the investment-q analysis of Section 5.7. Firm-fundamental controls are the standard finance set: $\ln AT$ (log total assets), leverage (long-plus-short-term debt over assets), return on assets, R&D intensity, capital-expenditure intensity, and $\ln(1 + age)$. Variable-by-variable construction recipes are in Appendix B.5. Continuous regressors are winsorised at the first and ninety-ninth percentiles within the estimation panel.

Table 3 reports descriptive statistics for the headline LSI measures, the firm-fundamental controls, and the external benchmarks and outcomes.

Three patterns shape the empirical strategy. First, the posting count per firm-year is heavily right-skewed: the median firm-year has 116 postings, the 75th percentile has 518, and the 99th percentile exceeds 24,000, reflecting very large multi-establishment employers (Amazon, Walmart, large healthcare systems). The mean-aggregator firm-year LSI is invariant to this scale effect because it averages within firm-year; the wage-weighted dollar flow scales with both posting volume and wages and is heavy-tailed by construction. I report both throughout Section 5 and use the mean as the primary aggregator. Second, the headline aggregator centres around 0.544 with a cross-firm-year standard deviation of 0.037; the four pillar means lie within a 0.04 band of each other but are not directly comparable across pillars in absolute terms because the underlying topic vocabulary differs in baseline prevalence. The within-pillar variation across firm-years is what drives the regression results. Third, the patent and Reprisk-incident outcomes are heavily right-skewed and zero-inflated (median firm-year: 0 green patents, 3 Reprisk incidents; maximum: several hundred green patents and 800-plus Reprisk incidents for the largest firms). These

Table 3: Descriptive statistics of key firm-year variables

Variable	N	Mean	SD	p25	Median	p75
<i>Panel A. LSI measures</i>						
LSI total_max (mean cosine)	42,734	0.544	0.037	0.522	0.548	0.569
LSI dollar flow (2017 USD, \$M)	41,977	36.65	190.3	0.83	3.71	16.13
E pillar mean	42,734	0.506	0.039	0.483	0.511	0.533
S pillar mean	42,734	0.536	0.037	0.515	0.540	0.561
G pillar mean	42,734	0.518	0.041	0.494	0.523	0.546
Sec pillar mean	42,734	0.495	0.038	0.472	0.498	0.521
<i>Panel B. Firm fundamentals (Compustat)</i>						
Total assets (\$M)	26,865	34,688	156,045	740	2,980	12,038
Sales (\$M)	26,851	12,661	40,201	399	1,645	6,520
Employees (thousands)	25,775	33.77	116.23	1.10	4.90	20.10
ln(Total assets)	26,856	8.03	2.17	6.61	8.00	9.40
Leverage = (DLTT + DLC) / AT	26,856	0.273	0.226	0.086	0.240	0.403
ROA = NI / AT	24,413	-0.009	0.163	-0.008	0.024	0.063
R&D intensity = XRD / AT	26,856	0.041	0.087	0.000	0.000	0.039
Capex intensity = CAPX / AT	26,767	0.034	0.038	0.009	0.023	0.047
ln(1 + firm age)	27,097	1.46	0.57	1.10	1.61	1.95
Tobin's q (standard)	25,761	2.25	60.7	0.76	1.23	2.09
Peters-Taylor q_{tot}	20,351	2.49	48.9	0.53	1.00	1.88
Peters-Taylor k_{int} (\$M)	24,138	8,234	24,366	225	883	3,993
<i>Panel C. External sustainability measures and outcomes</i>						
LSEG ESG combined score (0-100)	18,825	44.57	19.29	29.29	42.87	58.78
S&P Global ESG total score (0-100)	14,171	36.60	20.11	21.00	31.50	47.50
Reprisk incidents (count)	10,292	10.19	20.84	1.00	3.00	8.00
Patents per firm-year (count)	6,596	87.25	318.94	3.00	9.00	37.00
Green patents Y02 (count)	6,596	6.01	36.97	0.00	0.00	1.00
Green patents OECD ENV-TECH (count)	6,596	6.74	44.45	0.00	0.00	1.00
Sautner climate-change exposure	22,924	0.00099	0.00248	0.00012	0.00031	0.00074

Notes. Sample: 42,734 firm-years over 2013-2020 covering 10,232 firms. The number of observations varies across rows because not every firm has Compustat, LSEG, S&P Global, Reprisk, KPSS, or Sautner coverage in every year (see Table 2).

distributional features motivate the pseudo-Poisson maximum-likelihood specification (Silva and Tenreyro, 2006) that I use in the Reprisk-incident and green-patent analyses (Sections 5.4 and 5.5) rather than log-linear OLS.

Table 4 reports the industry distribution at the 2-digit NAICS level.

Table 4: Industry distribution of the estimation panel (top 12 NAICS-2 sectors)

NAICS-2	Sector	Firm-years	Firms	Share	Mean LSI
33	Manufacturing (durable, machinery, transport, electronics)	7,588	1,757	17.9%	0.555
51	Information (publishing, software, telecom, data)	6,972	1,598	16.5%	0.546
52	Finance and insurance	5,696	1,332	13.5%	0.541
54	Professional, scientific, technical services	3,625	887	8.6%	0.550
32	Manufacturing (non-durable, chemicals, plastics, paper)	3,606	934	8.5%	0.556
62	Health care and social assistance	1,609	390	3.8%	0.526
45	Retail trade (miscellaneous, sporting goods, general)	1,483	308	3.5%	0.519
42	Wholesale trade	1,355	317	3.2%	0.537
31	Manufacturing (food, beverage, textile, apparel)	1,273	298	3.0%	0.536
56	Administrative, support, waste management services	1,204	304	2.8%	0.545
72	Accommodation and food services	1,181	256	2.8%	0.505
21	Mining, quarrying, oil and gas extraction	1,100	270	2.6%	0.556
Top 12 total		36,692	8,651	86.7%	—
All 22 NAICS-2 sectors total		42,342	10,056	99.1%	0.544

Notes. Sectors are assigned at the firm level from LinkUp’s point-in-time `naics_code` field, truncated to the 2-digit NAICS-2 parent. Mean LSI is computed across all firm-years in the sector.

The estimation panel concentrates in three large sectors—manufacturing (NAICS 31–33), information (NAICS 51), and finance (NAICS 52)—which together account for 59.4 percent of firm-years; the top twelve sectors account for 86.7 percent. Cross-sectoral variation in mean LSI is modest (a 0.05-point range across the top twelve sectors), consistent with sustainability orientation being a within-sector firm-level characteristic rather than a sector-level fixed effect. This pattern motivates the industry-by-year fixed-effects specifications used throughout Section 5: industry-year fixed effects absorb sectoral and time-varying sectoral differences while preserving the within-industry-year firm-level variation in LSI that the construct captures. Firm-and-year fixed-effects specifications are reported as a within-firm robustness column where panel length supports them.

5 Empirical results

This section reports the empirical results across seven analyses. Each analysis is designed to answer a specific question about what LSI captures and what it predicts, and together they discharge the contribution claims set out in Section 1: that LSI converges with established firm-level sustainability measures at magnitudes small but reliably positive and well below the rater-to-rater range, consistent with non-redundant content (Section 5.1); that it tracks firm-level climate exposure in the expected direction (Section 5.2); that it carries information about subsequent ESG-rating changes that the ratings themselves do not contain (Section 5.3); that it is honestly bounded in what it captures about realised sustainability outcomes such as Reprisk incidents (Section 5.4); that it predicts substantive green-patenting output beyond what commercial ESG ratings explain, replicating the disconnect documented by Cohen et al. (2026) and extending its resolution to a labour-input measure following Darendeli et al. (2022)

(Section 5.5); that it behaves like an intangible-investment flow in its future-profitability and future-q signature (Section 5.6); and that the corresponding stock LSI-K is empirically incremental to the Peters-Taylor knowledge-capital and organisation-capital decomposition in the investment-q regression (Section 5.7).

Specification conventions. Unless stated otherwise, all regressions in this section use the firm-year panel constructed in Section 4 restricted to 2013-2020. The headline fixed-effect structure is industry-by-year (NAICS-2 $\times$ year), which absorbs sector-time trends while preserving within-industry firm-year variation in LSI. I also report firm-and-year fixed-effect specifications where panel length permits, as a within-firm robustness column. Standard errors are clustered at the firm level throughout, following the convention in the postings-and-investment literature (Babina et al., 2024; Darendeli et al., 2022). Control variables in the primary specification are the standard finance set: $\ln(AT)$, leverage, ROA, R&D intensity, capex intensity, and $\ln(1 + \text{age})$. Extended controls (Tobin’s q, $\ln(EMP)$, book-to-market, tangibility, advertising intensity) are added in robustness columns where they sharpen identification of the LSI coefficient, particularly in the rating-revision analysis (Section 5.3) where Tobin’s q controls for growth-opportunity proxying.

Estimation conventions. For continuous outcome variables I use OLS with the fixed-effect and clustering conventions above, implemented in R with the `fixest` package. For count outcome variables (Reprisk incident counts in Section 5.4, patent counts in Section 5.5), the primary specification is pseudo-Poisson maximum likelihood (PPML) following Silva and Tenreiro (2006), with $\log(1 + y)$ OLS reported as a robustness column. The PPML choice is motivated by the zero-inflated right-skewed distribution of the count outcomes (Section 4.4). For the first-incident event study (Section 5.4) I use the Sun-Abraham estimator with cohort-specific event-time coefficients to handle staggered treatment timing without the two-way fixed-effects bias in standard event-study TWFE specifications (Sun and Abraham, 2021).

Aggregator conventions. The headline firm-year aggregator is the mean of posting-level `total_max` across the firm’s postings in the calendar year, denoted $LSI_{i,t}$. Robustness columns report the wage-weighted aggregator $LSI_{i,t}^w$ from Section 3.7, the four pillar-specific means $E_{i,t}$, $S_{i,t}$, $G_{i,t}$, $Sec_{i,t}$, and the alt-corpus aggregator from Section 3.2 where the head-to-head with the curated corpus is informative. The intangible-stock variable $LSI-K_{i,t}$ is the perpetual-inventory cumulation of the wage-weighted dollar flow at depreciation rate $\delta = 0.20$, with $\delta \in \{0.15, 0.30\}$ reported as sensitivity (Appendix A.9).

5.1 Convergent validity with commercial ESG ratings and text-based climate exposure

Before any downstream contribution can be taken seriously, LSI must co-move with established firm-level sustainability benchmarks at plausible magnitudes. I use three: the LSEG (formerly Refinitiv) and S&P Global commercial ratings as the disclosure-anchored benchmarks, and the

Sautner et al. (2023) firm-level climate-exposure measure as the leading text-based benchmark and the natural external proxy for the environmental pillar.

The rating-divergence literature (Berg et al., 2022; Avramov et al., 2022; Christensen et al., 2022) documents pairwise Spearman correlations of 0.38–0.71 across the leading commercial raters, with roughly half of the divergence attributable to scope and weight choices. A non-redundant measure of a related but distinct construct slice should converge with the established raters at magnitudes meaningfully below this rater-to-rater range; convergence above the range would signal redundancy, the failure mode the new-measure literature is most often accused of (Raghunandan and Rajgopal, 2024). The expected pattern is a small but reliably positive correlation with the commercial ratings and a meaningfully stronger correlation with the text-based benchmark that shares LSI’s construction logic.

LSEG combined ESG. Table 5 regresses the LSEG combined ESG score on `lsi_mean` across three fixed-effect variants. The coefficient is stable across the FE ladder (pooled, NAICS-2 industry, industry-by-year) and significant at the 1 percent level throughout, on a sample of 16,171 firm-years with fixed effects. The stability is informative: most of the cross-sectional LSI-rating co-movement survives within-industry-year demeaning, so the convergence is not an artefact of industry-level sustainability differences or year-level rating trends. The economic magnitude is bounded; given an unconditional LSI standard deviation of 0.037 on the cosine scale (Section 4.4) and an LSEG-sample standard deviation of approximately 19 points, a one-SD shock to LSI moves the LSEG score by about 0.10 SDs, the order of magnitude expected from a measure that shares partial construct content with the headline rating but is not a substitute for it.⁸

S&P Global ESG total. Table 22 reports the analogous regressions against the S&P Global Corporate Sustainability Assessment total ESG score. The coefficient pattern mirrors the LSEG results and is significant at the 1 percent level across all three FE variants on 12,165 firm-years; a one-SD increase in LSI corresponds to a 1.85-point increase in the S&P Global total score, economically indistinguishable from the LSEG coefficient and reassuring on cross-rater stability. The smaller sample reflects S&P Global’s narrower assessed-firm coverage (33.2 vs. 44.1 percent of LSI firm-years per Table 2).

Table 23 drills into the S&P Global pillar decomposition. The environmental, social, and governance LSI pillar means (Section 3.6) all load positively and significantly on their S&P Global counterparts, with the G-pillar coefficient the smallest of the three. I read this as expected rather than anomalous: S&P Global’s Economic Governance pillar weighs board structure, ownership, and code compliance heavily, and these attributes are observed primarily at the firm-disclosure level rather than at the hiring level.

Sautner et al. climate-change exposure. Table 24 regresses `E_mean` on the four Sautner et al. (2023) components. All four coefficients are positive and significant at the 1 percent

⁸I do not report parallel LSEG pillar regressions because LSEG retired the legacy Refinitiv pillar methodology in 2022 and the current Workspace pull contains pre-2021 values only for the combined score, leaving too few firm-year observations on the 2013–2020 window to support reliable pillar estimation. The pillar-decomposition test is discharged through the S&P Global pillar regressions below, which have full pre-2021 coverage.

Table 5: LSI converges with the LSEG combined ESG score

Dependent Variable:	lseg_esg		
Model:	(1)	(2)	(3)
<i>Variables</i>			
lsi_mean	60.37*** (7.072)	51.55*** (7.381)	51.82*** (7.402)
ln_at	4.177*** (0.1717)	4.469*** (0.1822)	4.385*** (0.1835)
leverage	2.807** (1.163)	-0.1198 (1.249)	-0.3522 (1.251)
roa	24.14*** (1.955)	21.20*** (2.006)	19.77*** (2.015)
rd_int	22.55*** (3.889)	11.25*** (4.303)	11.58*** (4.320)
capx_int	19.45*** (7.036)	-15.26* (7.988)	-14.98* (8.055)
ln_age	4.180*** (0.3308)	4.427*** (0.3253)	12.77*** (1.392)
<i>Fixed-effects</i>			
naics2		Yes	Yes
year			Yes
<i>Fit statistics</i>			
Within R ²		0.24420	0.24782
R ²	0.22545	0.26344	0.26741
Observations	16,273	16,171	16,171

Clustered (company_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. The dependent variable is the LSEG combined ESG score, a continuous firm-year rating on the 0-100 scale, pulled from LSEG Workspace and described in Section 4.3. The regressor of interest is `lsi_mean`, the firm-year mean of posting-level `total_max` described in Section 3.6.

level; the climate-total and opportunity components carry the strongest fit, and the physical-exposure component the weakest, consistent with the documented fact that physical exposure is concentrated in a small subset of firms and is hard to detect from earnings-call language alone (Sautner et al., 2023).

Pooled and within-industry correlations. Table 6 consolidates the convergence evidence into ten pooled and within-industry Spearman/Pearson correlations (the within-industry correlation is computed after subtracting the 2-digit NAICS industry mean from each side). LSI total correlates with LSEG at pooled Spearman 0.10 (within-industry 0.09) and with S&P Global at pooled Spearman 0.06 (within-industry 0.04). The S&P pillar correlations are 0.04–0.07 pooled and 0.02–0.05 within-industry. The LSI–Sautner E-pillar correlation is substantially stronger: pooled Spearman 0.244 with `cc_expo` and within-industry 0.115.

Table 6: Pooled and within-industry correlations of LSI with external sustainability measures

Pair	N	Pooled		Within-industry	
		Spearman	Pearson	Spearman	Pearson
LSI total ↔ LSEG ESG	18,704	0.098	0.091	0.091	0.085
LSI total ↔ S&P ESG	14,060	0.059	0.066	0.036	0.045
E_mean ↔ LSEG E	0	—	—	—	—
E_mean ↔ S&P E	14,060	0.065	0.057	0.024	0.025
S_mean ↔ LSEG S	0	—	—	—	—
S_mean ↔ S&P S	14,060	0.042	0.050	0.023	0.029
G_mean ↔ LSEG G	0	—	—	—	—
G_mean ↔ S&P G	14,060	0.063	0.069	0.053	0.057
E_mean ↔ Sautner <code>cc_expo</code>	22,791	0.244	0.208	0.115	0.092
E_mean ↔ Sautner <code>rg_expo</code>	22,791	0.165	0.143	0.049	0.063

Notes. Cell entries are pairwise Spearman and Pearson correlations on the unconditional sample (*Pooled* columns) and on residuals after firm-year demeaning within 2-digit NAICS industry (*Within-industry* columns); *N* is the count of non-missing firm-year observations for the pair. The three LSEG pillar rows report no observations because LSEG retired the legacy Refinitiv pillar scores in 2022 and the current pull contains no pre-2021 pillar values; see Section 4.3 for the data-coverage discussion and the planned legacy-pillar back-fill.

These 0.05–0.10 commercial-rating correlations are an order of magnitude below Berg et al. (2022)’s 0.38–0.71 rater-pair range and well below Avramov et al. (2022)’s rater-pair distribution centred near 0.45, the expected pattern for a measure capturing a related but non-redundant slice of the construct: LSI reads hiring activity while the commercial raters read disclosed practices and outcomes. The lower magnitude with S&P Global than LSEG is consistent with S&P Global’s CSA placing heavier weight on disclosed metrics and board-level governance (Berg et al., 2022), both more distant from hiring than LSEG’s broader topical coverage.

The Sautner gap is the substantively informative fact. Two features align Sautner with LSI’s construction logic in ways the commercial ratings do not: both are text-based and built from unstructured prose rather than structured disclosure items, and both capture firm-level engagement with sustainability content (LSI through hiring decisions, Sautner through what management chooses to discuss) rather than disclosure quality or rating-assessor judgment. The stronger Sautner convergence is the predicted pattern under the claim that LSI captures action and exposure rather than disclosure.

Takeaway. LSI converges with the leading sustainability benchmarks at magnitudes consistent with capturing related but non-redundant content: reliably positive, significant at the 1 percent level under industry-by-year FE, and bounded at roughly 0.10 pooled Spearman against the commercial ratings; meaningfully stronger at pooled Spearman 0.24 against the [Sautner et al. \(2023\)](#) text-based environmental benchmark. The natural objection that the residual variation is noise rather than non-redundant signal is the substantive question for the downstream analyses: Section 5.3 delivers the cross-sectional lead-indicator pattern on subsequent rating revisions, Section 5.5 the incremental contribution to the [Cohen et al. \(2026\)](#) green-patent equation, and Section 5.7 the incremental loading in the Peters-Taylor investment-q decomposition. The convergence magnitudes here motivate; those three analyses discharge.

5.2 Climate exposure as a forcing variable for sustainability investment

Convergence with commercial ratings (Section 5.1) does not, by itself, demonstrate that LSI responds to the underlying economic forcing variable that should drive firm-level sustainability investment. The natural such forcing variable is firm-level climate exposure: firms more exposed to climate change through product markets, regulation, or physical operations have stronger ex-ante incentives to build internal sustainability capability ([Krueger et al., 2020](#); [Engle et al., 2020](#); [Pankratz et al., 2023](#)). A labour-input measure that did not respond to climate exposure in the expected direction would fail an important construct-validity test.

I use the [Sautner et al. \(2023\)](#) firm-year measure, derived from earnings-call transcripts and decomposed into total (`cc_expo`), opportunity (`op_expo`), regulatory (`rg_expo`), and physical (`ph_expo`) exposure. Sautner is the appropriate external benchmark for three reasons: it is constructed from a different data source (transcripts vs. postings), so co-movement reflects substantive orientation rather than common-vocabulary mechanics; the four-way decomposition lets me ask whether each kind of exposure moves the pillar that should logically respond to it; and the panel covers 53.6 percent of LSI firm-years, sufficient for firm-fixed-effect specifications.

Specification. The headline regression is

$$LSI_{i,t} = \alpha + \sum_{k \in \{cc, op, rg, ph\}} \beta_k \cdot \text{expo}_{i,t-1}^k + \gamma' X_{i,t} + \mu_{j(i),t} + \varepsilon_{i,t}, \quad (10)$$

where $X_{i,t}$ is the standard finance control set (ln AT, leverage, ROA, R&D intensity, capex intensity, $\ln(1 + \text{age})$) and $\mu_{j(i),t}$ is the fixed-effect structure I vary across four columns: pooled, NAICS-2 industry FE, industry-by-year FE, and firm-and-year FE. I lag exposure by one year both because the postings response to a change in earnings-call discourse should take time and because the lag is more conservative against contemporaneous-shock interpretations.

Headline result: cross-sectional loading absorbed by firm fixed effects

Table 7 reports the FE ladder. The coefficient on lagged total climate-change exposure is positive and significant in the cross-section, dropping from 3.22 pooled to roughly 1.5 with industry and industry-by-year FE (both at the 1 percent level). The industry-FE magnitude implies that a one-SD increase in lagged `cc_expo` is associated with an LSI increase of roughly 0.16

cross-sectional standard deviations, modest but well-identified. The pooled-to-industry drop tells me that part of the unconditional correlation is sectoral; the within-industry residual variation explains roughly half of the pooled relationship.

Table 7: LSI on lagged Sautner et al. climate-exposure components

Dependent Variable: Model:	(1)	lsi_mean (2)	(3)	(4)
<i>Variables</i>				
sautner_cc_expo_lag1	3.217*** (0.5730)	1.471*** (0.3902)	1.485*** (0.3912)	0.3519 (0.2971)
sautner_op_expo_lag1	-3.010*** (1.156)	-1.511* (0.8456)	-1.528* (0.8482)	-0.9571* (0.5285)
sautner_rg_expo_lag1	0.2188 (2.213)	1.269 (1.978)	1.042 (1.985)	-0.4986 (1.199)
sautner_ph_expo_lag1	2.125 (2.433)	1.558 (2.241)	1.527 (2.277)	3.248 (2.343)
ln_at	0.0014*** (0.0003)	0.0012*** (0.0003)	0.0012*** (0.0003)	0.0015* (0.0008)
leverage	-0.0083*** (0.0024)	-0.0016 (0.0023)	-0.0016 (0.0023)	-0.0007 (0.0024)
roa	-0.0025 (0.0033)	-0.0005 (0.0032)	0.0009 (0.0032)	0.0028 (0.0023)
rd_int	0.0633*** (0.0067)	0.0314*** (0.0071)	0.0322*** (0.0071)	0.0064 (0.0084)
capx_int	-0.0311** (0.0156)	-0.0218 (0.0147)	-0.0190 (0.0148)	0.0173 (0.0124)
ln_age	0.0083*** (0.0007)	0.0077*** (0.0007)	-0.0065 (0.0044)	0.0079 (0.0077)
<i>Fixed-effects</i>				
naics2		Yes	Yes	
year			Yes	Yes
company_id				Yes
<i>Fit statistics</i>				
Within R ²		0.02072	0.01121	0.00158
R ²	0.05645	0.16307	0.16552	0.83069
Observations	15,701	15,648	15,648	14,942

Clustered (company_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Dependent variable is the firm-year headline aggregator $LSI_{i,t}$, the firm-year mean of posting-level **total_max** (Section 3.6). The four right-hand-side covariates of interest are the one-year lags of the Sautner-van Lent-Vilkov-Zhang firm-level climate-exposure components (Sautner et al., 2023): total climate-change exposure (**cc_expo**), opportunity exposure (**op_expo**), regulatory exposure (**rg_expo**), and physical exposure (**ph_expo**), each measured at the firm-year level from earnings-call transcripts.

The within-firm specification (column 4) tells a different story: the coefficient collapses to 0.35 and loses significance. The null is informative, not a power problem: with 14,942 firm-years and SE 0.30, coefficients above approximately 0.6 would have been detected. The interpretation, which the rest of Section 5 reinforces, is that the cross-sectional and within-firm components do different work. The cross-sectional component (columns 2–3) shows that LSI sorts firms in the way the construct says it should, the firm-level analogue of the asset-pricing patterns in Krueger et al. (2020) and Engle et al. (2020). The within-firm null says that year-to-year LSI variation

within a firm is not driven by year-to-year fluctuations in earnings-call climate discourse: LSI captures a firm-level sustainability-orientation parameter that is largely a fixed feature of the firm, while transient conversational salience is noisier. This is the same cross-sectional-but-not-within-firm pattern that appears in the lead-lag analysis of Section 5.3 and is consistent with intangible investment being sticky and adjusting on long horizons.

The negative opportunity-exposure coefficient

The coefficient on lagged `op_expo` is negative across all four specifications, strongly significant pooled ($\beta = -3.01$) and weakly significant under FE, including within firm ($\beta = -0.96$, $p < 0.10$). Opportunity exposure is the only Sautner component whose effect survives firm-fixed-effect demeaning, and the sign is opposite to what a naive read of the construct suggests.

I read this cautiously. The most defensible interpretation is that firms describing climate change as an opportunity in their earnings calls may substitute that rhetoric for substantive operational hiring; the reading is consistent with the cheap-talk evidence in Bingle et al. (2022, 2024) and with the construct logic of the Sautner decomposition (opportunity language is the most affected by strategic communication incentives, whereas regulatory and physical components capture exposures closer to compulsory disclosure). Competing interpretations I cannot rule out at this stage are a substitution between revenue-generating climate-opportunity activity and the labour-input margin LSI captures, or a residual industry-mix effect within NAICS-2. The proper tests of the substitution interpretation are the green-patent regression in Section 5.5 and the future-profitability regression in Section 5.6.

Regulatory and physical exposure: largely insignificant

The coefficients on lagged `rg_expo` and `ph_expo` are insignificant in every specification, positive in sign under industry FE with SEs roughly 1.5 times the point estimates. Both components are individually much rarer in earnings-call discourse than total and opportunity exposure, and their substantive content is heavily concentrated in specific sectors (utilities and heavy manufacturing for regulatory; agriculture, insurance, coastal real estate for physical), thinning the within-industry identifying variation.

Pillar decomposition

Table 8 splits the dependent variable across the four LSI pillars under the industry-by-year FE specification. Three patterns are informative.

First, the E pillar carries by far the largest loading on lagged total climate exposure ($\beta_E = 2.63$, $p < 0.01$), nearly twice the headline magnitude; the S pillar also loads positively ($\beta_S = 1.18$, $p < 0.01$); the governance pillar is flat. This is the construct-validity pattern I would hope for: environmental hiring (climate-mitigation engineers, EHS officers, energy-efficiency specialists) responds most strongly to firm-level climate exposure, while governance hiring (board oversight, ethics, anti-corruption) is largely orthogonal to climate exposure as Sautner et al. measure it.

Second, the sectoral pillar loads positively and significantly ($\beta_{Sec} = 2.17$, $p < 0.01$) at a magnitude close to the E pillar. This is consistent with the pillar content rather than indicating

Table 8: Per-pillar LSI on lagged Sautner et al. climate-exposure components

Dependent Variables: Model:	E_mean (1)	S_mean (2)	G_mean (3)	Sec_mean (4)
<i>Variables</i>				
sautner_cc_expo_lag1	2.630*** (0.7521)	1.185*** (0.3609)	0.1856 (0.4464)	2.173*** (0.6758)
sautner_op_expo_lag1	-2.945** (1.323)	-1.551* (0.8085)	0.3434 (0.9407)	-1.535 (1.258)
sautner_rg_expo_lag1	0.9511 (2.488)	1.463 (1.923)	2.565 (2.242)	2.655 (2.459)
sautner_ph_expo_lag1	2.208 (2.628)	1.042 (2.477)	-1.394 (2.637)	2.688 (2.173)
ln_at	0.0006** (0.0003)	0.0012*** (0.0003)	0.0019*** (0.0003)	0.0015*** (0.0003)
leverage	0.0014 (0.0023)	-0.0022 (0.0023)	-0.0050** (0.0025)	-0.0009 (0.0023)
roa	-0.0005 (0.0033)	0.0025 (0.0032)	-0.0014 (0.0035)	0.0046 (0.0033)
rd_int	0.0178** (0.0074)	0.0282*** (0.0070)	0.0726*** (0.0079)	0.0286*** (0.0073)
capx_int	0.0163 (0.0153)	-0.0176 (0.0145)	-0.0604*** (0.0166)	0.0092 (0.0153)
ln_age	-0.0050 (0.0045)	-0.0066 (0.0045)	-0.0099** (0.0046)	-0.0090** (0.0046)
<i>Fixed-effects</i>				
naics2	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Within R ²	0.01243	0.00928	0.02987	0.02012
R ²	0.24498	0.14249	0.18793	0.21119
Observations	15,648	15,648	15,648	15,648
<i>Clustered (company_id) standard-errors in parentheses</i>				
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>				

Notes. Dependent variables are the firm-year pillar-mean aggregators $E_{i,t}$, $S_{i,t}$, $G_{i,t}$, and $Sec_{i,t}$, each the firm-year mean of the posting-level per-pillar maxes from Section 3.5. Right-hand-side specification matches Table 7 with NAICS-2-by-year fixed effects.

spillover: responsible finance and green buildings are mechanically climate-adjacent, and the within-pillar topic anchors are distinct from the E-pillar anchors (within-pillar cosine 0.71 in Table 60).

Third, the negative opportunity loading is concentrated in the E and S pillars ($\beta_E^{op} = -2.94$, $\beta_S^{op} = -1.55$) and absent from governance ($\beta_G^{op} = +0.34$). This is the pillar pattern the strategic-communication interpretation predicts: substitution between earnings-call talk and operational hiring would mechanically show up in the E and S pillars, while governance hiring responds to disclosure mandates rather than to voluntary strategic framing. The pattern strengthens the cautious interpretation above rather than weakening it.

The regulatory and physical components remain insignificant across all four pillar regressions; the substantive content load of climate exposure is carried by the `cc_expo` aggregate.

Takeaway

LSI tracks firm-level climate exposure in the expected direction across firms, with the strongest loading concentrated on the environmental pillar where construct validity predicts it. The within-firm null is informative rather than fatal: LSI captures a firm-level orientation parameter that correlates with permanent climate exposure, while year-to-year LSI variation within a firm is not driven by year-to-year fluctuations in earnings-call discourse, a pattern that recurs in Section 5.3. The negative opportunity coefficient is reported transparently and interpreted cautiously, with the proper test deferred to the green-patent and future-profitability analyses below. Taken together, the climate-exposure evidence discharges the second contribution claim from Section 1: LSI responds to the firm-level economic forcing variable the climate-finance literature has identified as the relevant driver of corporate sustainability behaviour.

5.3 LSI as a leading indicator of subsequent ESG-rating revisions

This section asks whether LSI carries information about subsequent ESG-rating revisions that the ratings themselves do not contain at the same point in time. Three tests organise the analysis. Section 5.3.1 decomposes the variance of LSI and partitions the LSI-and-rating panel into four divergence quadrants. Section 5.3.2 reports the headline temporal-precedence result. Section 5.3.3 reports realised outcomes by quadrant.

This is the paper’s most distinctive empirical contribution because it is where LSI does something the commercial ratings do not: it picks up persistent firm-level sustainability orientation that the ratings progressively recognise over the next one to three years. The literature has documented rating divergence (Berg et al., 2022; Christensen et al., 2022; Avramov et al., 2022) and rating sluggishness as a qualitative property, but no prior measure I am aware of formalises the lead-indicator pattern at the firm-year level on a labour-input source. The cross-sectional-versus-within-firm asymmetry in the precedence result is genuine; I report it honestly and develop the interpretation below.

5.3.1 C1: variance decomposition and divergence quadrants

Where the variation in LSI lives. The first descriptive question is where the variation sits. I compute the variance decomposition on the 42,342 firm-year panel with non-missing LSI and NAICS-2. Table 25 reports: the firm component carries 80.5% of the total sum of squares, the industry component 11.7%, and the year component 1.1%. The qualitative ordering is unambiguous: LSI is overwhelmingly a between-firm-within-industry-year measure. The 80.5% firm share is the empirical analog of the design choice (Sections 3.6–3.7) to build the firm-year measure from posting-level content rather than from industry-level imputations, and it is what makes the industry-by-year FE specification absorb only a modest share of LSI variation. The 11.7% industry share motivates the always-on industry FE throughout Section 5.

The LSI-by-rating quadrant split. The 80.5% firm-share result means there is substantial within-industry-year cross-firm variation in LSI to align with or differ from the commercial rating. I sort firm-years into four quadrants on the joint within-panel-standardised distribution: HL_HR, HL_LR, LL_HR, LL_LR. The diagonals are agreement cells; the off-diagonals are divergence cells. The split is constructed separately for LSEG (Table 9) and S&P Global (Table 26) because the two ratings differ in scale, coverage, and methodology.

Table 9: Firm characteristics by LSI-vs-rating quadrant, LSEG combined ESG score

Quadrant	N	Mean LSI	Mean rating	Mean assets	Mean ROA	Mean R&D	Mean leverage
HL_HR (high LSI, high rating)	5,378	0.573	61.86	64,174	0.044	0.026	0.282
HL_LR (high LSI, low rating)	5,154	0.571	29.54	36,020	−0.016	0.052	0.267
LL_HR (low LSI, high rating)	3,508	0.516	61.29	46,740	0.047	0.014	0.299
LL_LR (low LSI, low rating)	4,785	0.514	29.07	49,940	0.017	0.024	0.275

Notes. Quadrants are defined on the firm-year panel with non-missing LSI and non-missing LSEG combined ESG score. Each firm-year is assigned to one of four quadrants by the signs of its within-panel standardised LSI and standardised rating: HL_HR is above-mean on both; HL_LR is above-mean on LSI and below-mean on the rating; LL_HR is below-mean on LSI and above-mean on the rating; LL_LR is below-mean on both.

All four cells are well populated. Under the LSEG rating, each cell carries at least 18% of the 18,825-firm-year merged panel; under S&P Global each carries at least 16% of 14,171. Roughly four firm-years in ten in the LSEG panel and five in ten in S&P Global fall into a divergence cell. This is the aggregate-confusion finding of Berg et al. (2022) reappearing against the LSI labour-input measure rather than across rater pairs.

The divergence cells differ systematically from the agreement cells on observables. Under the LSEG split, HL_HR firms are the largest by mean assets and HL_LR firms the smallest; HL_LR has the lowest mean ROA (the only negative cell mean) and the highest R&D intensity, suggesting that high-LSI-but-low-rating firms are growth-investing rather than currently profitable, and that the rating may be downweighting substantive sustainability work because the firm has not yet built the disclosure infrastructure or financial track record the rating rewards. The same asset-size ordering holds under S&P Global.

Which firm characteristics drive absolute divergence. To formalise this, I regress the absolute standardised divergence $|\text{LSI}_{i,t}^{\text{std}} - \text{rating}_{i,t}^{\text{std}}|$ on the standard controls with NAICS-2 and year FE. The two coefficients sign-stable and significant across both providers are ROA

and firm age, both negative: LSI and the commercial ratings agree most for established, profitable firms and disagree more for younger, less profitable firms, a pattern consistent with the rating-versus-action gap in [Raghunandan and Rajgopal \(2024\)](#) being concentrated where sustainability orientation has not yet been validated by profitability or a long disclosure track record. The $\ln(\text{AT})$ coefficient flips sign across providers (negative for LSEG, positive for S&P Global), plausibly explained by different coverage tilts; leverage is insignificant under both. The within- R^2 is small (0.016 LSEG, 0.006 S&P), so neither coefficient should be over-interpreted.

5.3.2 C2: temporal precedence (headline)

The headline result is that LSI carries information about subsequent ESG-rating changes that the ratings themselves do not contain. I regress the k -year-ahead change in the rating, $\Delta \text{rating}_{i,t+k} = \text{rating}_{i,t+k} - \text{rating}_{i,t}$, on contemporaneous LSI and the contemporaneous rating level, conditioning on the standard controls and industry-by-year FE. Including the contemporaneous rating isolates the LSI lead-indicator content: the LSI coefficient picks up the increment in subsequent rating change associated with one unit of LSI holding the current rating fixed. I run this at three horizons ($k = 1, 2, 3$) against two providers (LSEG, S&P Global), giving six headline coefficients.

Table 10: Temporal precedence: LSI predicts subsequent changes in the LSEG combined ESG score at horizons $k = 1, 2, 3$

Dependent variable	ΔLSEG_{t+1}	ΔLSEG_{t+2}	ΔLSEG_{t+3}
Model	(1)	(2)	(3)
<i>Variables</i>			
LSI mean ($\text{LSI}_{i,t}$)	9.62*** (2.36)	20.00*** (4.07)	25.84*** (5.96)
LSEG ESG (t)	-0.157*** (0.006)	-0.227*** (0.009)	-0.290*** (0.012)
$\ln(\text{AT})$	0.442*** (0.053)	0.443*** (0.091)	0.457*** (0.130)
Leverage	0.533 (0.382)	1.610** (0.678)	2.166** (1.025)
ROA	4.356*** (0.884)	7.680*** (1.332)	8.131*** (1.947)
R&D intensity	0.187 (1.482)	1.003 (2.505)	3.138 (3.721)
Capex intensity	-4.220* (2.525)	-5.596 (4.500)	-14.661** (6.585)
$\ln(1 + \text{age})$	-1.804* (0.925)	-2.763* (1.436)	-5.716** (2.245)
<i>Fixed effects</i>			
NAICS-2 $\times$ Year	Yes	Yes	Yes
Observations	12,312	9,090	6,393
R^2	0.090	0.129	0.164
R^2 within	0.081	0.118	0.151

Notes. Each column reports an OLS regression of the change in the LSEG combined ESG score over the next k calendar years ($\Delta \text{LSEG}_{i,t+k} = \text{LSEG}_{i,t+k} - \text{LSEG}_{i,t}$) on contemporaneous LSI, the contemporaneous LSEG rating level, and the standard firm-fundamental control set ($\ln(\text{AT})$, leverage, ROA, R&D intensity, capex intensity, $\ln(1 + \text{age})$). The fixed-effect structure is NAICS-2 by year, which absorbs sector-time trends and the mechanical mean-reversion in rating changes that varies with sector-year.

Table 10 reports the LSEG result. The LSI coefficient is positive and significant at the

1% level at all three horizons, rising monotonically from 9.62 ($k = 1$) to 25.84 ($k = 3$). The economic magnitude is sizeable: a one-SD increase in LSI is associated with a roughly 0.36-point increase in the next-year LSEG combined score, scaling to roughly 0.96 points at the three-year horizon against a 0–100 unconditional range. Within- R^2 rises from 0.081 to 0.151 across horizons, consistent with the rating progressively recognising what LSI captured earlier. Table 27 reports the parallel S&P Global result, with coefficients of 6.34, 12.97, and 21.67 across the three horizons, smaller in absolute terms because of scale differences but with the same rank pattern. The cross-sectional precedence result generalises across the two leading commercial ratings.

The fixed-effect progression and the within-firm null. A four-column FE progression at $k = 1$ (Tables 28 and 29) makes the identifying variation explicit. The LSI coefficient is stable across pooled, industry-FE, and industry-by-year specifications (9–12 for LSEG, 8–9 for S&P Global, all significant at 1%); it collapses with firm-and-year FE to 6.35 ($p \approx 0.23$) for LSEG and 0.46 (cleanly indistinguishable from zero) for S&P Global. The within-firm null is therefore not specific to one provider.

I read the cross-sectional-versus-within-firm asymmetry as informative about what LSI is rather than as a failure. LSI is a cross-sectional leading indicator of subsequent rating changes, not a within-firm causal driver of revisions. Across firms, those with higher current LSI are upgraded more over the next one to three years conditional on their current rating; within a firm, year-to-year LSI variation does not push the rating up the following year. This is internally consistent: the within-firm component is mechanically small (the 80.5% firm-share variance decomposition implies the within-firm component is roughly 20% of the cross-sectional SD, and smaller after partialling out year effects), and a one-year fluctuation in posting-level orientation is unlikely to register with analysts who reassess annually and weight long-run disclosure infrastructure heavily. The persistent cross-sectional component is what analysts are progressively recognising, not the year-to-year noise. This aligns with the qualitative pictures in Berg et al. (2022), Christensen et al. (2022), and Avramov et al. (2022); what I add is a real-time, vendor-independent labour-input measure the rating measurably catches up to.

Extended-control robustness. The natural concern is that LSI proxies for an omitted growth or scale variable. Tables 30 and 31 add Tobin’s q , $\ln(\text{EMP})$, book-to-market, tangibility, and advertising intensity. The cross-sectional industry-by-year LSI coefficient is essentially unchanged (10.67 LSEG, 5.92 S&P Global, both at 1%), and the within-firm column remains null. The q coefficient itself is small, so the LSI loading is not absorbed by the growth-opportunity control; the within-firm null persisting under extended controls rules out the alternative that it is driven by omitted slowly-changing scale variables.

5.3.3 C3: outcomes by quadrant

The C2 coefficients establish that LSI predicts subsequent rating changes. The next question is whether the four quadrants also correspond to systematic differences in realised outcomes downstream. I look at two outcomes: forward Reprisk ESG incidents (count of firm-level ESG news incidents in the k -year-ahead window, by severity), and forward green patents (count of

KPSS patents classified under the CPC Y02 family). The quadrant-by-outcome tables (Tables 11 and 32) are descriptive rather than identified.

Table 11: Realised outcomes by LSI-vs-rating quadrant, LSEG combined ESG score

Quadrant	N	Reprisk $t + 1$	Reprisk $t + 2$	Reprisk sev. $t + 1$	Green pat. $t + 1$	Non-green pat. $t + 1$
HL_HR (high LSI, high rating)	5,378	11.66	11.70	35.45	12.08	164.82
HL_LR (high LSI, low rating)	5,154	11.29	11.18	34.15	5.09	59.36
LL_HR (low LSI, high rating)	3,508	9.89	10.05	30.18	5.13	141.97
LL_LR (low LSI, low rating)	4,785	13.60	13.99	40.63	4.23	84.54

Notes. The four-quadrant assignment is identical to Table 9. Each cell reports the firm-year mean of the outcome variable across the quadrant: Reprisk $t + k$ is the count of Reprisk ESG incidents flagged for the firm in the k -year forward window; Reprisk severity $t + 1$ is the severity-weighted incident index (severity $\times$ reach $\times$ novelty summed across incidents) one year ahead; green patents $t + 1$ is the count of KPSS patents issued in $t + 1$ that are classified as green under the CPC Y02 family; non-green patents $t + 1$ is the residual count of non-green patents in $t + 1$.

The cleanest pattern is on green patents. Under LSEG, HL_HR firms produce a mean of 12.08 green patents in the year ahead against 5.09, 5.13, and 4.23 for the other three cells; the two divergence cells produce nearly identical means, suggesting that LSI and the rating each contribute partially and that the joint signal is what discriminates HL_HR. Under S&P Global the same story appears with sharper magnitudes (22.90 HL_HR vs. 1.88 LL_LR, a 12-to-1 ratio), though LL_HR (12.79) sits above HL_LR (3.04), the opposite of the LSEG ordering. The combined picture is that the LSI-and-rating combination discriminates green-patenting firms most sharply, with relative weight depending on provider, the cross-rater divergence of Berg et al. (2022) reappearing on the outcome margin. I take this as an argument for using LSI as a complement to rather than a substitute for commercial ratings. The pattern is consistent with the labour-input-to-innovation channel in Darendeli et al. (2022) and the rating-to-innovation disconnect in Cohen et al. (2026).

The Reprisk columns do not behave the way a literal reading would predict. Under LSEG, the highest-incident cell is LL_LR (13.60), not HL_HR (11.66); under S&P Global the two H-rating cells are highest. Neither split produces the clean "high orientation $\rightarrow$ fewer incidents" gradient because Reprisk counts scale strongly with firm size and media coverage, both higher for large rated firms. I discuss this at length in Section 5.4; the quadrant table makes the descriptive limitation visible, while the green-patent column survives the firm-size adjustment and shows the expected gradient.

Takeaway. This section is the paper’s distinctive empirical contribution. LSI carries persistent firm-level sustainability orientation the commercial ratings have not yet recognised, and the ratings progressively recognise it over the next one to three years. The cross-sectional lead-indicator coefficient is positive and significant at every horizon under both providers, survives standard and extended controls, and does not survive within-firm FE, a feature of what LSI is (a between-firm persistent-orientation measure) rather than a failure of the cross-sectional claim. The four-way quadrant split produces economically meaningful divergence cells in both providers, and the joint LSI-and-rating signal most sharply discriminates firms that subsequently produce green-patenting output. The picture aligns with the literature’s qualitative claim that ratings are sluggish trailing indicators and formalises that claim quantitatively at the firm-year level on a labour-input measure independent of, and prior to, the rating itself.

5.4 Realised ESG incidents: a discriminant-validity boundary

This block asks two complementary questions about LSI and realised firm-level sustainability outcomes, operationalised by counts of ESG incidents in Reprisk (Section 4.3). The reactive question: do firms step up sustainability hiring after a salient ESG incident? The forward-looking question: does higher firm-year LSI predict fewer subsequent Reprisk incidents, conditional on commercial ESG ratings and standard finance controls? Both directions yield null results, and I set the nulls out carefully before turning to their construct-validity implications in Section 5.4.4.

5.4.1 Reactive direction: Sun-Abraham event study around the first Reprisk incident

If LSI captures the labour-input flow of sustainability investment and firms respond to incidents by reinforcing the capability whose absence let the incident occur, a salient incident should be followed by an observable uptick in sustainability hiring. The reactive effect should be cleanest at the firm's first incident, where marginal informational content is largest (subsequent incidents are partially anticipated by repeat-offender dynamics; see Glossner (2021)).

I run the test on the firm-quarter LSI panel (2013Q1-2020Q4), identifying the first Reprisk incident as the minimum `incident_date` after matching to LinkUp via primary ISIN; the control cohort is never-treated firms over the same window. The estimator is the Sun-Abraham cohort-by-event-time interacted event study (Sun and Abraham, 2021), implemented via `fixest::sunab()`, which avoids the negative-weight problem of staggered-DiD TWFE (the closely related Callaway and Sant'Anna (2021) estimator would be a natural cross-check). I include firm and calendar-quarter fixed effects, cluster at the firm level, and take $q = -1$ as the reference event-time.

Figure 2 plots the cohort-averaged event-time coefficients over a window of 29 pre-quarters through 21 post-quarters. The series is essentially flat: the largest pre-period t -statistic is -1.94 at $q = -7$ ($p = 0.053$), and the largest post-period t -statistic is $+1.25$ at $q = +10$ ($p = 0.213$), with no monotone post-incident trend.

Table 12: Sun-Abraham overall ATT of the first Reprisk incident on firm-quarter LSI


	Estimate	Std. error	t -stat	p -value
Overall ATT (firm-quarter LSI)	0.000223	0.001613	0.14	0.890

Notes. The overall ATT is the cohort-weighted average of the post-treatment Sun-Abraham coefficients reported in Figure 2, aggregated via the `agg = "ATT"` option in `fixest::sunab()`. The dependent variable is firm-quarter $LSI_{i,q}$, with a sample-period mean of approximately 0.544 and a within-firm standard deviation of approximately 0.020 at the quarter frequency (Section 4.4).

The cohort-aggregated overall ATT in Table 12 is 0.000223 (clustered SE 0.001613, $t = 0.14$, $p = 0.89$). The point estimate is two orders of magnitude smaller than the within-firm-quarter standard deviation of LSI (around 0.020), and the 95% CI of roughly $[-0.0030, +0.0034]$ excludes any economically meaningful reactive response.

The null does not show that firms ignore incidents; it shows that whatever they do does not flow through to the postings stream at quarterly frequency in a way the firm-quarter LSI mean can detect. Several mechanisms are consistent with this: large firms have substantial hiring-cycle inertia; triggered hires may be specialised compliance, audit, or EHS roles concentrated in

Figure 2: Sun-Abraham event-study estimates of firm-quarter LSI around the first Reprisk incident



Notes. Estimates from the Sun-Abraham cohort-by-event-time interacted estimator implemented in `fixest::sunab()`, following [Sun and Abraham \(2021\)](#). The dependent variable is firm-quarter LSI_{*i,q*} (the firm-quarter mean of posting-level `total_max` from Section 3.6).

a single function rather than distributed across the broad role base driving the firm-quarter mean; and any response may lag the incident by quarters or years. What the flat profile is not compatible with is the strong reactive-hiring channel that motivates the test.

5.4.2 Forward-looking direction: PPML horse race of LSI on subsequent incidents

The forward-looking test asks whether firm-year mean LSI carries predictive content for the next-year Reprisk incident count beyond commercial ESG ratings and standard finance controls. I estimate pseudo-Poisson maximum-likelihood (PPML) regressions following [Silva and Tenreyro \(2006\)](#), the appropriate estimator for the zero-inflated right-skewed count distribution (Section 4.4). The dependent variable is `rep_incidents_lead1`. Column (1) includes LSI with finance controls; column (2) adds the LSEG combined ESG score; column (3) adds S&P Global total ESG. NAICS-2-by-year fixed effects throughout; firm-clustered standard errors.

Table 13 reports the main result. The LSI coefficient is positive but insignificant across all three columns ($\beta = 0.505, 0.670, 0.510$ with SEs 0.755, 0.745, 0.775). The sign opposes the "LSI predicts fewer incidents" hypothesis but is too noisy to distinguish from zero either way. The log-linear OLS counterpart in Table 37 gives the same null.

Three features characterise where the predictive content sits. First, the LSEG combined ESG score loads negatively and significantly throughout ($\beta = -0.0076$ to -0.0083 , $p < 0.01$): this is the well-documented controversy-adjustment effect, since LSEG penalises firms for adverse media events that overlap heavily with the Reprisk-incident set ([Glossner \(2021\)](#)). The S&P Global total ESG score, which lacks the same adjustment, loads at $\beta = 0.001$ (insignificant),

Table 13: PPML horse race of LSI on next-year Reprisk incident count

Dependent Variable:	rep_incidents_lead1		
Model:	(1)	(2)	(3)
<i>Variables</i>			
lsi_mean	0.5049 (0.7552)	0.6700 (0.7450)	0.5099 (0.7747)
ln_at	0.7227*** (0.0193)	0.7395*** (0.0196)	0.7454*** (0.0243)
leverage	0.5311*** (0.1570)	0.5783*** (0.1672)	0.6441*** (0.1811)
roa	1.038*** (0.4024)	1.771*** (0.4211)	1.806*** (0.4880)
rd_int	5.828*** (0.6951)	5.366*** (0.7429)	5.432*** (0.7793)
capx_int	2.885*** (0.9472)	2.947*** (0.9820)	3.358*** (1.059)
ln_age	-1.203*** (0.3141)	-1.972*** (0.4852)	-1.604*** (0.5222)
lseg_esg		-0.0076*** (0.0017)	-0.0083*** (0.0017)
sp_esg_total			0.0013 (0.0019)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Squared Correlation	0.66479	0.68283	0.68074
Pseudo R ²	0.60579	0.62406	0.62293
BIC	66,225.1	56,686.1	50,384.5
Observations	6,936	5,797	4,784
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. The dependent variable is the firm-year count of Reprisk incidents in year $t + 1$ (`rep_incidents_lead1`), constructed from the Reprisk WRDS extract described in Section 4.3. Estimation is by pseudo-Poisson maximum likelihood (PPML) following [Silva and Tenreyro \(2006\)](#), which handles the zero-inflated right-skewed distribution of the count outcome without requiring a $\log(1 + y)$ transformation.

confirming this interpretation. Second, the firm-fundamental controls absorb roughly 60% of the variation in next-year incident counts: $\ln(AT)$ at $\beta = 0.74$, leverage at $\beta = 0.58$, R&D intensity at $\beta = 5.37$, capex intensity at $\beta = 2.95$ (all $p < 0.01$ in column (2)). What predicts incidents within industry-year is operational scale and investment, not sustainability orientation. Third, with 5,797 firm-years in column (2), the null is not driven by power: the implied 95% CI on the lead-1 elasticity runs from -0.80 to $+2.14$, ruling out large negative effects.

5.4.3 Robustness across severity, firm fixed effects, horizons, and pillars

The headline null could in principle hide an effect along four margins that the main specification averages out: severity (high-severity exposure may matter more than raw counts), within-firm identification, horizon (capability mechanisms may take years to build), and pillar-on-pillar effects (cross-pillar aggregation may mask diagonal effects).

Severity-weighted outcomes. Table 33 replaces raw counts with the severity-weighted index $\sum_j \text{severity}_j \times \text{reach}_j \times \text{novelty}_j$ and additionally reports the high-severity count (severity ≥ 3). The LSI coefficient remains positive and insignificant on both the severity-weighted index ($\beta = 0.73$ and 0.58) and the high-severity count ($\beta = 1.18$, SE 1.77).

Firm and year fixed effects. Table 34 replaces NAICS-2-by-year with firm and year fixed effects. The LSI coefficient flips sign ($\beta = -0.45$ and -0.55) but stays far from significance, and the within- R^2 collapses to roughly 0.016. The sign flip is informative: it is consistent with the small positive cross-sectional loading in the main spec reflecting a residual size-and-coverage correlation (larger firms with more visible sustainability programmes are also more visible to Reprisk’s monitoring of over 100,000 news sources) rather than a substantive orientation channel.

Horizon robustness. Table 35 reports horizons $k = 1, \dots, 5$. The LSI coefficient drifts upward from $\beta = 0.67$ at $k = 1$ to $\beta = 1.41$ at $k = 5$ while standard errors widen, and all horizons fail to reach significance. The drift’s direction rules out the "capability takes time to build out" rescue, which would require the loading to turn negative and significant at long horizons.

Pillar decomposition. Table 36 estimates the pillar-on-pillar diagonal. The pillar-specific LSI loadings are all positive and insignificant (E : $\beta = 1.43$; S : $\beta = 0.79$; G : $\beta = 1.47$, weakly significant at 10% only). Pillar-specific sustainability orientation does not predict pillar-specific incidents either.

Taken together, the null on next-year incident counts is robust to severity weighting, within-firm identification, horizon length, and pillar decomposition. There is no rescue specification that recovers a substantive LSI-predicts-fewer-incidents effect.

5.4.4 Takeaway: a boundary condition for the LSI measure

This section establishes a boundary condition. The reactive direction (Section 5.4.1) shows posting-derived LSI does not measurably rise after firms’ first Reprisk incident, with a tight 95% confidence band of approximately $[-0.003, +0.003]$ on the ATT. The forward-looking direction

(Section 5.4.2) shows current LSI does not predict next-year (or longer-horizon) incidents in any specification I tested. I report both nulls in the main text rather than the appendix because the interpretation is informative for the rest of the paper.

The interpretation is that LSI captures forward-looking sustainability effort and capability accumulation rather than the high-frequency realisation of sustainability outcomes Reprisk records. This reading is consistent with two patterns established elsewhere. First, the within-versus-cross-sectional split in Section 5.3 shows LSI carries cross-sectional but not within-firm lead-indicator content for the LSEG combined ESG score; the same geometry appears in Table 34. LSI is about persistent firm orientation rather than year-on-year actions that mechanically prevent incidents. Second, the intangible-investment interpretation built up in Section 3.1 and tested in Section 5.7 frames LSI as a labour-input flow into a capability stock: investment flows produce capability over time, not deterministic prevention of operational risk realisations.

The pattern is also consistent with the corporate-misconduct literature. Glossner (2021) documents strong incident recidivism driven by firm and industry characteristics; Krüger (2015) shows the market response to corporate ESG events is shaped by firm characteristics other than stated sustainability orientation; Krueger et al. (2020) characterises institutional investors as concerned about climate risks but unable to derive precise signals from sustainability information. Across these papers, the operative mechanisms driving Reprisk-style incidents are operational and structural variables (size, leverage, R&D, capex, sector exposure), not sustainability orientation in the sense LSI captures.

The boundary is a positive statement about scope rather than a negative one about the measure. LSI is a forward-looking labour-input measure of sustainability capability, and it carries the lead-indicator content for ESG ratings (Section 5.3), the convergent-validity content against LSEG and S&P Global (Section 5.1), the climate-exposure tracking (Section 5.2), and the green-patenting predictive content (Section 5.5) one would expect from such a measure. It does not carry predictive content for the high-frequency Reprisk incident stream in either direction, and the reader should not interpret the rest of the paper as claiming that it does.

5.5 Predicting green innovation beyond commercial ESG ratings

This section tests whether LSI predicts substantive green-innovation output, measured by next-year counts of climate-change-mitigation patents. It is the headline outcome analysis for the paper's "beyond ESG ratings" claim. Green patents are the outcome where Cohen et al. (2026) explicitly document that leading commercial ESG scores fail: their ratings-versus-innovation disconnect is the title of the paper. Green patenting is also a measurable output rather than a disclosure intermediate, so a finding that LSI predicts green patents while ratings do not is robust to the disclosure-versus-action critique that motivates the construct (Bingler et al., 2022, 2024; Raghunandan and Rajgopal, 2024). The postings-to-patents path has a recent instantiation in Darendeli et al. (2022), who show that green hiring predicts subsequent green patents and frame the postings-as-investment reading that I extend to the broader capability construct. The design therefore allows close replications of both leading papers on the same panel.

I organise the section around six analyses: a baseline (5.5.1), a Darendeli replication (5.5.2), the headline Cohen-Gurun-Nguyen replication (5.5.3), a horse race against ratings (5.5.4), an

R&D-complementarity and green-versus-non-green placebo (5.5.5), and a pillar decomposition (5.5.6).

The headline finding is the Cohen et al. (2026) replication: in their exact specification, LSEG ESG enters at the margin of significance with negligible economic magnitude, while LSI loads at $\beta = 3.56$ ($p < 0.01$) and adds 0.7 percentage points of within- R^2 on the green-patent prediction. The analysis replicates the Cohen disconnect on my panel and extends its resolution to a labour-input measure.

5.5.1 Baseline PPML on next-year green patents

I begin with a PPML baseline of next-year green-patent counts on $LSI_{i,t}$ with the standard finance control set ($\ln(AT)$, leverage, ROA, R&D intensity, capex intensity, $\ln(1 + \text{age})$) under year, industry-by-year, and firm-and-year fixed effects. The dependent variable is the next-year count of patents carrying at least one CPC Y02 or Y04S classification, matching Cohen et al. (2026) and Darendeli et al. (2022). Estimator is PPML following Silva and Tenreyro (2006).

Table 14: PPML baseline of next-year green patents on LSI, CPC Y02+Y04S green flag

Dependent Variable:	patents_green_lead1		
Model:	(1)	(2)	(3)
<i>Variables</i>			
lsi_mean	17.17*** (5.129)	5.546 (4.017)	0.9613 (1.792)
ln_at	0.8523*** (0.0793)	1.102*** (0.0461)	0.0710 (0.1556)
leverage	1.414* (0.7517)	-0.1038 (0.4866)	-0.2434 (0.5012)
roa	3.062*** (0.9424)	-0.1643 (0.7482)	-0.4109 (0.4796)
rd_int	10.33*** (1.296)	10.28*** (1.209)	-0.9350 (1.870)
capx_int	9.092*** (2.748)	2.133 (2.019)	-3.002 (3.362)
ln_age	0.7062 (0.7411)	0.4519 (0.7273)	1.737*** (0.6541)
<i>Fixed-effects</i>			
year	Yes	Yes	Yes
naics2		Yes	
company_id			Yes
<i>Fit statistics</i>			
Squared Correlation	0.25505	0.73179	0.93079
Pseudo R^2	0.51585	0.77036	0.91769
BIC	102,961.4	48,913.1	17,657.8
Observations	6,332	6,272	3,154

Clustered (company_id) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Dependent variable is the count of patents filed in year $t + 1$ that carry at least one CPC Y02 or Y04S classification, the climate-change-mitigation and smart-grid flag described in Section 4.3 and used in Cohen et al. (2026) and Darendeli et al. (2022). The regressor of interest is $LSI_{i,t}$, the firm-year mean of posting-level `total_max` from Section 3.6.

Table 14 reports the three columns. The unconditional year-only column delivers $\beta = 17.17$ ($t = 3.35$), implying that a one-standard-deviation move in LSI ($\sigma = 0.037$) raises the next-year green-patent rate by about 89%. Adding industry-by-year fixed effects shrinks the loading to $\beta = 5.55$ (no longer significant at 10%): roughly two-thirds of the cross-sectional gradient is sectoral. The firm-and-year column collapses the within-firm loading to $\beta = 0.96$ (insignificant). This is the recurring pattern: LSI captures persistent firm-level orientation that predicts green-patent output cross-sectionally, but year-to-year within-firm variation does not push year-to-year green output.

Table 38 re-estimates on the broader OECD ENV-TECH 2023 classification. The industry-by-year loading is $\beta = 11.05$ ($p < 0.01$), about twice the Y02+Y04S magnitude; the firm-and-year loading is again insignificant. The pattern is robust to classification breadth. The headline replications below use Y02+Y04S to remain comparable with the published specifications.

5.5.2 Replication of Darendeli, Law, and Shen (2022) with LSI added

Darendeli et al. (2022) estimate $\ln(1 + \text{green patents}_{i,t+1})$ on firm controls ($\ln(\text{EMP})$, $\ln(\text{SALE})$, R&D/sales, capex/sales, advertising/sales, leverage, ROA, book-to-market, $\ln(1 + \text{age})$, tangibility) with industry-by-year fixed effects and firm-clustered standard errors. Their headline finding is that green hiring (from Burning Glass green-task keywords) predicts subsequent green patents. I hold their right-hand-side fixed and substitute my LSI for their green-hiring measure on the same firms. LSI differs from their measure in posting source (LinkUp), scoring (continuous cosine similarity over a 31-topic action-language reference rather than discrete green-task counts), and construct breadth (full sustainability capability rather than only environmental tasks).

Table 39 reports five columns. Column 1 is the Darendeli baseline without LSI; controls behave as expected. Column 2 adds $\text{LSI}_{i,t}$ at $\beta = 2.15$ (SE 0.81, $p < 0.01$). Column 3 adds LSEG ESG on the LSEG sub-sample: the LSI loading rises to $\beta = 3.88$ ($p < 0.01$) while LSEG is indistinguishable from zero. The rise is mechanically a partial-out effect: the LSEG sub-sample has denser ESG-reporting infrastructure, and the variation surviving partialling LSEG out is the variation uncorrelated with reporting density, which is what LSI is designed to capture. Columns 4-5 add firm-and-year fixed effects: the within-firm LSI loading is $\beta = 0.22$ (insignificant), confirming the cross-sectional-versus-within-firm split.

Table 40 reports the PPML sensitivity. Point estimates are in the same direction ($\beta = 5.02$ and 4.87) but with wider standard errors and no statistical significance: PPML on the raw count is dominated by the right tail of large patenters, while the log transformation in Table 39 balances small and large patenters. I follow Cohen et al. (2026) and Darendeli et al. (2022) in reporting log-linear OLS as the primary specification for replication tables.

5.5.3 Replication of Cohen, Gurun, and Nguyen (2024): the headline finding

The Cohen-Gurun-Nguyen replication is the headline analysis. Cohen et al. (2026) document that leading commercial ESG scores are essentially uncorrelated with green patenting, interpreting the disconnect as evidence that disclosure-based ratings fail to capture substantive sustainability-innovation activity. Their specification is a log-linear OLS of $\ln(1 + \text{green patents}_{i,t+1})$ on the LSEG ESG score with a tight valuation control set (log market cap, book-to-market, log age,

R&D intensity, ROA) and industry-by-year fixed effects. I replicate that specification on my panel and add LSI as a marginal regressor, both overall and decomposed by pillar.

Table 15: (headline): Cohen-Gurun-Nguyen (2020) ESG-Innovation Disconnect replication with LSI added

Dependent Variable: Model:	ln_green_lead1		
	(1)	(2)	(3)
<i>Variables</i>			
lseg_esg	0.0036* (0.0019)	0.0034* (0.0019)	0.0030 (0.0019)
ln_mktcap	0.2978*** (0.0346)	0.2983*** (0.0343)	0.2964*** (0.0341)
btm	0.1903*** (0.0450)	0.1928*** (0.0449)	0.1922*** (0.0455)
ln_age	-0.0928 (0.1253)	-0.1009 (0.1242)	-0.0829 (0.1192)
rd_int	1.382*** (0.3747)	1.351*** (0.3775)	1.443*** (0.3752)
roa	-0.4517** (0.2075)	-0.4574** (0.2083)	-0.5295*** (0.2049)
lsi_mean		3.558*** (1.054)	
E_mean			-3.976 (2.661)
S_mean			-0.1306 (2.833)
G_mean			-2.132 (1.966)
Sec_mean			9.930*** (2.298)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Within R ²	0.22895	0.23588	0.24601
R ²	0.33522	0.34119	0.34993
Observations	4,556	4,556	4,556

Clustered (company_id) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. This table replicates the headline specification of [Cohen et al. \(2026\)](#) on the LSI panel and adds LSI as a marginal regressor in column 2 and the four LSI pillars in column 3. The dependent variable is $\ln(1 + \text{green patents}_{i,t+1})$ on the CPC Y02+Y04S flag, matching their construction.

Table 15 reports three columns. Column 1 is the pure Cohen specification on my sample: $\beta_{\text{LSEG}} = 0.0036$ ($p < 0.10$), economically modest (a one-SD LSEG move corresponds to a 7% movement in expected next-year green patents). Within- R^2 is 0.229. This reproduces the Cohen disconnect on my sample.

Column 2 adds $\text{LSI}_{i,t}$. LSI loads at $\beta = 3.56$ (SE 1.05, $p < 0.01$). The economic magnitude is large: a one-SD LSI move corresponds to a 14% movement in expected next-year green patents, twice the LSEG-implied magnitude. Within- R^2 rises from 0.229 to 0.236, a 0.7 percentage point

gain at constant valuation controls. The LSEG coefficient is essentially unchanged at $\beta = 0.0034$, so the LSI contribution is incremental. This is the central finding. In the specification where Cohen et al. document the disconnect, a labour-input measure adds significant explanatory power. The disconnect is therefore a property of rating-based measurement, not of the green-patent outcome per se; a labour-input proxy that bypasses the disclosure channel contributes incremental power that the rating, by construction, does not.

Column 3 replaces overall LSI with the four pillar means. The sectoral pillar carries most of the loading ($\beta = 9.93$, $p < 0.01$); E, S, G enter individually with negative point estimates not distinguishable from zero. The pillar pattern must be read against the multicollinearity caveat from Section 3.9 (within-pillar mean anchor cosine of 0.71, Table 60): the four pillar means carry overlapping variation, so individual coefficients in a four-pillar regression represent residual-after-netting variation only. The sectoral-pillar dominance is consistent with the substantive content of that pillar (Sec1 responsible-finance stewardship, Sec2 green buildings) overlapping mechanically with renewables, electrification, and built-environment efficiency work in Y02.

5.5.4 Horse race against LSEG and S&P Global ratings (simpler-control specification)

I re-run the comparison against LSEG and S&P Global ESG ratings under the paper’s standard six-variable control set and industry-by-year fixed effects, for comparability with the rest of the paper’s results sections.

Table 41 reports three columns: LSI plus LSEG, LSI plus S&P, and LSI plus both. The LSI loading is $\beta = 5.92$, 5.39, and 5.87, similar in magnitude to the baseline ($\beta = 5.55$) but not statistically distinguishable from zero at 10%. The LSEG and S&P coefficients are essentially zero ($|\beta| < 0.01$) and equally insignificant.

The contrast with the Cohen-specification result is important to flag honestly. Two differences explain it. First, Cohen uses log market cap rather than log total assets, which absorbs a different part of the cross-sectional size variation relevant for innovation output. Second, Cohen includes book-to-market, the Tobin’s-q-style growth-opportunity control standard in the innovation literature. The richer Cohen and Darendeli control sets absorb more sectoral and growth confounding than the paper’s standard baseline; the LSI loading survives both but the simpler-control horse race cannot statistically distinguish it from zero. I read the simpler-control horse race as setting the precision floor and the literature-standard replications as the sharper tests; neither makes the LSI loading negative or contradicts the cross-sectional pattern.

5.5.5 R&D-complementarity test and green-versus-non-green placebo

The R&D-complementarity test asks whether LSI and R&D combine through a joint production function: if so, the LSI-by-R&D interaction on green output should be positive and significantly larger in absolute value than on non-green output. This is the green-innovation analogue of the directed-technical-change mechanism in Aghion et al. (2016).

Table 42 reports the interaction at $\beta = -56.93$ ($p < 0.01$) on green and $\beta = -17.88$ (insignificant) on non-green. The sign does not support the sustainability-specific complementarity hypothesis: at higher R&D intensity the marginal effect of LSI on green patenting is smaller,

more consistent with a substitution or saturation reading in which high-R&D firms produce green innovation independently of their sustainability-labour-input intensity. The per-pillar robustness in Table 43 confirms: no pillar shows a green-specific complementarity pattern sharper than its non-green analogue.

The green-versus-non-green placebo is a falsification test: LSI is designed to predict sustainability-oriented output, not innovation in general. Table 44 reports LSI loading at $\beta = 5.55$ on green and $\beta = -1.86$ on non-green under the same industry-by-year specification. The direction is consistent with the placebo prediction (positive on green, negative on non-green; cross-coefficient gap of approximately 7.4); statistically, neither loading is significant at 10% on its own. I report the placebo as directionally consistent but not a strong stand-alone significance test, and I read it together with the Cohen replication, where the literature-standard control set produces a sharp loading.

5.5.6 Pillar decomposition on green and non-green patents

The final analysis decomposes the LSI loading by pillar means under the paper’s standard baseline controls and industry-by-year fixed effects.

Table 45 reports the decomposition. On green, the sectoral pillar dominates at $\beta = 23.52$ ($p < 0.01$); G enters negatively at $\beta = -13.24$ ($p < 0.05$); E and S are insignificant. On non-green, E enters negatively at $\beta = -15.94$ ($p < 0.01$); the sectoral pillar drops to $\beta = 4.40$ (insignificant), an order of magnitude below its green loading.

Three points, with a caveat. The caveat is that the four pillar means share substantial multicollinearity (within-pillar mean cosine of 0.71, Table 60), so individual coefficients identify residual-after-netting variation only and are not directly comparable to the overall LSI coefficient. First, the across-outcome pattern for the sectoral pillar (positive on green, an order of magnitude smaller and insignificant on non-green) is the strongest signal, consistent with Sec1 stewardship and Sec2 green-buildings content overlapping with Y02 renewables, electrification, and built-environment work, though best read as suggestive given multicollinearity. Second, the E-pillar pattern (negative on non-green, insignificant on green) is best read as a partial-residual artefact rather than as evidence environmental hiring substitutes for general patenting. Third, the G-pillar negative loading on green is consistent with the broader finding that governance vocabulary is least connected to substantive sustainability action: governance postings concentrate in compliance and reporting, and the Walk-versus-Talk benchmark in Table 56 flagged G as the one pillar where both measures show a negative Walk-Talk gap. A univariate pillar-by-pillar regression is deferred to the job-market version with the longer 2013-2023 panel.

Takeaway. This section delivers the paper’s cleanest "beyond ESG ratings" finding through the Cohen replication. In the exact specification Cohen et al. (2026) use to document the disconnect, LSI loads at $\beta = 3.56$ ($p < 0.01$) while the LSEG ESG score is at the margin of significance and explains very little. The Darendeli replication confirms LSI carries the green-patent signal in the published hiring-to-patents specification of Darendeli et al. (2022): the within-industry-year LSI loading is $\beta = 2.15$ ($p < 0.01$) and rises to $\beta = 3.88$ once LSEG is controlled. The simpler-control horse race produces estimates of the same sign and magnitude but absorbed statistical

significance, which I read as a precision floor while the literature-standard replications carry the sharp test. Two honest boundaries: the within-firm null on green patents ($\beta = 0.22$ in the Darendeli firm-and-year column) confirms the cross-sectional-versus-within-firm split; and the R&D-interaction test does not support a sustainability-specific complementarity reading. Taken together, these results extend the Darendeli et al. (2022) hiring-to-patents path to the broader capability construct and contribute incremental cross-sectional explanatory power to the Cohen et al. (2026) green-patent equation through a labour-input proxy that bypasses the disclosure channel (Bingler et al., 2022; Raghunandan and Rajgopal, 2024), while bounding the claim to cross-sectional rather than within-firm variation.

5.6 Future profitability: the intangibles signature

The construct definition in Section 3.1 commits LSI to a sharp empirical prediction: a labour-input intangible-investment flow should depress near-term return-on-assets (because the labour cost is expensed against current operating income) and elevate long-horizon market valuation (because the resulting capability accrues through the discounted-cash-flow channel). Edmans (2011), Eisfeldt and Papanikolaou (2013), Peters and Taylor (2017), and Crouzet et al. (2022) document this joint signature for employee satisfaction, organisation capital, and the broader intangibles family. This section asks whether LSI traces the same arc.

Specification. For outcome $y \in \{\text{ROA}, \ln(\text{sales}), q, \ln(\text{labour productivity})\}$ and horizon $k \in \{1, 3, 5\}$, I estimate

$$y_{i,t+k} = \alpha + \beta \cdot \text{LSI}_{i,t} + \phi \cdot y_{i,t} + \gamma' X_{i,t} + \mu_{j(i)} + \lambda_t + \varepsilon_{i,t+k}, \quad (11)$$

where $X_{i,t}$ is the standard finance control set ($\ln(\text{AT})$, leverage, R&D intensity, capex intensity, $\ln(1 + \text{age})$), $\mu_{j(i)} + \lambda_t$ are NAICS-2 industry and year fixed effects, and standard errors are clustered at the firm level. The lagged-self term $\phi \cdot y_{i,t}$ strips out mechanical persistence so that β identifies the marginal response of y_{t+k} to current LSI.

5.6.1 ROA dynamics: the short-run cost of intangible investment

Table 16 reports the future-ROA regressions. The headline is the one-year-ahead loading: $\beta = -0.0564$ ($p < 0.01$, $N = 18,467$). A one-standard-deviation increase in LSI (0.037) lowers next-year ROA by approximately 21 basis points, conditional on current ROA and the standard controls. The coefficient decays in absolute magnitude across horizons (-0.0564 at $k = 1$, -0.0443 at $k = 3$, -0.0259 at $k = 5$) and loses significance by $k = 3$, consistent with a cost concentrated in the period the investment is undertaken; I do not push the long-horizon point estimates further because the lead- k panel mechanically shrinks ($N = 4,196$ at $k = 5$). The near-term negative loading is what the intangible-capital framework predicts: the wages of sustainability-capability hires and the associated training and integration time map mechanically into a near-term reduction in operating income, in the same way that the sustainability component of SG&A enters Peters and Taylor (2017)’s organisation-capital decomposition. The magnitude is in line with the short-run cost estimates for AI and green hiring in Babina et al. (2024) and Darendeli

et al. (2022).

Table 16: LSI on future ROA at horizons $k \in \{1, 3, 5\}$

Dependent Variables: Model:	roa_lead1 (1)	roa_lead3 (2)	roa_lead5 (3)
<i>Variables</i>			
lsi_mean	-0.0564*** (0.0194)	-0.0443 (0.0327)	-0.0259 (0.0536)
roa	0.6441*** (0.0151)	0.4408*** (0.0248)	0.4046*** (0.0330)
ln_at	0.0044*** (0.0004)	0.0053*** (0.0007)	0.0061*** (0.0011)
leverage	0.0066 (0.0048)	0.0192** (0.0087)	0.0032 (0.0121)
rd_int	-0.1649*** (0.0232)	-0.1349*** (0.0447)	0.0207 (0.0611)
capx_int	0.0335 (0.0264)	0.0662 (0.0423)	-0.0620 (0.0685)
ln_age	0.0119** (0.0049)	0.0122 (0.0112)	0.0244 (0.0218)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Within R ²	0.54078	0.30551	0.19094
R ²	0.57196	0.34051	0.23015
Observations	18,467	9,790	4,196
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. Dependent variable is firm-year ROA k years ahead, $ROA_{i,t+k}$, with $k \in \{1, 3, 5\}$ across columns. The right-hand-side variable of interest is the contemporaneous headline aggregator $LSI_{i,t}$ (Section 3.6).

Table 46 (appendix) reports the parallel $\ln(\text{sales})$ regressions. The LSI coefficient is statistically insignificant at every horizon and changes sign across columns. The null is informative for the construct: LSI is not a contemporaneous scale or demand-shock measure, and the dollar volume of sales is sticky enough at the one-to-five-year horizon that a labour-input capability flow does not move it conditional on the highly persistent lagged outcome.

5.6.2 Tobin's q and labour productivity: the long-horizon benefit side

Table 17 reports the future-Tobin's-q regressions. The headline is the five-year-ahead loading: $\beta = 1.98$ ($p < 0.05$, $N = 4,053$). A one-standard-deviation increase in LSI implies a long-horizon q premium of approximately 0.073, on the order of five percent of one cross-sectional standard deviation of q, in the same economic-magnitude range as the long-horizon equity premium for the Best Companies to Work For in Edmans (2011) and the organisation-capital return premium in Eisfeldt and Papanikolaou (2013). The short-horizon q loadings are not load-bearing: at $k = 1$ the autoregressive control on contemporaneous q absorbs almost all the available cross-sectional variation, and the residual identification window only opens up enough by $k = 5$ for LSI to load precisely. The increasing-precision-with-horizon pattern for q

is the mirror image of the decreasing-precision-with-horizon pattern for ROA, exactly as the costly-investment-with-long-horizon-payoff narrative predicts.

Table 17: LSI on future Tobin's q at horizons $k \in \{1, 3, 5\}$

Dependent Variables:	tobins_q_std_lead1	tobins_q_std_lead3	tobins_q_std_lead5
Model:	(1)	(2)	(3)
<i>Variables</i>			
lsi_mean	16.58 (14.77)	0.1887 (0.7116)	1.978** (0.9804)
tobins_q_std	0.1237*** (0.0079)	1.383** (0.6149)	0.6159*** (0.0729)
ln_at	-1.334 (1.203)	-0.0030 (0.0303)	-0.0495*** (0.0189)
leverage	8.083 (7.266)	0.0305 (0.1766)	-0.0231 (0.2254)
roa	-30.26 (33.38)	-4.646 (4.803)	1.093* (0.5969)
rd_int	-36.78 (45.89)	-7.346 (9.732)	4.803*** (1.249)
capx_int	-20.78 (22.86)	-3.530 (3.253)	-1.249 (0.9364)
ln_age	6.675 (7.663)	0.1330 (0.9775)	0.2714 (0.2324)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Within R ²	0.02144	0.28103	0.35880
R ²	0.02395	0.30152	0.42428
Observations	17,752	9,445	4,053

Clustered (company_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Dependent variable is firm-year standardised Tobin's q k years ahead, $q_{i,t+k}$, from the Peters-Taylor data (Peters and Taylor, 2017), with $k \in \{1, 3, 5\}$ across columns. The right-hand-side variable of interest is the contemporaneous headline aggregator $LSI_{i,t}$.

The joint reading of Tables 16 and 17 is the central empirical fact of this section. LSI depresses next-year ROA at $\beta = -0.056$ and elevates Tobin's q five years ahead at $\beta = +1.98$. This is the classic intangibles signature: short-run profitability cost paired with long-horizon valuation premium, conditional on the lagged self of the outcome. That the same arc appears for a measure built from job postings rather than from survey data, SG&A allocations, or compensation disclosures supports the construct's interpretation as a labour-input intangible-investment flow.

Table 47 (appendix) reports the parallel $\ln(\text{labour productivity})$ regressions. The LSI coefficient is insignificant at every horizon. As with the sales result, the null is the right one for the construct: LSI is a capability-accumulation measure, not a contemporaneous productivity-shock measure, and the long-horizon market-valuation premium delivered by the q regression is the substantive benefit-side result.

5.6.3 Capex-intensity interaction: where the medium-horizon cost is concentrated

Table 48 (appendix) adds the LSI-times-capex interaction to the ROA_{t+3} regression. The motivation is the joint-adjustment-cost reading in Peters and Taylor (2017) and Crouzet et al. (2022): tangible and intangible investment flows co-move, and the joint cost of accumulating both at once may exceed the sum of the individual costs. The interaction coefficient is $\beta_{LSI \times capx} = -1.49$ ($p < 0.05$), with the LSI main effect approximately zero and the capex main effect positive. The implied marginal LSI loading on ROA_{t+3} is close to zero at zero capex intensity, approximately -0.03 at the sample-median capex intensity (0.03), and approximately -0.26 at the 95th percentile (0.18). The medium-horizon cost of LSI accumulation is concentrated in firms simultaneously expanding physical capital.

I read this cautiously and attach two competing interpretations without committing to one. The first is joint-adjustment-cost complementarity: integrating new sustainability capability into new physical plant (energy-efficient process design, EHS protocols, training engineering hires for new assets) carries a coordination cost that loads on the operating-income line at the medium horizon. The second is substitution: firms with high capex may invest in physical capital in lieu of sustainability capability, and the negative interaction may pick up diminishing returns rather than joint adjustment cost. The cleanest discriminating test would compare the long-horizon response of LSI in high-capex versus low-capex sub-samples; the consortium-draft sample is too thin in the relevant cells and I defer this to the job-market version. The relevant point is that the future-ROA loading is concentrated where it should mechanically be largest.

Takeaway

LSI depresses near-term ROA at $\beta = -0.056$ one year ahead and elevates Tobin's q at $\beta = +1.98$ five years ahead, conditional on the lagged self of each outcome. The pattern is the joint signature documented by Edmans (2011) for employee satisfaction, by Eisfeldt and Papanikolaou (2013) and Peters and Taylor (2017) for organisation capital, and by Crouzet et al. (2022) for the broader intangible-capital family. The nulls on $\ln(\text{sales})$ and $\ln(\text{labour productivity})$ confirm that LSI is not capturing a generic positive-shock or productivity-shock signal, and the capex-interaction locates the medium-horizon cost where it should mechanically concentrate. The future-profitability evidence supports the construct's interpretation as a labour-input intangible-investment flow and motivates the formal Peters-Taylor decomposition in Section 5.7.

5.7 Integration into the intangible-capital catalog: LSI-K and the investment-q regression

If LSI is genuinely an intangible-capital flow, then its perpetual-inventory cumulation LSI-K should fit inside the Peters-Taylor framework alongside knowledge capital and organisation capital. I therefore put LSI-K into the augmented investment-q regression that Peters and Taylor (2017) use to validate their q against the intangible-capital base, and ask whether LSI-K is empirically incremental to their two existing components. I view this as the load-bearing intangibles-literature contribution of the paper. The natural worry is that SG&A already records sustainability and ESG-reporting payroll costs, so a sustainability-hiring measure might simply

be relabelling part of organisation capital; the Peters-Taylor decomposition in Table 18 discharges this worry directly. I offer LSI-K here as a framework-integration step rather than as a definitive intangible stock: the depreciation rate δ for sustainability capital is not pinned down by either theory or prior estimates, and the perpetual-inventory output is correspondingly conditional on the central case I adopt.

5.7.1 Constructing LSI-K via perpetual inventory

The full construction is in Section 3.8; I summarise here only what matters for the regressions. The flow input is the wage-weighted dollar-denominated LSI flow from Section 3.7, deflated to 2017 dollars with the CPI-U. The cumulation rule is the standard Hall (1990) perpetual inventory used by Corrado et al. (2009) and Peters and Taylor (2017): $K_t = (1 - \delta)K_{t-1} + I_t$, with initial condition $K_{t_0} = I_{t_0}/(g + \delta)$. I take $g \approx 0.014$ from the pooled panel and $\delta = 0.20$ as the central case, matching Peters and Taylor (2017); sensitivities at $\delta \in \{0.15, 0.30\}$ are in Appendix A.9. The output is a firm-year stock LSI- $K_{i,t}$ in millions of 2017 dollars, unit-comparable to Compustat's k_int . The share I use below, $lsik_share_{i,t} = LSI-K_{i,t}/(k_int_{i,t} + LSI-K_{i,t})$, is the fraction of total intangible capital devoted to sustainability orientation.

5.7.2 Investment-q with LSI-K share: the headline Peters-Taylor decomposition

The headline regression uses the Peters-Taylor investment-rate $IK_{i,t}^{PT} = CAPX_{i,t}/k_int_{i,t}$ as the outcome and augmented Tobin's q (q_tot) as the baseline right-hand-side variable, adding the two Peters-Taylor components and then LSI-K share successively. Industry (NAICS-2) and year fixed effects throughout; firm-clustered standard errors. Table 18 reports the four-column decomposition on the held-fixed sample of $N = 19,737$ firm-years.

The substantive content is in the contrast between columns (1)-(3) and column (4). The baseline q regression and the two-component Peters-Taylor specifications produce a within- R^2 below 0.001: q_tot loads near zero, the two intangible-capital intensities are negative and small, and the existing catalog leaves the within-cell residual essentially untouched. Adding LSI-K share in column (4) changes the picture qualitatively: $\beta_{lsik_share} = 47.80$ (SE = 18.12, $p < 0.01$) and the within- R^2 jumps to 0.046. The coefficient on q_tot flips and becomes significant, and the coefficients on k_int_org and k_int_know shift substantially (organisation-capital intensity nearly triples in magnitude, knowledge-capital intensity switches sign). Collinear regressors would inflate standard errors and leave partial coefficients qualitatively unchanged; these large directional shifts are the signature of a regressor that carries variation orthogonal to the existing set and sharpens the partial associations of the other components. LSI-K share is identifying a distinct intangible-capital channel that the Peters-Taylor decomposition was not previously controlling for.

The pattern is robust to the choice of denominator. Table 49 (appendix) varies the denominator of IK^{PT} and reproduces the result ($\beta_{lsik_share} = 47.35$, within- $R^2 = 0.045$ in the relevant column). The unweighted-intensity robustness in Table 50 (appendix) establishes the converse: cumulating the dimensionless $LSI_{i,t}$ aggregator through the same perpetual inventory produces a small and insignificant coefficient ($\beta = -0.65$). Wage-weighting is what brings the measure

Table 18: (headline): Peters-Taylor decomposition with LSI-K share

Dependent Variable: Model:	IK_PT			
	(1)	(2)	(3)	(4)
<i>Variables</i>				
q_tot	-0.0001 (0.0007)	0.0002 (0.0008)	0.0001 (0.0008)	-0.0335** (0.0144)
lsik_share				47.80*** (18.12)
<i>Fixed-effects</i>				
naics2	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Within R ²	1.62×10^{-8}	7.27×10^{-5}	7.77×10^{-5}	0.04599
R ²	0.01446	0.01453	0.01453	0.05978
Observations	19,737	19,737	19,737	19,737

Clustered (company_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Dependent variable across all four columns is the Peters-Taylor investment-rate $IK_{i,t}^{PT} = CAPX_{i,t}/k_{int_{i,t}}$. The four columns add intangible-capital components successively: column (1) regresses on augmented Tobin's q (q_tot) only; column (2) adds organisation-capital intensity k_{int_org}/AT ; column (3) further adds knowledge-capital intensity k_{int_know}/AT , which together exhaust the Peters-Taylor intangible-capital decomposition; column (4) adds the LSI-K share of total intangible capital $lsik_share_{i,t} = LSI-K_{i,t}/(k_{int_{i,t}} + LSI-K_{i,t})$.

into the intangible-capital framework, and only the dollar-denominated cumulation produces a stock that is unit-comparable to k_{int} and that can enter the regression as a meaningful share.

5.7.3 Orthogonalised LSI-K: a methodological refinement

A natural concern is that LSI-K share might be picking up the sustainability-allocated component of organisation capital rather than identifying a separate channel. I orthogonalise $LSI-K_{i,t}$ on k_{int_org} , $\ln(AT)$, and leverage with NAICS-2-by-year fixed effects, and put the standardised residual into the second-stage investment-q regression (Table 51, appendix). The coefficient is small and insignificant. I read this as a methodological refinement rather than a substantive contradiction: the orthogonalisation operates on the LSI-K *level*, which scales with firm size and is dominated by the size component, so the residual variation is small and noisy after both size and organisation capital are absorbed. The share form is bounded in $[0, 1]$ and answers the right economic question (the proportion of intangible capital allocated to sustainability), in the same way that the R&D literature uses R&D-over-sales rather than R&D dollar levels. I therefore report LSI-K throughout as a share or ratio rather than as a residualised level.

5.7.4 Cross-sectional quintile sorts: the intangibles signature

If LSI-K is a genuine intangible-capital component, then firms in its right tail should display the classic intangibles signature: higher Tobin's q (the option value of the asset is in the market price) and lower contemporaneous ROA (the accumulating expense is recorded against current earnings). Eisfeldt and Papanikolaou (2013) document this for organisation capital and Eisfeldt

et al. (2022) extend it through a broader HML-style intangible factor. Within each year I form quintiles on $LSI-K_{i,t}/AT_{i,t}$ and report firm-year cell means.

Table 19: firm characteristics by LSI-K / AT quintile (within-year sorts)

LSI-K/AT quintile	N	Mean LSI-K/AT	Mean ROA	Mean Tobin's q	Mean q_tot	Mean R&D intensity	Mean CAPX intensity
Q1 (lowest)	5,309	0.0004	+0.028	1.23	0.85	0.013	0.033
Q2	5,306	0.0027	+0.019	1.34	1.43	0.015	0.032
Q3	5,307	0.0090	+0.010	1.93	2.79	0.039	0.033
Q4	5,306	0.0237	-0.023	2.22	3.88	0.063	0.033
Q5 (highest)	5,309	–	-0.079	4.71	2.97	0.073	0.041

Notes. Within-year quintile sort on the ratio $LSI-K_{i,t}/AT_{i,t}$, with LSI-K in millions of 2017 dollars and AT in millions of dollars (the units cancel; the ratio is dimensionless). Firm-year cell averages of total assets normalised LSI-K, return on assets (ROA), standard Tobin's q (`tobins_q_std`), augmented Peters-Taylor Tobin's q (`q_tot`), R&D intensity, and CAPX intensity.

The pattern is monotone across all five outcomes in the predicted direction. Mean ROA falls from +0.028 in Q1 to -0.079 in Q5; mean standard Tobin's q rises from 1.23 to 4.71 (a factor of 3.8x); augmented Peters-Taylor q rises from 0.85 to a Q4 peak of 3.88 (Q5 at 2.97, the right tail being predominantly small-cap firms with compressed augmented-q denominators); R&D intensity rises monotonically from 1.3 to 7.3 percent, and CAPX intensity from 3.3 to 4.1 percent. High-LSI-K firms are not substituting away from R&D or tangible investment in favour of sustainability hiring; they are doing more of all three, the same intangibles-tangibles complementarity that Corrado et al. (2009) document.

The Q5 cell entry for mean LSI-K/AT is shown as missing because the within-year ratio in Q5 is heavily right-skewed (a tail of small-cap firms with mechanically small AT denominators); the cross-sectional ROA and q patterns are unaffected because they are computed directly on the firm-year cell. The robustness sort in Table 52 (appendix) cuts on $LSI-K_{i,t}/k_{int_org_{i,t}}$ instead, which is scale-free, and reproduces the same monotone pattern. The intangibles signature is in the data, not in the choice of denominator. LSI-K does for sustainability orientation what the organisation-capital sort in Eisfeldt and Papanikolaou (2013) does for management capability: both slice on a labour-input-derived intangible stock, and both produce a familiar intangibles-heavy growth-firm profile in the high quintile.

Takeaway

LSI-K is a genuinely new intangible-capital component, not a duplicate of organisation or knowledge capital, and it carries the classic intangibles cross-sectional signature. The column (4) result in Table 18 is the load-bearing finding: $\beta_{lsik_share} = 47.80$ ($p < 0.01$), with the within- R^2 rising from below 0.001 to 0.046 after both Peters-Taylor components are in the regression. The orthogonalised null is a functional-form refinement (share, not residualised level), and the quintile sorts confirm that the intangibles signature is robust to the choice of scaling denominator. Together with the intangible-investment warrant for the LSI flow in Section 5.6, this evidence places the construct inside the intangible-capital framework alongside R&D and organisation capital. Peters and Taylor (2017), Crouzet et al. (2022), and Sichel (2025) flag sustainability orientation as a category whose measurement infrastructure is still maturing; LSI-K offers a labour-input integration into that catalog, conditional on the central-case depreciation rate.

5.8 Synthesis across the empirical results

The empirical sections together describe LSI as a persistent firm-level intangible-investment flow: convergent with commercial ESG ratings at non-redundant magnitudes (Section 5.1), tracking climate exposure (Section 5.2), leading ESG-rating revisions over one to three years (Section 5.3), predicting green-patenting output beyond ratings in the Cohen replication (Section 5.5), depressing near-term ROA and elevating long-horizon Tobin’s q (Section 5.6), and entering the Peters-Taylor decomposition as a genuinely incremental intangible component (Section 5.7).

Two cross-cutting patterns deserve attention before the discussion. First, the headline lead-indicator and green-patent results are cross-sectional but not within-firm: both are robust to industry-by-year fixed effects but absorbed by firm-and-year fixed effects, which I read honestly as LSI capturing persistent firm-level sustainability orientation rather than year-to-year operational shocks. Second, the null on Reprisk incidents (Section 5.4) and the null on near-term sales growth (Section 5.6) jointly pin down the discriminant-validity boundary: LSI is not a measure of realised operational risk control and not a measure of current-period productivity, but a labour-input intangible-investment flow.

6 Discussion and Conclusion

This paper introduces LSI, a firm-year measure of labour-input sustainability capability constructed from job postings. The central claim is not that LSI measures realised sustainability performance. It measures firms’ revealed labour-market investment in sustainability-related capability: the extent to which firms allocate labour demand toward environmental, social, governance, and sectoral sustainability work.

The empirical results support this interpretation. LSI is positively but modestly correlated with commercial ESG ratings and with firm-level climate exposure, suggesting that it captures sustainability-related content without being redundant with existing benchmarks. More importantly, LSI predicts subsequent LSEG and S&P Global rating revisions over one-to-three-year horizons even after controlling for current ratings. This relationship is primarily cross-sectional rather than driven by high-frequency within-firm variation. I interpret this pattern as evidence that commercial ratings under-weight persistent firm-level sustainability orientation that LSI captures and gradually recognise it over time.

The green-patent results point in the same direction. Replicating the [Cohen et al. \(2026\)](#) specification, LSI adds predictive power for subsequent green patenting in the setting where commercial ESG ratings have limited explanatory content. This confirms and generalises the labour-input insight of [Darendeli et al. \(2022\)](#): sustainability-related hiring contains information about future green innovation that ratings do not fully capture. The firm-fixed-effect attenuation is not a contradiction of this interpretation. It indicates that the relevant signal is persistent between-firm sustainability capability rather than year-to-year hiring shocks. That time scale is consistent with innovation as an output of accumulated labour, routines, and knowledge rather than a contemporaneous response to annual posting fluctuations.

The valuation results complete the intangible-capital interpretation. LSI is associated with lower near-term ROA and higher long-horizon Tobin’s q , the pattern expected when firms

incur short-run costs to build a capability that markets capitalise over longer horizons. When wage-weighted and cumulated into the perpetual-inventory stock LSI-K, the measure remains informative after controlling for Peters-Taylor knowledge capital and organisation capital. This is the paper’s strongest link to the own-account intangible-capital framework of [Corrado et al. \(2005, 2009\)](#) and to the firm-level implementation in [Peters and Taylor \(2017\)](#). Sustainability capability appears to be an intangible category not fully spanned by standard R&D and organisation-capital measures.

The null results are equally important because they define the boundary of the construct. LSI does not predict fewer subsequent Reprisk environmental, social, or governance incidents across horizons, severity weightings, pillar splits, or fixed-effect specifications. Nor does LSI rise materially after a firm’s first incident in the Sun-Abraham event-study design ([Sun and Abraham, 2021](#)). These nulls indicate that LSI is not a measure of realised incident avoidance, short-horizon ESG risk reduction, or reactive crisis management. This is not a failure of the measure; it is construct discipline. Reprisk incidents capture realised adverse events shaped by operational exposure, media visibility, compliance failures, and scale. LSI captures ex ante labour-market investment in sustainability capability. The separation is useful because it prevents the measure from being confused with an operational-risk or controversy proxy.

The contribution is therefore best understood as a measurement contribution with three implications. First, for the intangible-capital literature, the paper adds a firm-year, dollar-denominated measure of sustainability labour investment that can be cumulated into a stock and compared with existing intangible-capital components ([Peters and Taylor, 2017](#); [Crouzet et al., 2022](#); [Sichel, 2025](#)). Second, for the ESG-measurement literature, it provides an input-side, capability-oriented benchmark that is independent of voluntary disclosure choices and rating-vendor aggregation rules ([Berg et al., 2022](#); [Christensen et al., 2022](#); [Avramov et al., 2022](#); [Bingler et al., 2022](#); [Raghunandan and Rajgopal, 2024](#)). Third, for empirical sustainability finance, it shows that labour-input sustainability capability predicts green innovation and future rating revisions where commercial ESG ratings are incomplete ([Cohen et al., 2026](#); [Darendeli et al., 2022](#)).

Several limitations qualify the interpretation. First, the sample is restricted to listed firms with LinkUp coverage and to the external panels available for validation and outcome tests. The construct itself does not require listed status, but the current empirical implementation does. Second, postings measure labour demand rather than realised hiring. This follows the convention in postings-based work ([Hershbein and Kahn, 2018](#); [Deming and Kahn, 2018](#); [Babina et al., 2024](#); [Darendeli et al., 2022](#)), but it means that LSI should be read as a hiring-intent and capability-investment proxy rather than as realised employment. Third, the reference corpus is built from sustainability standards, many of which originate in disclosure regimes. The action-language curation and Walk-versus-Talk validation mitigate this concern, but residual disclosure-register leakage may remain. Fourth, the rating analysis currently uses two commercial providers, LSEG and S&P Global. A broader provider panel would allow a sharper decomposition of which rating methodologies under-weight the LSI signal most. Fifth, the paper is descriptive and predictive, not causal. I do not claim that increasing LSI causes green patents, valuation changes, or rating upgrades.

In short, LSI is not a measure of whether firms are already sustainable. It is a measure of whether firms are investing labour into becoming capable of sustainability-related work. That distinction is the paper's main conceptual payoff. It explains why LSI leads ratings and predicts green innovation, why it behaves like an intangible investment flow, and why it does not mechanically predict fewer high-frequency ESG incidents. The paper therefore contributes a new input-side measurement infrastructure for sustainability finance: one that complements ratings, disclosure text, and incident databases by measuring the labour firms commit to building sustainability capability.

A Measurement appendix

This appendix contains the source-family inventory, the prompt template and an example topic anchor, the alt-corpus construction recipe, the encoder-comparison placeholder, and the validation-battery numbers that support the construct claims in Section 3. Together they make the measurement pipeline transparent without requiring access to an external repository.

A.1 Posting language distribution

A.2 Source-family inventory

The curated corpus draws on the following 26 source families. The per-source bibliographic references and download URLs are listed at the end of this appendix; the per-topic source mapping is described in Section A.3 below.

GRI suite Universal Standards (GRI 1, 2, 3), all Topic Standards (200, 300, 400 series), all 41 Sector Standards.

SASB All 11 sector frameworks covering approximately 77 industries.

ISO family 14001 (Environmental Management), 45001 (Occupational Health and Safety), 50001 (Energy Management), 26000 (Social Responsibility guidance), 37001 (Anti-bribery), and the 2024 climate-change-mitigation amendment.

TCFD Final recommendations and implementing guidance.

ILO Fundamental conventions (Forced Labour, Freedom of Association, Discrimination, Child Labour, Occupational Safety and Health, Wage Setting), Decent Work Agenda, Just Transition Guidelines, and tripartite declarations on multinationals.

OECD MNE Guidelines (2023 revision), Due Diligence Guidance for Responsible Business Conduct, OECD Anti-Bribery Convention, G20/OECD Principles of Corporate Governance.

GHG Protocol Corporate Accounting and Reporting Standard, Scope 2 Guidance, Scope 3 Standard.

UNGPs UN Guiding Principles on Business and Human Rights.

UNGC UN Global Compact Ten Principles.

IFC Performance Standards 1-8 and Guidance Notes.

B Corp 2025 topical standards (9 files covering Workers, Community, Environment, Customers, Governance).

CDP Climate Change, Water Security, and Forests corporate questionnaires.

ESRS European Sustainability Reporting Standards: ESRS 1 (general principles), ESRS 2 (general disclosures), E1-E5, S1-S4, G1.

EU Taxonomy Environmental Delegated Act covering technical screening criteria for objectives 3 through 6.

PCAF Parts A (Financed Emissions), B (Facilitated Emissions), C (Insurance-Associated Emissions).

TNFD Taskforce on Nature-related Financial Disclosures recommendations and LEAP approach.

SBTN Science Based Targets Network methodology for nature.

SBTi Corporate Net-Zero Standard.

NZBA Net-Zero Banking Alliance Guidelines.

PRB / PSI UN Principles for Responsible Banking, Principles for Sustainable Insurance.

PRI UN Principles for Responsible Investment.

SA8000 Social Accountability standard.

LEED Leadership in Energy and Environmental Design v4.1 reference guides.

Equator Principles Fourth edition.

WBA World Benchmarking Alliance core social indicators.

CBD GBF Convention on Biological Diversity Global Biodiversity Framework.

IRENA Renewable Energy and Jobs annual reports.

Exclusions. Two source families are held out for use as external validation benchmarks rather than as anchor sources: ESCO (European Skills, Competences, Qualifications and Occupations) and O*NET (the Occupational Information Network green-occupation classification of [Dierdorff et al., 2009](#)). Including either in the reference corpus would render the validation tests in Section 3.9 circular. ISSB Standards S1 (general sustainability disclosures) and S2 (climate-related disclosures) are excluded on time-period coherence grounds. Both were finalised in 2023 and introduce disclosure vocabulary that did not exist during most of the 2013-2020 sample window. Including them would risk anchoring the measure on language that postdates the empirical activity being scored. A trimmed-anchor robustness analysis using only pre-2020 vocabulary is planned for the job-market version.

A.3 Curation prompt template

Every topic anchor is generated by a single canonical prompt template instantiated with topic-specific parameters from a CSV table. The full prompt for any topic is reconstructible from the template below plus the topic’s CSV row. The rationale for each design choice (action-language register, source-fidelity rule, acronym convention, per-topic isolation) is described in Section 3.2; this subsection records the literal template artefact for reproducibility. The curation model is Claude Opus 4.7 with a one-million-token context, used because it can read all source documents for a topic in a single pass.

1. Role and task. You are assisting in the construction of a reference corpus for a finance research project that measures sustainability orientation in job postings via embedding-based similarity scoring (bge-m3). The corpus is organised into a taxonomy of topic subcorpora. For each topic, you produce one synthesised prose file that captures the operational vocabulary and concept space of that topic, drawing on the listed source documents. This file will be embedded as a single dense vector and used as an anchor against which job postings are scored via cosine similarity. The objective is not to reproduce any source document, not to write a disclosure-style summary, and not to write an aspirational ESG narrative. The objective is to produce a dense operational description of the work being done in the topic domain, written in the register most likely to maximise semantic alignment with the language used in actual job postings for roles that perform this work.

2. Voice and register. Write in present-tense, agentive, descriptive prose that describes work being performed. The reader should be able to imagine a worker or team doing the activities described. In-register examples: *“Safety officers maintain a register of applicable legal, regulatory, and industry-specific OHS requirements, map each requirement to operational controls, and update the register when regulations change.”* *“Investigators reconstruct the sequence of events using witness interviews, scene examination, equipment records, and procedure reviews.”* *“Asset-integrity programmes inspect, test, and maintain pressure vessels, piping, rotating equipment, instrumented protective systems, and relief devices on a risk-based schedule.”* Out-of-register registers to avoid: aspirational or promissory (*“the company commits to leading the industry in safety excellence”*), disclosure or reporting (*“the organization shall disclose its OHS performance”*), standard-definition (*“occupational health and safety is defined as the discipline concerned with. . .”*), marketing or values (*“we believe in a culture of safety”*), and policy-declarative (*“workers have the right to refuse unsafe work”* as a standalone normative claim; acceptable when describing a documented procedural element). When source documents use out-of-register language, translate the underlying operational content into in-register prose: the framework says “the organization shall identify hazards”; the anchor reads “safety officers identify hazards through job hazard analyses that decompose each task into steps.”

3. Source-fidelity rule. You are synthesising across the listed sources for each topic, not extracting from them. The output is new prose that captures the operational vocabulary and concept structure that emerges from the union of sources for that topic. You may use operational vocabulary from the sources, recognised technical terms, named methodologies (HAZOP, LOPA, HACCP, RULA, REBA, NIOSH lifting equation, SBTi, FPIC, LEAP, MACC, double materiality), and standard procedural elements drawn from the sources. You may not reproduce verbatim passages, and you may not invent operational practices with no basis in the listed sources. When sources use different terminology for the same concept (lockout/tagout vs. energy isolation vs. control of hazardous energy), prefer the term most likely to appear in actual job postings; when in doubt, include both. When sources disagree substantively, default to the operational consensus and omit framing disagreements.

4. Time-period vocabulary handling. The job-posting sample spans 2013–2023; several admitted source documents (ESRS 2023, TNFD 2023, PCAF Parts B and C 2023–2025, NZBA v4 2025, SBTi 2026 versions) were finalised within or near the sample tail. These were admitted because they codify pre-existing operational practice. When such sources use language that is itself novel disclosure framing (e.g., “double materiality assessment” framed as a regulated reporting step rather than as the underlying stakeholder and risk analysis work), prefer the underlying operational content over the novel framing. Operational practice that existed prior to the source’s finalisation should sound contemporaneous with the full sample window, not exclusively contemporary.

5. Length. Target 1,700–2,000 tokens measured by the bge-m3 tokenizer (approximately 1,250–1,500 English words; will vary). Push toward the upper end only when the topic genuinely spans more sub-concepts than fit cleanly in the target range; do not pad. Going under 1,500 tokens suggests insufficient concept variation; expand. Do not exceed 4,096 tokens (hard cap; matches posting-side truncation).

6. Structure. No markdown headers in the body. No bolding, no bullet points, no numbered lists, no YAML frontmatter. Paragraph breaks only. Aim for 8–14 paragraphs, each a coherent operational sub-topic. The first paragraph should establish the management-system or organisational backbone of the topic; subsequent paragraphs should cover specific operational practices, each centred on one sub-domain. The order should flow naturally but does not need to follow any prescribed schema.

7. Jargon and technical depth. Use recognised operational vocabulary at the level a competent practitioner in the field would use. Named methodologies, standard acronyms, recognised procedural elements, and technical terms are encouraged. The intent is to score generic vocabulary (“*we care about safety*”) low and operational depth (“*manage HAZOP reviews, oversee LOTO compliance, coordinate Tier 2 process-safety-event investigations*”) high. Do not over-explain technical terms—postings for substantive roles assume the reader knows what HAZOP means. When a topic has both technical-industrial content and office/services-applicable content (OHS has industrial process safety and office ergonomics; DEI has both manufacturing-floor diversity work and professional-services inclusive recruitment), include both in the same operational register. Office-applicable content should use recognised operational vocabulary for that domain (DSE assessments, RULA, REBA, ERGs, structured interview rubrics, pay-equity regression analyses) rather than soft framing.

8. Industry-specific content. When operational content is meaningfully industry-specific (process safety in OHS for high-hazard industries; tailings management in environmental topics for mining; financed-emissions methodology in finance), include it within the topic file in the relevant paragraph rather than relegating it to sector rollups. The taxonomy uses max-of-topic-scores aggregation at scoring time, so sector-flavoured topic content does not double-count with separate sector files.

9. Variables filled at runtime. The following fields are substituted at generation time from `per-topic-parameters.csv`: `{{topic_id}}` (e.g., E1, S10, Sec1), `{{topic_name}}` (kebab-case slug, e.g., `climate-mitigation-ghg-accounting`), `{{pillar}}` (one of `environmental`, `social`, `governance`, `sectoral`), `{{concept_focus}}` (one-paragraph description of the operational domain), `{{source_files}}` (list of source-document paths), and `{{topic_notes}}` (topic-specific boundary clarifications and editorial caveats, may be blank).

10. Source documents for this topic. For each listed source in `{{source_files}}`, treat the file as the topic-relevant primary material. Where the parameters file annotates a source with an extraction note (e.g., “OECD MNE Guidelines—competition chapter only”), respect the scope. Where a source appears in multiple topic anchors, extract only the portion relevant to this topic and do not duplicate content that belongs to another topic.

11. Concept focus and boundary. `{{concept_focus}}`. `{{topic_notes}}`.

12. Output. Produce only the synthesised prose file. Do not include explanatory preamble, source citations within the prose, methodological commentary, frontmatter, or summary. The file should read as a self-contained operational description of the topic domain. Output path: `corpus-curated/topics/{{topic_id}}-{{topic_name}}.md`.

13. Acronym convention. For any acronym or initialism used in the topic file, include both the expanded form and the abbreviation at least once. Canonical convention: full form on first use with abbreviation in parentheses (Net-Zero Banking Alliance (NZBA), sustainable aviation fuel (SAF), marginal abatement cost curves (MACC)). After the first paired use, either form may appear; mixing in subsequent sentences is acceptable. Rationale: postings may use the full form, the abbreviation, or both, and the anchor must be retrievable against either. Exceptions accept the canonical form without expansion: standard chemical symbols (CO_2 , CH_4 , N_2O , SF_6 , NO_x , SO_x); standard identifiers used as proper nouns (ISO 14001, ISO 45001, GRI 305, ESRS E1, ISAE 3000, API RP 754); acronyms with no informative full form in current use (CDP); and universally known scientific units (GHG, GDP, R&D).

14. Per-topic isolation. When generating multiple topic files in the same working session, ground each new file in its own source list, concept focus, and topic notes. Do not pattern-match register, paragraph structure, opening sentence, or transitional phrasing to previously generated topic files. Each topic should be drafted as if from a cold start within the constraints of this template. Template-level conventions apply to every file (third-person operational voice, dense single-block format, action-verb register, acronym discipline); topic-level idioms do not transfer. Recycling idioms from previous topics inflates between-anchor cosine similarity in ways that depress discriminating power between adjacent topics.

15. Quality checklist. Before delivering, verify: first paragraph establishes the management-system or organisational backbone; 8–14 paragraphs, each a coherent sub-domain; length within 1,500–4,096 tokens, target 1,700–2,000; no aspirational verbs (commit, aim, strive, aspire,

pledge); no disclosure scaffolding (“shall report”, “Disclosure X-Y requires”); no first-person (“we”, “our”) or future tense; named methodologies and acronyms used at practitioner depth, not over-explained; every acronym appears at least once paired with its full form; office/services register included where the topic naturally carries it; industry-specific operational vocabulary woven into relevant paragraphs; no markdown headers, bullets, bold, or frontmatter in the output; output filename matches `{{topic_id}}-{{topic_name}}.md`; no idiom, opening, or transitional phrasing recycled from a previously generated topic file in this session.

A.4 Example topic anchor: S1 Occupational Health and Safety

The S1 anchor was the original pilot anchor under which the prompt template was developed. I reproduce a six-paragraph extract below as an illustration of the action-language register, covering the management-system backbone, hazard identification, permit-to-work systems, worker training, incident investigation, and the process-safety paragraph for high-hazard industries. The full anchor runs thirteen paragraphs at approximately 1,900 bge-m3 tokens; the omitted paragraphs cover contractor safety, occupational health surveillance, PPE programmes, emergency preparedness, psychosocial-risk management, manual-handling and ergonomic exposures (industrial and office, with RULA/REBA/NIOSH/DSE assessments), transport and fleet safety, and management review.

Occupational health and safety (OHS) work begins with a documented management system that defines policy, scope, roles, and accountability lines. Senior managers approve the policy, allocate budget, and assign ownership of OHS performance to operating-line managers rather than to a centralised compliance function. A joint health and safety committee with elected worker representatives meets on a fixed cadence, reviews incident data, signs off on risk-assessment outputs, and contributes to control selection. Workers retain the right to refuse work they reasonably believe is dangerous and exercise that right without retaliation; the procedure for raising and escalating refusals is documented in onboarding materials and posted at workstations. Safety officers maintain a register of applicable legal, regulatory, and industry-specific OHS requirements, map each requirement to operational controls, and update the register when regulations change.

Hazard identification covers routine and non-routine activities, planned changes, emergency situations, and the activities of contractors and visitors on site. Job hazard analyses (JHAs) decompose each task into steps, identify the hazards at each step (physical, chemical, biological, ergonomic, psychosocial), and assign controls. Risk assessments evaluate each identified hazard against likelihood and severity, account for the number and vulnerability of exposed workers, and rank residual risk after controls. The hierarchy of controls drives selection: elimination of the hazard at source, substitution with a less hazardous material or process, engineering controls such as machine guarding, local exhaust ventilation, isolation, interlocks, and fall-arrest anchorages, administrative controls such as safe-work procedures, permits, signage, shift-pattern adjustment, and rotation, and personal protective equipment (PPE) as the last layer. Engineering controls are designed in during process and facility design reviews and during management-of-change procedures, not retrofitted after deployment.

Permit-to-work systems govern high-risk activities. Hot-work, confined-space entry, working at height, energised-electrical work, excavation, lifting with cranes, and line-breaking each require a documented hazard review, verified atmospheric or other measurements, isolation

of relevant energy sources, a defined permit holder, time-bounded validity, and sign-off at issue and at close-out. Lockout/tagout (LOTO) procedures isolate hazardous energy before maintenance, with each isolation point individually locked, each lock owned by the worker performing the maintenance, and verification that stored energy has been dissipated before work begins. Confined-space entries gas-test before entry and at defined intervals during entry, post an attendant outside the space with continuous communication, and stage rescue equipment within reach.

Worker training is task-based, competency-verified, and refreshed on a defined cycle. Induction training covers site hazards, evacuation routes, reporting channels, and the right to refuse unsafe work. Task-specific training is delivered before a worker is permitted to perform the task, with demonstration of competence by the worker and sign-off by a qualified trainer. Refresher training is scheduled annually for high-hazard tasks and after any incident or near-miss involving the task. Records of training delivery and competence verification are retained and made available to workers on request. Contractor workers complete site induction before mobilising, and their employers provide evidence of equivalent task-specific training for the work they will perform.

Incident investigation triggers on any work-related injury, occupational illness, fatality, or near-miss with reasonable potential for harm. Investigators reconstruct the sequence of events using witness interviews, scene examination, equipment records, and procedure reviews, identify immediate and underlying causes through structured techniques such as five-whys or causal-tree analysis, and propose corrective actions assigned to named owners with target dates. Lessons learned are communicated to all workers performing similar tasks and to the joint health and safety committee. The OHS data system tracks lagging indicators (fatalities, lost-time injuries, recordable-injury rate, occupational illness rate, days lost) and leading indicators (near-miss reporting rate, hazard observations per worker, safety-walk completion rate, training compliance rate, corrective-action closure rate, and permit-audit pass rate).

High-hazard industries operate process safety programmes alongside occupational safety. Process hazard analyses (hazard and operability studies (HAZOP), what-if, layer-of-protection analysis (LOPA)) identify deviations from design intent and the consequences of equipment failure or human error. Asset-integrity programmes inspect, test, and maintain pressure vessels, piping, rotating equipment, instrumented protective systems, and relief devices on a risk-based schedule. Management-of-change procedures review modifications to equipment, materials, procedures, organisation, and software for OHS and environmental implications before implementation. Process-safety event rates (Tier 1 and Tier 2 events under API RP 754 or equivalent) are tracked alongside personal-injury metrics. [continues for six more paragraphs covering contractor safety, occupational health surveillance, PPE, emergency preparedness, psychosocial risk, ergonomics, transport and fleet safety, and management review]

The excerpt illustrates the three editorial rules described in Section 3.2: third-person agentive prose, dense use of named methodologies at practitioner depth, and integration of industrial and office-services content in the same operational register.

A.5 Alt-corpus construction

The alt-corpus is the methodological control against which the prompted curation is benchmarked. It uses the same 26 source families and the same 31-topic taxonomy as the curated corpus, but

skips the prompted-synthesis editorial step. Three reasons motivate building it. First, it identifies how much of the construct work is done by the source-document selection and taxonomy versus by the editorial register: anything the alt-corpus already does is attributable to the selection step, and anything the curated corpus does over and above the alt-corpus is attributable to the editorial step. Second, it provides a falsification benchmark: if the curated corpus performs no better than the alt-corpus on the substantive validation tests, the editorial step is not earning its complexity. Third, it gives readers a mechanical baseline they can re-construct without large-language-model dependencies, which makes the curation step’s incremental value visible.

The pipeline has three stages. Stage 1 extracts plain text from each source PDF with `pdfplumber` and applies mechanical cleaning: repeated headers and footers (lines appearing on more than 40 percent of pages with a lowercase-and-trim normalisation) are dropped across all pages; copyright notices and standalone page numbers are dropped; runs of three or more blank lines are collapsed to one; inline section identifiers such as "GRI 305-1", "ESRS E1-1", or "IFC PS6" are preserved.

Stage 2 chunks each cleaned text file against the bge-m3 fast tokenizer with a token cap of 3,800 (leaving 7 percent headroom under the 4,096 cap to absorb special-token overhead). Splitting priority is: pack whole paragraphs (split on double-newlines) until adding the next paragraph would exceed the cap; if a single paragraph is over the cap, split at sentence boundaries; if a single sentence is over the cap (rare, mostly tables rendered as one line), fall back to a hard token-window split. No overlap is allowed between chunks. The output is one JSON-Lines file per source with one row per chunk.

Stage 3 maps each chunk to one or more topic IDs via the same per-topic source-file table used by the curated curation, and writes a single parquet file with columns `topic_id`, `pillar`, `chunk_id`, `source_file`, `chunk_index`, `n_tokens`, and `text`. The parquet contains 2,160 chunks across the 31 topics.

The alt-corpus is then embedded with the same bge-m3 configuration as the curated anchors, and two aggregations are produced. Variant A is mean-pooled: chunk embeddings within each topic are L2-normalised and averaged, yielding 31 topic-level anchors that are directly comparable to the curated 31 anchors. Variant B is max-over-chunks: each chunk is treated as its own sub-anchor, and the topic-level similarity to a posting is the maximum cosine across that topic’s chunks. Variant A is the more natural structural parallel to the curated corpus; Variant B mechanically maximises coverage by retaining every chunk as a candidate match. Both are reported in the Stage 1 validation tables (Section A.7). The mechanical cleaning is deliberate: the alt-corpus does not attempt to filter aspirational language, disclosure scaffolding, or other registers we curated against, so any text the cleaning preserves is intentionally retained. This is what makes the head-to-head against the curated corpus an informative test of the editorial filter’s contribution.

A.6 Encoder comparison

Planned for the job-market version: a head-to-head encoder comparison against one domain-adapted alternative (an ESG-BERT variant used as an embedding backbone), to quantify how much of the LSI signal is robust to the encoder choice versus specific to the general-purpose

multilingual bge-m3 encoder used here.

A.7 Validation battery results

The Stage 1 validation battery establishes the construct from five independent angles: external discrimination against an occupation-level green-flag benchmark (T1.1), substantive discrimination between action-oriented and disclosure-oriented green occupations (T1.1.b), text-level value-add over occupation-level imputation (T1.4), head-to-head dominance over published keyword dictionaries (T1.3), and internal coherence of the topic-anchor geometry (T5 and T6.1). All validation tests run on a stratified posting-level sample of 250,000 postings balanced across calendar years 2013-2023.

The headline reading of the battery is the curated-versus-alt-corpus trade-off documented in Tables 55 and 56. The mechanical alt-corpus achieves slightly higher AUC against the O*NET binary green flag on the aggregated score (a -0.009 to -0.014 gap on `total_max`), but the curated corpus produces a Walk-versus-Talk discrimination gap approximately twice as large on every pillar with a positive overall gap. The O*NET green-flag benchmark is itself biased toward postings that share disclosure register with the source documents, so the alt-corpus over-performs precisely on the dimension where both share the same confound. The curated corpus does worse on this confounded benchmark and substantially better on the substantive Walk-versus-Talk benchmark, which I interpret as evidence that the action-language editorial filter is doing the discrimination work it was designed to do (Section 3.2).

A.8 Discriminant-validity supporting tables

The PPML horse race in Table 13 has a log-linear OLS counterpart, reproduced here for robustness. The full per-event-time Sun-Abraham coefficient table (over fifty cohort-by-event-time interactions) is summarised by the overall ATT reported in Table 12 and the event-time path plotted in Figure 2.

A.9 LSI-K sensitivity to depreciation rate

The headline G2 result in Section 5.7 uses a perpetual-inventory depreciation rate of $\delta = 20\%$, which matches the convention Peters and Taylor (2017) use for the organisation-capital stock built from capitalised SG&A. Their knowledge-capital stock, built from cumulated R&D, uses $\delta = 15\%$ as the slower-depreciating counterpart. Both rates are commonly adopted in the broader intangibles literature, with $\delta = 30\%$ as the upper-bound robustness for high-turnover intangibles. The natural sensitivity exercise is therefore to verify that the headline result survives the full set $\{15\%, 20\%, 30\%\}$. Table 53 reports the headline PT-decomposition regression under all three rates.

The qualitative pattern is robust across the three rates. The LSI-K share coefficient is positive and highly significant in every column, rising from $\beta = 42.64^{***}$ at $\delta = 15\%$ to $\beta = 47.80^{***}$ at $\delta = 20\%$ (the headline) and $\beta = 56.20^{***}$ at $\delta = 30\%$. The monotone rise is mechanical rather than substantive: a higher δ depreciates the stock faster, so a given share value corresponds to a smaller absolute stock, and the regression compensates with a larger per-share marginal

coefficient. The Peters-Taylor decomposition coefficients (q_{tot} , k_{int_org}/at , k_{int_know}/at) are stable across the columns, as expected because they do not depend on the depreciation rate applied to the LSI flow. The within- R^2 rises monotonically from 0.041 at $\delta = 15\%$ to 0.053 at $\delta = 30\%$, consistent with faster depreciation producing a sharper time-series signal of current sustainability orientation rather than a smoother long-run-average measure.

B Data sources and variable construction

This appendix records the vendor-specific implementation detail (database identifiers, field codes, pull procedures, crosswalk pipeline, and variable-construction recipes) that supports the conceptual data description in Section 4. The main text describes *what* each data source provides and *how* the regression specifications use it; this appendix records *exactly where the data come from* and *exactly how each variable is constructed* for transparency.

B.1 Posting cleaning pipeline

The cleaning pipeline applied to each LinkUp posting before embedding is implemented in `clean_postings.py`. The script concatenates the posting title and description (with a double-newline separator), then applies three light tiers of normalisation. Tier 0 (encoding normalisation) replaces Windows line endings, expands escaped-unicode sequences (`\u00e9` $\rightarrow$ `é`), unescapes HTML entities twice (`&amp;` $\rightarrow$ `&` $\rightarrow$ `&`), and applies a fixed table of mojibake corrections that handle the common Latin-1 $\rightarrow$ UTF-8 transcoding errors (e.g. the byte sequence rendered as `Ãe` maps back to `é`, the sequence rendered as `a-tilde-euro-trademark` maps back to an apostrophe, and similar fixes for accented characters and curly quotes). It then strips non-breaking spaces, soft hyphens, the Unicode replacement character (`\ufffd`), zero-width characters (`\u200b`–`\u200f`, `\u2028`–`\u202f`, `\u2060`–`\u206f`, `\ufeff`), and ASCII control characters (`\x00`–`\x08`, `\x0b`, `\x0c`, `\x0e`–`\x1f`, `\x7f`). Tier 1 (structural-noise removal) strips HTML tags, URLs (`http/https/www`), email addresses, North American and international phone numbers, US ZIP codes (preserving the leading state abbreviation), UK postcodes, and Canadian postcodes. Tier 4 (whitespace tidying) trims leading and trailing whitespace per line, collapses runs of spaces or tabs, removes lines composed only of decorative symbols (`*=_•|`), removes orphan navigation lines such as “Apply Now”, “Share this Job”, “Return to Job Search”, and collapses runs of three or more blank lines to a single blank-line separator. A final pass strips any unpaired UTF-16 surrogates that would otherwise break the parquet write.

The cleaning is deliberately light. Two heavier tiers are skipped to preserve substantively informative text: Tier 2 (HRIS labels such as “Job Code:”, “Posting ID:”, “Apply Online ID:”) is skipped because these labels rarely carry sustainability signal but their selective removal risks accidentally stripping substantive content with similar surface form; Tier 3 (boilerplate language about equal-opportunity statements, benefits descriptions, “About the Company” blocks) is skipped because the boilerplate often contains sustainability-relevant content (diversity statements, environmental commitments) that the construct should pick up. The length filter is minimal: keep any posting with at least one non-whitespace character after cleaning. The cleaning is applied year by year and produces the cleaned posting panel that is the input to the embedding pipeline. The retention rate from raw to cleaned is approximately 99 percent across all years; the dropped 1 percent are postings that were empty, all-HTML, or all-boilerplate after stripping.

B.2 LinkUp universe funnel

Table 1 in Appendix C reports the full funnel from the raw LinkUp universe to the estimation panel. The 2007–2023 LinkUp pull contains 257.9 million posting records across 61,519 unique employer entities. Approximately 22 percent of records are scrape-failure records for which LinkUp did not capture the full posting text before the original posting page was removed or expired, leaving 200.4 million description-bearing postings across 50,291 entities. Restricting to the 2013–2020 estimation window (the binding constraint on the consortium-draft compute budget; pre-2013 has insufficient coverage and post-2020 was deferred) yields the 60.7 million posting sample for the 10,232 ever-listed firms that comprise the estimation panel. The ever-listed restriction uses LinkUp’s point-in-time reference file (`pit_company`) which records the listed-status history of each entity throughout the 2007–2023 window. The 10,232 firms collectively published 42,734 firm-years with at least one description-bearing posting; this is the unit of analysis throughout Section 5. No exclusions are applied on country, financial-sector membership, or within-window delisting, to avoid survivorship bias in the forward-looking outcome regressions.

B.3 External data sources: vendor codes and pull details

Compustat (financial fundamentals). I pull annual fundamentals from `comp.funda` (North America) and `comp.g_funda` (Global). The variables pulled include the accounting items `AT`, `SALE`, `EMP`, `NI`, `DLTT`, `DLC`, `CAPX`, `XRD`, `XAD`, `PPENT`, `CEQ`, `CHE`, and `XSGA`; the market-valuation variables `MKVALT`, `PRCC_F`, and `CSHO`; and the Peters-Taylor intangible-capital variables `tobins_q_std`, `q_tot`, `k_int_know`, `k_int_org`, and `k_int`. Match to the LinkUp universe is via the LinkUp-to-gvkey crosswalk described in §B.4. Coverage: 6,671 firms across 26,865 firm-years (62.9 percent of LSI firm-years).

LSEG (Refinitiv) ESG. I pull LSEG ESG scores from LSEG Workspace using the Refinitiv Instrument Code (RIC) for each firm. The variables include the headline LSEG combined ESG score (`TR.TRESGScore`), the ESG controversies overlay (`TR.TRESGCControversiesScore`), the per-pillar environmental, social, and governance scores (`TR.TRESGE*PillarScore`, `TR.TRESGS*PillarScore`, `TR.TRESGCG*PillarScore`), and the CSR sustainability-reporting flag. LSEG currently distributes two related products: the LSEG ESG Score series (the rebranded continuation of the legacy Refinitiv methodology, available across the full 2013–2025 window for the combined score) and the newer Sustainability Ratings and Data series (launched with FY-2022 data and not back-filled). In the Workspace pull used here, the per-pillar topic-decomposition scores are present only from 2021 onwards, while the combined score covers the full window. This restricts the LSEG pillar-level convergent-validity tests (Section 5.1) to a post-2021 sub-sample too small for reliable estimation. The LSEG headline regressions in Section 5 therefore use the combined score on the full 2013–2020 window, and the pillar-level analyses use the S&P Global pillars (which do have pre-2021 coverage). Re-pulling the legacy pillar series through `TR.TRESG*PillarScore` is planned for the job-market version. Coverage: 4,839 firms across 18,825 firm-years (44.1 percent of LSI firm-years).

S&P Global Corporate Sustainability Assessment. I pull CSA scores from WRDS (`ciq_sustainability`), which provides question-level, criterion-level, dimension-level, and total ESG scores. I aggregate to the firm-year by extracting the dimension-level scores (Environmental, Social, Economic Governance) and the total ESG score for each firm-year-assessment combination. Coverage: 4,301 firms across 14,171 firm-years (33.2 percent of LSI firm-years). Pillar coverage extends through the full 2013–2020 window.

Reprisk ESG incidents. I use the Reprisk database of corporate ESG incidents, downloaded through WRDS. Reprisk monitors more than 100,000 sources in 23 languages and records each incident’s date, severity (1–3), reach (1–3), novelty (1–2), ESG-pillar classification (Environmental, Social, Governance, Cross-cutting), and detailed topic tags including a greenwashing flag. Match to the LinkUp universe is via primary ISIN. Coverage: 3,298 firms across 10,292 firm-years (24.1 percent of LSI firm-years). I construct per firm-year an incident count, a severity-weighted incident index ($\sum \text{severity} \times \text{reach} \times \text{novelty}$), pillar-specific incident counts (E, S, G, cross-cutting), and a high-severity count ($\text{severity} \geq 3$).

KPSS patents and green-patent classification. I use the [Kogan et al. \(2017\)](#) extended patent file (KPSS 2024 vintage), which covers all USPTO-issued patents from 1926 through 2024 with firm identifiers (CRSP `permno`) and patent economic values. I restrict to patents issued during 2013–2023 (the lead-1 outcome at the 2020 horizon requires 2021 patent data) and match to the LinkUp universe via the `permno-to-cid` crosswalk. Coverage: 1,658 firms across 6,596 firm-years (15.4 percent of LSI firm-years), necessarily restricted to US-listed patenting firms. Two parallel green-patent classifications are attached to each patent: the Cooperative Patent Classification (CPC) Y02 plus Y04S family (climate-change mitigation across adaptation, buildings, capture and storage, energy generation, industry, transport, waste; plus smart grids), which is the classification used by [Cohen et al. \(2026\)](#) and [Darendeli et al. \(2022\)](#); and the OECD ENV-TECH 2023 classification, which is broader and covers climate mitigation plus environmental management (pollution control, waste), water-related adaptation, and biodiversity, mapped to approximately 13,500 CPC codes. Both classifications are reported in the green-patent analysis (Section 5.5).

Sautner-et-al. climate-change exposure. I use the firm-quarter and firm-year climate-change-exposure measures of [Sautner et al. \(2023\)](#), downloaded from the authors’ OSF replication archive. The variables are total climate-change exposure (`cc_expo`) and the opportunity, regulatory, and physical sub-components (`op_expo`, `rg_expo`, `ph_expo`). Match is via `gvkey`. Coverage: 5,653 firms across 22,924 firm-years (53.6 percent of LSI firm-years).

BLS Occupational Employment and Wage Statistics. The wage-weighting procedure described in Section 3.7 uses the national mean annual wage series at the detailed 6-digit Standard Occupational Classification (SOC) level from the U.S. Bureau of Labor Statistics OEWS (formerly OES) annual release. This is the same wage series used in [Acemoglu et al. \(2022\)](#) and the broader postings-as-investment literature. The LinkUp postings carry an 8-digit O*NET occupation code; the mapping to 6-digit SOC is deterministic via truncation (the 8-digit O*NET

code is a child of a single 6-digit SOC parent). Approximately 98 percent of posting-firm-years receive a wage match; the residual 2 percent have O*NET codes that do not map to a detailed SOC (recently added occupations or non-standard custom codes). Nominal wages are deflated to constant 2017 dollars using the CPI-U series (annual averages, base 1982–1984 = 100). The choice of 2017 as the base year is the midpoint of the 2013–2020 sample window, which minimises the average deflation distance.

B.4 Identifier crosswalks

The LinkUp company identifier (`cid`) has no direct correspondence to any standard market-data identifier. Each external panel requires a crosswalk from `cid` to `gvkey` (Compustat), `permno` (CRSP), primary RIC (LSEG), and the supporting identifier set (ISIN, CUSIP, SEDOL, LEI, Organisation PermID). I construct four canonical crosswalk files in a three-phase pipeline. Phase 1 queries WRDS directly using each `cid`’s most recent primary ticker and exchange (`comp.security`, `comp.g_security`, `ciq.wrds_isin`, `ciq.wrds_cusip`, `crsp.ccmxpf_linktable`). Phase 2 uses LSEG Workspace to retrieve the canonical identifier set for each `cid`’s RIC, then feeds the LSEG-returned ISIN, CUSIP, or SEDOL back into WRDS for a second-pass match. Phase 3 applies fuzzy-name matching with the `rapidfuzz` library at token-set-ratio threshold 65 for residual unmatched `cids`, followed by manual review of borderline candidates.

The final crosswalk covers 7,622 of 10,232 `cids` (74.5 percent) with at least one validated `gvkey`, mapping to 5,961 unique `gvkeys`. The `cid`-to-`gvkey` ratio of 1.28 reflects the prevalence of parent-subsidiary structures in LinkUp’s employer-entity identification: multiple `cids` can correspond to the same Compustat parent (for example, separate `cids` for IBM Corporation, IBM Research, and IBM Global Services link to `gvkey` 006066). The 25.5 percent of unmatched `cids` fall into three categories: (i) firms identified as listed in LinkUp’s reference file but with no active record in WRDS or LSEG during 2013–2025 (typically private companies briefly listed in earlier decades or LinkUp reference errors); (ii) firms on minor international exchanges absent from WRDS Global coverage; and (iii) firms whose listing periods predate Compustat’s coverage of their market.

B.5 Variable construction recipes

Firm-fundamental controls are constructed at the firm-year level from Compustat as follows. Total assets is `AT`, in millions of dollars; the log is $\ln AT$. Leverage is $(DLTT + DLC) / AT$. Return on assets (`roa`) is NI / AT . R&D intensity (`rd_int`) is $XRD / SALE$ with zero imputation for missing `XRD`. Capital-expenditure intensity (`capx_int`) is $CAPX / AT$. Firm age ($\ln(1 + age)$) is computed from the first non-missing fiscal-year-end appearance in Compustat. Tobin’s q (`tobins_q_std`) and augmented Peters-Taylor q (`q_tot`) are taken directly from the Peters-Taylor intangible-capital release. Labour productivity ($\ln labour\ prod$) is $\ln(SALE/EMP)$ for firm-years with non-missing `EMP`.

Continuous regressors are winsorised at the 1st and 99th percentiles within the estimation panel to limit the influence of extreme values, following the convention in the postings-as-investment literature. The Peters-Taylor decomposition share variable used in the investment- q analysis (Section 5.7) is $lsik_share = LSI-K_{i,t} / (k_int_{i,t} + LSI-K_{i,t})$, a well-defined ratio on

$[0, 1]$. The LSI-K stock is computed from the wage-weighted dollar flow described in Section 3.7 via the perpetual-inventory rule $K_t = (1 - \delta)K_{t-1} + I_t$ with $K_{t_0} = I_{t_0}/(g + \delta)$, where $g \approx 0.014$ is the pooled mean log-growth of the flow series and $\delta = 0.20$ is the headline depreciation rate (sensitivity to $\delta \in \{0.15, 0.30\}$ is reported in Appendix A.9).

C Supplementary Tables and Figures

This appendix collects the robustness tables, alternative-classification results, alternative fixed-effect specifications, and pillar-decomposition tables that are referenced from the main text but not displayed inline. Headline tables and figures appear inline within Sections 4 and 5.

Construct and validation

Table 20: Reference corpus: 31 sustainability topic anchors

ID	Pillar	Topic
E1	Environmental	Climate-change mitigation and GHG accounting
E2	Environmental	Climate physical and transition risk
E3	Environmental	Energy management
E4	Environmental	Water and marine stewardship
E5	Environmental	Biodiversity, ecosystems, nature
E6	Environmental	Circular economy, materials, waste
E7	Environmental	Pollution prevention, chemicals
E8	Environmental	Environmental management systems
S1	Social	Occupational health and safety
S2	Social	Workforce development and training
S3	Social	Diversity, equity, and inclusion
S4	Social	Labour rights and collective bargaining
S5	Social	Forced and child labour, modern slavery
S6	Social	Human-rights due diligence
S7	Social	Community engagement, indigenous rights
S8	Social	Supply-chain social due diligence
S9	Social	Customer and product safety, responsible marketing
S10	Social	Data privacy, information security, AI governance
S11	Social	Just transition
S12	Social	Employment quality and working conditions
G1	Governance	Anti-corruption and bribery
G2	Governance	Fair competition, market conduct
G3	Governance	Tax responsibility and transparency
G4	Governance	Business conduct, ethics, culture
G5	Governance	Supply-chain governance
G6	Governance	Sustainability governance, board oversight
G7	Governance	Political engagement, lobbying disclosure
G8	Governance	ESG risk management, internal controls
G9	Governance	Materiality, stakeholder engagement
Sec1	Sectoral	Responsible finance stewardship
Sec2	Sectoral	Green buildings

Notes. The 31 topics are organised into four anchor groups (E, S, G, Sectoral). Per-topic anchor texts are curated from a set of 26 sustainability-standard source families (full inventory in Appendix A.2) under the action-language editorial rule defined in Section 3.2.

Table 21: Face validity: illustrative top-decile and bottom-decile postings by `total_max`

Year	Job title	Excerpt (first 200 characters)	LSI
<i>Panel A. Top-decile postings (substantive sustainability work)</i>			
2022	Senior Diversity, Equity, and Inclusion Program Manager (Amazon)	DEI Program Manager. We are looking for someone who can build inclusive cultures, design programmes that move our DEI strategy forward...	0.751
2021	Risk - Climate Risk - Scenario Analysis Lead, Director	Lead Independent Risk Management's efforts to develop Climate Risk Scenario Analysis; develop and implement scenario analysis...	0.739
2021	Sustainability Data and Reporting Lead (GSK Consumer Health)	Lead the Sustainability Data and Reporting agenda; define a Sustainability strategy for GSK CH; drive the reporting...	0.736
2022	Corporate Environmental, Social, and Governance Analyst (Spectrum Brands)	Support ESG strategy, data collection, and reporting across the enterprise; partner with Finance, Legal, IT, HR...	0.735
2020	Manager, Environmental, Health and Safety (LGC Group)	Manage the EHS programme across multiple sites; ensure compliance with regulatory requirements; lead incident investigations...	0.708
<i>Panel B. Bottom-decile postings (no substantive sustainability content)</i>			
2015	KILZ Brand Ambassador (Shawnee, MO)	KILZ Brand Ambassador.	0.260
2018	Overnight Baker	Overnight Baker 601539, Chicago IL - North Michigan Avenue.	0.280
2021	Swabbing Couriers	Swabbing Couriers, BRL-21-9368, Chicago IL.	0.280
2016	Fry, TXRH New Hartford	TXRH New Hartford is looking to hire a Fry.	0.313
2016	Porter (Columbus)	Porter, Columbus.	0.301

Notes. Illustrative postings drawn from the top-decile and bottom-decile of the cross-sectional `total_max` distribution on the stratified 250,000-posting validation sample. Top postings are systematically substantive sustainability roles spanning climate-risk strategy, EHS management, DEI programmes, sustainability reporting, and corporate ESG analysis.

Convergent validity

Table 22: LSI converges with the S&P Global ESG total score

Dependent Variable: Model:	sp_esg_total		
	(1)	(2)	(3)
<i>Variables</i>			
lsi_mean	58.67*** (7.732)	49.55*** (8.185)	50.04*** (8.139)
ln_at	5.395*** (0.2010)	5.839*** (0.2141)	5.618*** (0.2174)
leverage	0.1152 (1.258)	-2.318* (1.292)	-1.994 (1.287)
roa	11.14*** (2.205)	8.103*** (2.191)	6.181*** (2.191)
rd_int	29.37*** (4.882)	16.70*** (5.134)	17.54*** (5.176)
capx_int	40.10*** (7.969)	14.27 (9.682)	15.28 (9.674)
ln_age	-8.201*** (0.3527)	-7.654*** (0.3467)	10.86*** (1.878)
<i>Fixed-effects</i>			
naics2		Yes	Yes
year			Yes
<i>Fit statistics</i>			
Within R ²		0.36004	0.28124
R ²	0.33613	0.38131	0.40756
Observations	12,265	12,165	12,165

Clustered (company_id) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. The dependent variable is the S&P Global Corporate Sustainability Assessment total ESG score (**sp_esg_total**), pulled from WRDS **ciq_sustainability** and described in Section 4.3. The regressor of interest is **lsi_mean**.

Table 23: Pillar-by-pillar convergence of LSI with S&P Global ESG pillars

Dependent Variables: Model:	sp_esg_e (1)	sp_esg_s (2)	sp_esg_g (3)
<i>Variables</i>			
E_mean	53.89*** (10.70)		
ln_at	7.975*** (0.2849)	5.583*** (0.2524)	4.295*** (0.1726)
leverage	-3.576** (1.786)	-1.159 (1.432)	0.1784 (1.071)
roa	8.787*** (3.019)	4.452* (2.450)	5.799*** (1.819)
rd_int	21.87*** (7.426)	10.18* (5.620)	9.709** (4.166)
capx_int	24.12* (13.00)	20.67* (10.84)	6.369 (7.793)
ln_age	12.37*** (2.586)	9.574*** (1.779)	10.87*** (1.668)
S_mean		47.31*** (9.095)	
G_mean			39.58*** (6.211)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Within R ²	0.29952	0.23976	0.24670
R ²	0.35857	0.36388	0.47252
Observations	12,165	12,165	12,165
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. Each column regresses one S&P Global ESG pillar score (Environmental, Social, Economic Governance) on the corresponding LSI pillar mean (E_mean, S_mean, G_mean, constructed as the firm-year mean of the posting-level per-pillar max as described in Section 3.6). Specifications carry industry-by-year fixed effects and the headline control set, with firm-clustered standard errors in parentheses.

Table 24: LSI E-pillar converges with Sautner et al. (2023) climate-change exposure

Dependent Variables: Model:	sautner_cc_expo (1)	sautner_op_expo (2)	sautner_rg_expo (3)	sautner_ph_expo (4)
<i>Variables</i>				
E_mean	0.0056*** (0.0009)	0.0017*** (0.0004)	0.0004*** (7.43×10^{-5})	8.23×10^{-5} *** (3.11×10^{-5})
ln_at	2.51×10^{-5} * (1.46×10^{-5})	1.5×10^{-5} ** (6.35×10^{-6})	3.67×10^{-6} ** (1.48×10^{-6})	-2.21×10^{-6} * (1.26×10^{-6})
leverage	-0.0001 (9×10^{-5})	-5.69×10^{-5} (4.39×10^{-5})	-1.3×10^{-5} * (6.93×10^{-6})	4.27×10^{-6} (4.82×10^{-6})
roa	-0.0004** (0.0002)	-0.0002** (7.78×10^{-5})	-4.69×10^{-5} *** (1.17×10^{-5})	1.33×10^{-5} * (7.26×10^{-6})
rd_int	-0.0019*** (0.0004)	-0.0005*** (0.0002)	-0.0002*** (3.38×10^{-5})	-6.01×10^{-5} *** (1.67×10^{-5})
capx_int	0.0041*** (0.0010)	0.0020*** (0.0005)	7.94×10^{-5} (6.99×10^{-5})	-2.43×10^{-5} (2.95×10^{-5})
ln_age	0.0001 (0.0002)	1.29×10^{-5} (0.0001)	2.38×10^{-5} (1.48×10^{-5})	5.02×10^{-6} (5.33×10^{-6})
<i>Fixed-effects</i>				
naics2	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Within R ²	0.01795	0.00990	0.00975	0.00270
R ²	0.37819	0.29659	0.16087	0.01236
Observations	20,397	20,397	20,397	20,397

Clustered (company_id) standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Notes. The dependent variables are the four Sautner et al. (2023) firm-year climate-exposure measures derived from earnings-call transcripts: total climate-change exposure (cc_expo), and its opportunity, regulatory, and physical components (op_expo, rg_expo, ph_expo). The regressor of interest is E_mean, the firm-year mean of the posting-level environmental-pillar maximum cosine.

Lead-indicator divergence: alternative specifications and pillars

Table 25: Variance decomposition of firm-year LSI across firm, industry, and year components

Component	Sum of squares	Share of total
Industry (NAICS-2)	6.788	0.117
Year	0.644	0.011
Firm	46.934	0.805
Total	58.276	1.000

Notes. The decomposition is computed on the 2013-2020 firm-year panel restricted to observations with non-missing LSI and NAICS-2 (the LSI base of 42,734 firm-years collapses to 42,342 once the NAICS-2 filter is applied). The component sum of squares is computed group-by-group as $\sum_g n_g(\bar{x}_g - \bar{x})^2$, where $\bar{x}_g$ is the group mean (industry, year, or firm), $\bar{x}$ is the panel-wide mean, and n_g is the group's firm-year count.

Table 26: Firm characteristics by LSI-vs-rating quadrant, S&P Global total ESG score

Quadrant	N	Mean LSI	Mean rating	Mean assets	Mean ROA	Mean R&D	Mean leverage
HL_HR (high LSI, high rating)	3,138	0.574	58.60	119,608	0.047	0.023	0.278
HL_LR (high LSI, low rating)	4,833	0.572	23.80	28,279	0.020	0.037	0.288
LL_HR (low LSI, high rating)	2,236	0.518	57.32	99,426	0.048	0.015	0.286
LL_LR (low LSI, low rating)	3,964	0.515	23.12	47,604	0.037	0.020	0.292

Notes. Quadrants are defined on the firm-year panel with non-missing LSI and non-missing S&P Global total ESG score, following the same standardisation convention as Table 9. The four-quadrant population is 14,171 firm-years, the full S&P-Global-merged panel.

Table 27: Temporal precedence: LSI predicts subsequent changes in the S&P Global total ESG score at horizons $k = 1, 2, 3$

Dependent variable Model	$\Delta\text{S\&P}_{t+1}$ (1)	$\Delta\text{S\&P}_{t+2}$ (2)	$\Delta\text{S\&P}_{t+3}$ (3)
<i>Variables</i>			
LSI mean ($\text{LSI}_{i,t}$)	6.34*** (2.23)	12.97** (5.24)	21.67*** (7.99)
S&P ESG total (t)	-0.021*** (0.004)	-0.038*** (0.009)	-0.058*** (0.016)
$\ln(\text{AT})$	-0.017 (0.048)	0.047 (0.111)	0.241 (0.186)
Leverage	-0.067 (0.317)	0.554 (0.769)	0.682 (1.251)
ROA	1.396* (0.738)	3.654** (1.749)	4.919* (2.807)
R&D intensity	1.067 (1.235)	5.419* (3.290)	9.321* (5.247)
Capex intensity	0.969 (2.225)	-1.786 (5.036)	-7.567 (8.113)
$\ln(1 + \text{age})$	-0.316 (1.465)	4.161 (5.216)	11.291 (7.238)
<i>Fixed effects</i>			
NAICS-2 $\times$ Year	Yes	Yes	Yes
Observations	8,427	5,763	4,200
R^2	0.348	0.264	0.093
R^2 within	0.007	0.011	0.015

Notes. The specification mirrors Table 10 but uses the S&P Global Corporate Sustainability Assessment total ESG score in place of the LSEG combined ESG score. The pattern is qualitatively identical: the LSI coefficient is positive and significant at every horizon and rises monotonically from 6.34 at $k = 1$ to 21.67 at $k = 3$.

Table 28: Fixed-effect progression at horizon $k = 1$: cross-sectional precedence versus within-firm null, LSEG combined ESG

Dependent variable: ΔLSEG_{t+1} Model	Pooled OLS (1)	+ Industry FE (2)	+ Ind $\times$ Year FE (3)	+ Firm + Year FE (4)
<i>Variables</i>				
LSI mean ($\text{LSI}_{i,t}$)	11.51*** (2.20)	9.38*** (2.37)	9.62*** (2.36)	6.35 (5.29)
LSEG ESG (t)	-0.151*** (0.005)	-0.159*** (0.006)	-0.157*** (0.006)	-0.757*** (0.015)
$\ln(\text{AT})$	0.368*** (0.049)	0.456*** (0.053)	0.442*** (0.053)	1.362*** (0.460)
Leverage	1.157*** (0.361)	0.629 (0.386)	0.533 (0.382)	-0.143 (1.300)
ROA	4.579*** (0.859)	4.064*** (0.924)	4.356*** (0.884)	2.523* (1.467)
R&D intensity	2.502* (1.325)	0.307 (1.494)	0.187 (1.482)	16.206*** (4.706)
Capex intensity	-0.151 (2.017)	-3.861 (2.545)	-4.220* (2.525)	5.515 (6.113)
$\ln(1 + \text{age})$	0.194 (0.153)	0.292* (0.155)	-1.804* (0.925)	14.578*** (4.708)
<i>Fixed effects</i>				
NAICS-2		Yes	Yes	
Year			Yes	Yes
Firm				Yes
Observations	12,360	12,312	12,312	11,766
R^2	0.077	0.083	0.090	0.447
R^2 within		0.081	0.081	0.349

Notes. Each column reports an OLS regression of the one-year-ahead change in the LSEG combined ESG score on contemporaneous LSI, the contemporaneous rating level, and the standard control set. The four columns vary only in the fixed-effect structure: column (1) is pooled OLS with no fixed effects; column (2) adds NAICS-2 industry fixed effects; column (3) is the headline NAICS-2 by year specification reported in column (1) of Table 10; column (4) replaces industry-by-year with firm and year fixed effects, which removes the cross-sectional between-firm variation and identifies the LSI coefficient from within-firm year-to-year variation only.

Table 29: Fixed-effect progression at horizon $k = 1$: cross-sectional precedence versus within-firm null, S&P Global ESG

Dependent variable: $\Delta S\&P_{t+1}$ Model	Pooled OLS (1)	+ Industry FE (2)	+ Ind $\times$ Year FE (3)	+ Firm + Year FE (4)
<i>Variables</i>				
LSI mean (LSI _{<i>i,t</i>})	8.47*** (2.18)	8.62*** (2.40)	6.34*** (2.23)	0.46 (4.33)
S&P ESG total (<i>t</i>)	−0.055*** (0.004)	−0.056*** (0.004)	−0.021*** (0.004)	−0.378*** (0.017)
ln(AT)	−0.005 (0.049)	0.001 (0.053)	−0.017 (0.048)	0.260 (0.409)
Leverage	0.381 (0.344)	0.526 (0.365)	−0.067 (0.317)	0.478 (1.243)
ROA	−0.807 (0.770)	−0.455 (0.791)	1.396* (0.738)	2.274 (1.763)
R&D intensity	1.141 (1.313)	1.549 (1.453)	1.067 (1.235)	0.944 (7.389)
Capex intensity	1.918 (1.833)	3.845 (2.456)	0.969 (2.225)	3.833 (5.461)
ln(1 + age)	0.482*** (0.117)	0.481*** (0.117)	−0.316 (1.465)	8.999* (5.351)
<i>Fixed effects</i>				
NAICS-2		Yes	Yes	
Year			Yes	Yes
Firm				Yes
Observations	8,472	8,427	8,427	7,258
R^2	0.036	0.039	0.348	0.588
R^2 within		0.036	0.007	0.177

Notes. The specification mirrors Table 28 but uses the S&P Global total ESG score. The same pattern obtains under the second rating provider: the LSI coefficient is significant at the 1% level across columns (1) to (3) at magnitudes 6-9 and collapses to an insignificant 0.46 in column (4) once firm fixed effects absorb the between-firm variation.

Table 30: Extended-control robustness at horizon $k = 1$: LSEG combined ESG

Dependent variable: ΔLSEG_{t+1} Model	Extended (Ind $\times$ Yr) (1)	Extended (Firm + Yr) (2)
<i>Variables</i>		
LSI mean ($\text{LSI}_{i,t}$)	10.67*** (2.46)	5.71 (5.41)
LSEG ESG (t)	-0.166*** (0.006)	-0.766*** (0.016)
$\ln(\text{AT})$	0.191** (0.088)	1.255** (0.618)
$\ln(\text{EMP})$	0.435*** (0.096)	0.387 (0.645)
Leverage	0.270 (0.405)	0.006 (1.349)
ROA	3.238*** (0.948)	2.585* (1.445)
R&D intensity	-0.533 (1.674)	15.504*** (4.861)
Capex intensity	-4.615 (3.262)	3.591 (6.206)
Advertising intensity	-4.037 (3.499)	10.672 (17.578)
Tangibility	0.034 (0.693)	0.111 (2.722)
Book-to-market	-0.105* (0.054)	-0.192 (0.201)
$\ln(1 + \text{age})$	-1.619 (0.989)	20.115*** (6.067)
Tobin's q	0.114** (0.050)	-0.063 (0.119)
<i>Fixed effects</i>		
NAICS-2 $\times$ Year	Yes	
Firm		Yes
Year		Yes
Observations	11,850	11,314
R^2	0.095	0.451
R^2 within	0.085	0.354

Notes. The extended control set adds Tobin's q (a growth-opportunity proxy), $\ln(\text{EMP})$ (a scale proxy independent of book assets), book-to-market (a value-versus-growth proxy), tangibility (PPE over assets), and advertising intensity (XAD over sales) to the standard control set in Table 28. Column (1) is the NAICS-2 by year specification; column (2) is the firm and year specification.

Table 31: Extended-control robustness at horizon $k = 1$: S&P Global ESG

Dependent variable: $\Delta S\&P_{t+1}$ Model	Extended (Ind $\times$ Yr) (1)	Extended (Firm + Yr) (2)
<i>Variables</i>		
LSI mean (LSI $_{i,t}$)	5.92*** (2.22)	0.85 (4.45)
S&P ESG total (t)	-0.022*** (0.004)	-0.383*** (0.018)
ln(AT)	-0.020 (0.072)	0.484 (0.611)
ln(EMP)	-0.005 (0.074)	-0.138 (0.519)
Leverage	-0.157 (0.326)	0.392 (1.376)
ROA	1.750** (0.846)	1.886 (1.769)
R&D intensity	1.552 (1.556)	-3.217 (7.688)
Capex intensity	2.061 (2.870)	-0.494 (5.490)
Advertising intensity	-6.958** (2.819)	24.289 (14.769)
Tangibility	0.122 (0.571)	4.434 (2.717)
Book-to-market	-0.038 (0.052)	-0.147 (0.113)
ln(1 + age)	-0.819 (1.476)	9.965* (5.651)
Tobin's q	-0.025 (0.046)	0.094 (0.115)
<i>Fixed effects</i>		
NAICS-2 $\times$ Year	Yes	
Firm		Yes
Year		Yes
Observations	8,140	7,016
R^2	0.352	0.593
R^2 within	0.009	0.180

Notes. Extended-control specification with the same variable set as Table 30, applied to the S&P Global total ESG score. The LSI coefficient in column (1) is 5.92 and significant at the 1% level, marginally below the 6.34 standard-control coefficient in Table 29 column (3).

Table 32: Realised outcomes by LSI-vs-rating quadrant, S&P Global total ESG score

Quadrant	N	Reprisk $t + 1$	Reprisk $t + 2$	Reprisk sev. $t + 1$	Green pat. $t + 1$	Non-green pat. $t + 1$
HL_HR (high LSI, high rating)	3,138	16.72	16.67	51.05	22.90	258.19
HL_LR (high LSI, low rating)	4,833	7.12	7.18	21.35	3.04	58.67
LL_HR (low LSI, high rating)	2,236	15.86	15.89	48.59	12.79	257.34
LL_LR (low LSI, low rating)	3,964	10.00	10.59	29.31	1.88	65.38

Notes. The four-quadrant assignment is identical to Table 26. The green-patent separation is even sharper under the S&P Global split than under the LSEG split: HL_HR firms produce a mean of 22.9 green patents in the year ahead, roughly 7.5 times the HL_LR mean of 3.0 and 12 times the LL_LR mean of 1.9.

Reprisk-incident robustness

Table 33: Severity-weighted Reprisk outcomes, PPML

Dependent Variables: Model:	rep_sev_index_lead1 (1)	rep_high_sev_lead1 (2)	rep_high_sev_lead1 (3)
<i>Variables</i>			
lsi_mean	0.7345 (0.7087)	0.5775 (0.7353)	1.182 (1.766)
lseg_esg	-0.0076*** (0.0016)	-0.0085*** (0.0016)	-3.91×10^{-5} (0.0032)
ln_at	0.7130*** (0.0183)	0.7171*** (0.0224)	0.7697*** (0.0474)
leverage	0.4768*** (0.1608)	0.5495*** (0.1736)	-0.3067 (0.2837)
roa	1.976*** (0.4146)	2.082*** (0.4755)	4.571*** (1.062)
rd_int	5.120*** (0.6645)	5.205*** (0.7033)	1.785 (1.411)
capx_int	2.730*** (0.9371)	3.115*** (1.012)	0.4460 (1.793)
ln_age	-2.012*** (0.4577)	-1.645*** (0.4967)	-2.213** (0.9818)
sp_esg_total		0.0015 (0.0018)	
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Squared Correlation	0.67585	0.67426	0.27809
Pseudo R ²	0.65715	0.65554	0.26983
BIC	138,350.3	122,841.0	4,638.6
Observations	5,797	4,784	5,699

Clustered (company_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Dependent variables: columns (1) and (2) report the severity-weighted incident index $\sum_j \text{severity}_j \times \text{reach}_j \times \text{novelty}_j$ summed over Reprisk incidents in year $t + 1$ (**rep_sev_index_lead1**); column (3) reports the count of high-severity incidents ($\text{severity} \geq 3$) in year $t + 1$ (**rep_high_sev_lead1**). The severity-weighting allows the test to attach more weight to incidents that the Reprisk methodology classifies as high consequence, addressing the concern that a null on raw counts might mask a substantive null on severity-weighted exposure.

Table 34: Firm and year fixed-effects (within-firm-over-time) PPML

Dependent Variable:	rep_incidents_lead1	
Model:	(1)	(2)
<i>Variables</i>		
lsi_mean	-0.4457 (0.4281)	-0.5494 (0.4349)
lseg_esg	-0.0009 (0.0008)	-0.0012 (0.0008)
ln_at	0.3830*** (0.0636)	0.3875*** (0.0683)
leverage	0.0389 (0.1335)	0.0986 (0.1379)
roa	0.3176* (0.1885)	0.3192 (0.2007)
rd_int	1.464 (1.416)	2.573 (1.691)
capx_int	0.5618 (0.7110)	0.6271 (0.7212)
ln_age	-0.3023 (0.2131)	-0.1135* (0.0579)
sp_esg_total		0.0022 (0.0014)
<i>Fixed-effects</i>		
company_id	Yes	Yes
year	Yes	Yes
<i>Fit statistics</i>		
Squared Correlation	0.95982	0.95982
Pseudo R ²	0.82814	0.83026
BIC	35,058.2	29,419.0
Observations	5,257	4,312
<i>Clustered (company_id) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Notes. The specification replaces the NAICS-2-by-year fixed effects of Table 13 with firm and year fixed effects, so the identifying variation is within-firm-over-time. The dependent variable is the firm-year count of Reprisk incidents in year $t + 1$.

Table 35: Horizon robustness, PPML on Reprisk incident count at lead $k = 1, \dots, 5$

Dependent Variables: Model:	rep_incidents_lead1 (1)	rep_incidents_lead2 (2)	rep_incidents_lead3 (3)	rep_incidents_lead4 (4)	rep_incidents_lead5 (5)
<i>Variables</i>					
lsi_mean	0.6700 (0.7450)	0.9659 (0.7825)	1.192 (0.9111)	1.118 (1.007)	1.409 (1.121)
lseg_esg	-0.0076*** (0.0017)	-0.0076*** (0.0018)	-0.0067*** (0.0022)	-0.0073*** (0.0026)	-0.0084*** (0.0028)
ln_at	0.7395*** (0.0196)	0.7366*** (0.0209)	0.7318*** (0.0233)	0.7363*** (0.0268)	0.7441*** (0.0309)
leverage	0.5783*** (0.1672)	0.5035*** (0.1843)	0.3982* (0.2109)	0.2547 (0.2558)	0.0511 (0.3148)
roa	1.771*** (0.4211)	1.883*** (0.4936)	1.730*** (0.5512)	2.343*** (0.6574)	1.507* (0.7798)
rd_int	5.366*** (0.7429)	5.663*** (0.8688)	6.059*** (0.9592)	6.531*** (1.256)	6.830*** (1.478)
capx_int	2.947*** (0.9820)	3.233*** (1.033)	3.696*** (1.247)	3.126** (1.392)	3.101 (1.930)
ln_age	-1.972*** (0.4852)	-1.257** (0.6167)	-0.5031 (0.6096)	-1.125*** (0.2182)	-0.7238*** (0.2138)
<i>Fixed-effects</i>					
naics2	Yes	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>					
Squared Correlation	0.68283	0.68247	0.66829	0.65958	0.64866
Pseudo R ²	0.62406	0.61913	0.61114	0.59861	0.58681
BIC	56,686.1	44,286.4	33,984.0	25,023.7	17,254.5
Observations	5,797	4,511	3,333	2,329	1,515

Clustered (company_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Each column estimates the same PPML horse-race specification as column (2) of Table 13 but moves the dependent variable to the firm-year count of Reprisk incidents at horizon $t + k$ for $k = 1, \dots, 5$. The horizon extension addresses the concern that the one-year-ahead null might mask a slower causal effect of LSI on subsequent incidents working through capability build-out or operational implementation.

Table 36: Pillar-decomposition PPML, pillar-specific incident counts on pillar-specific LSI

Dependent Variables: Model:	rep_e_count_lead1 (1)	rep_s_count_lead1 (2)	rep_g_count_lead1 (3)
<i>Variables</i>			
E_mean	1.430 (1.091)		
lseg_esg	0.0016 (0.0023)	-0.0057*** (0.0020)	-0.0113*** (0.0017)
ln_at	0.7947*** (0.0322)	0.6847*** (0.0238)	0.7829*** (0.0211)
leverage	0.6433** (0.2502)	0.2781 (0.1931)	0.3938** (0.1808)
roa	1.265** (0.5047)	1.957*** (0.5138)	2.107*** (0.5657)
rd_int	-1.797 (1.997)	4.987*** (0.9156)	6.336*** (0.5847)
capx_int	3.556*** (1.253)	3.322*** (1.083)	1.738* (1.053)
ln_age	-0.7423 (0.5649)	-2.273*** (0.5300)	-1.970*** (0.5502)
S_mean		0.7852 (0.8428)	
G_mean			1.465* (0.7492)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Squared Correlation	0.67326	0.58679	0.69996
Pseudo R ²	0.53907	0.50820	0.56635
BIC	27,806.5	41,800.7	32,153.8
Observations	5,797	5,797	5,797
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. The specification matches the cross-pillar diagonal: $E_{i,t}$ is regressed on the lead-1 count of environmental incidents (**rep_e_count_lead1**), $S_{i,t}$ on the lead-1 count of social incidents (**rep_s_count_lead1**), and $G_{i,t}$ on the lead-1 count of governance incidents (**rep_g_count_lead1**). The pillar-specific decomposition addresses the concern that the headline null in Table 13 might mask a substantive pillar-on-pillar effect that gets averaged out by the cross-pillar aggregation.

Table 37: log-linear OLS counterpart to the PPML horse race (robustness)

Dependent Variable:	ln_inc1		
Model:	(1)	(2)	(3)
<i>Variables</i>			
lsi_mean	0.1652 (0.4598)	0.2088 (0.4918)	-0.0562 (0.5478)
ln_at	0.4399*** (0.0126)	0.5047*** (0.0156)	0.5410*** (0.0185)
leverage	0.0942 (0.0852)	0.2170** (0.1029)	0.4601*** (0.1190)
roa	0.0704 (0.1601)	0.5175** (0.2134)	1.049*** (0.2771)
rd_int	2.559*** (0.4206)	2.523*** (0.4862)	2.496*** (0.5787)
capx_int	1.295** (0.5220)	1.533** (0.6382)	1.600** (0.7845)
ln_age	-0.4340*** (0.1024)	-0.9746*** (0.2585)	-0.9548** (0.4180)
lseg_esg		-0.0045*** (0.0010)	-0.0069*** (0.0012)
sp_esg_total			0.0030** (0.0014)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Within R ²	0.48656	0.50925	0.51848
R ²	0.50677	0.53668	0.55428
Observations	6,936	5,797	4,784
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. Robustness counterpart to the PPML specification in Table 13. The dependent variable is $\log(1 + \text{rep_incidents}_{i,t+1})$ rather than the raw count.

Green patents: alternative classification, replications, and decompositions

Table 38: PPML baseline on next-year OECD ENV-TECH patents

Dependent Variable: Model:	patents_oecd_lead1 (1)	(2)
<i>Variables</i>		
lsi_mean	11.05*** (3.317)	1.729 (2.101)
ln_at	1.067*** (0.0556)	0.2711* (0.1441)
leverage	0.2330 (0.6182)	-0.5111 (0.6489)
roa	-1.480** (0.7311)	-0.6928 (0.4953)
rd_int	6.303*** (1.179)	0.3206 (1.310)
capx_int	3.807 (2.469)	-1.992 (2.286)
ln_age	0.2091 (0.4388)	3.529*** (0.6697)
<i>Fixed-effects</i>		
naics2	Yes	
year	Yes	Yes
company_id		Yes
<i>Fit statistics</i>		
Squared Correlation	0.68783	0.93424
Pseudo R ²	0.72805	0.91944
BIC	76,703.4	21,691.9
Observations	9,068	3,766
<i>Clustered (company_id) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Notes. Dependent variable is the count of patents filed in year $t+1$ that match the OECD ENV-TECH 2023 environmental-technology classification described in Section 4.3. The OECD list is substantially broader than the CPC Y02+Y04S list: it adds environmental management (pollution control, waste), water-related adaptation, and biodiversity classes on top of the climate-change-mitigation overlap with Y02.

Table 39: Darendeli-Law-Shen (2022) replication with LSI added, log-linear OLS

Dependent Variable: Model:	(1)	(2)	ln_green_lead1 (3)	(4)	(5)
<i>Variables</i>					
ln_emp	-0.0011 (0.0377)	0.0023 (0.0376)	-0.0654 (0.0468)	0.0952** (0.0398)	0.0952** (0.0397)
ln_sale	0.2740*** (0.0438)	0.2697*** (0.0436)	0.3706*** (0.0572)	0.0325 (0.0200)	0.0327 (0.0200)
rd_sale	0.0005*** (0.0002)	0.0005*** (0.0002)	0.0003** (0.0001)	6.4×10^{-6} (3.89×10^{-5})	9.94×10^{-6} (3.98×10^{-5})
capx_sale	0.0026 (0.0018)	0.0023 (0.0017)	0.0193* (0.0107)	-0.0003 (0.0007)	-0.0003 (0.0007)
ad_sale	-0.6828 (0.5211)	-0.5389 (0.5252)	-0.2332 (0.7727)	0.1280 (0.4265)	0.1348 (0.4270)
leverage	-0.3076*** (0.1167)	-0.3204*** (0.1163)	-0.4266*** (0.1565)	0.1880*** (0.0690)	0.1875*** (0.0690)
roa	-1.142*** (0.1607)	-1.137*** (0.1600)	-1.427*** (0.2211)	-0.0285 (0.0683)	-0.0293 (0.0683)
btm	0.0574 (0.0440)	0.0588 (0.0439)	0.0508 (0.0445)	0.0281 (0.0180)	0.0281 (0.0181)
ln_age	-0.1582* (0.0943)	-0.1628* (0.0943)	-0.2758** (0.1380)	0.1827 (0.2829)	0.1738 (0.2838)
tangibility	-0.2773 (0.1977)	-0.2563 (0.1982)	-0.3191 (0.2532)	0.2243 (0.2584)	0.2244 (0.2583)
lsi_mean		2.155*** (0.8091)	3.880*** (1.084)		0.2174 (0.4301)
lseg_esg			0.0018 (0.0021)		
<i>Fixed-effects</i>					
naics2	Yes	Yes	Yes		
year	Yes	Yes	Yes	Yes	Yes
company_id				Yes	Yes
<i>Fit statistics</i>					
Within R ²	0.22090	0.22400	0.22836	0.01076	0.01083
R ²	0.28876	0.29159	0.33310	0.88898	0.88899
Observations	6,205	6,205	4,490	5,854	5,854
<i>Clustered (company_id) standard-errors in parentheses</i>					
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>					

Notes. I replicate the headline log-linear specification of Darendeli et al. (2022) on the LSI panel. The dependent variable is $\ln(1 + \text{green patents}_{i,t+1})$ on the CPC Y02+Y04S flag; controls match Darendeli et al. exactly: $\ln(\text{EMP})$, $\ln(\text{SALE})$, R&D-to-sales, capex-to-sales, advertising-to-sales, leverage, ROA, book-to-market, $\ln(1 + \text{age})$, and tangibility.

Table 40: Darendeli-Law-Shen (2022) replication recast as PPML on raw counts

Dependent Variable:	patents_green_lead1		
Model:	(1)	(2)	(3)
<i>Variables</i>			
ln_emp	-0.2467 (0.1753)	-0.2535 (0.1763)	-0.1875 (0.2002)
ln_sale	1.279*** (0.1434)	1.267*** (0.1505)	1.248*** (0.1646)
rd_sale	0.0002 (0.0024)	0.0003 (0.0024)	-0.0111*** (0.0028)
capx_sale	0.0297* (0.0158)	0.0290* (0.0158)	0.3085*** (0.0597)
ad_sale	1.440 (4.593)	1.401 (4.621)	2.083 (5.103)
leverage	0.4228 (0.5658)	0.3279 (0.5423)	0.4394 (0.5280)
roa	-2.363** (1.204)	-2.212* (1.187)	-1.353 (1.596)
btm	0.0376 (0.0357)	0.0431 (0.0358)	0.0405 (0.0332)
ln_age	-0.7447 (0.6059)	-0.7525 (0.6063)	-0.8205 (0.7021)
tangibility	-1.133 (1.059)	-1.114 (1.062)	-1.200 (1.130)
lsi_mean		5.023 (4.203)	4.873 (4.489)
lseg_esg			-0.0014 (0.0057)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Squared Correlation	0.66795	0.66605	0.67476
Pseudo R ²	0.71538	0.71696	0.72729
BIC	60,173.8	59,850.5	51,742.5
Observations	6,147	6,147	4,434
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. Same right-hand-side specification as Table 39 columns 1-3 (Darendeli controls, NAICS-2-by-year fixed effects), but with the raw count of next-year Y02+Y04S patents as the dependent variable and PPML as the estimator following [Silva and Tenreyro \(2006\)](#). The point estimate on LSI ($\beta = 5.02$ in column 2, $\beta = 4.87$ in column 3) is in the same direction as the log-linear specification but the standard errors are inflated relative to OLS so the loading is not statistically distinguishable from zero.

Table 41: PPML horse race of LSI against commercial ESG ratings on next-year green patents

Dependent Variable:	patents_green_lead1		
Model:	(1)	(2)	(3)
<i>Variables</i>			
lsi_mean	5.917 (4.240)	5.387 (4.161)	5.870 (4.295)
lseg_esg	0.0007 (0.0039)		0.0048 (0.0046)
ln_at	1.120*** (0.0511)	1.165*** (0.0737)	1.199*** (0.0706)
leverage	-0.1268 (0.5286)	-0.1656 (0.5080)	-0.1256 (0.5022)
roa	-0.2958 (0.9242)	-0.7207 (1.075)	-0.9808 (1.130)
rd_int	10.75*** (1.357)	11.35*** (1.365)	11.82*** (1.372)
capx_int	1.830 (2.122)	2.515 (1.878)	2.760 (1.856)
ln_age	0.4064 (0.8648)	1.077 (1.334)	1.013 (1.395)
sp_esg_total		-0.0061 (0.0056)	-0.0090 (0.0060)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Squared Correlation	0.73214	0.73661	0.74234
Pseudo R ²	0.77670	0.77787	0.78246
BIC	42,615.9	38,842.9	37,418.2
Observations	4,499	3,468	3,313
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. Same outcome and estimator as Table 14 (PPML on green patents_{*i,t+1*} on the CPC Y02+Y04S flag) with the simpler baseline control set (ln(AT), leverage, ROA, R&D intensity, capex intensity, ln(1 + age)) and NAICS-2-by-year fixed effects. The horse race adds the LSEG combined ESG score (column 1), the S&P Global total ESG score (column 2), or both (column 3) as competing predictors against LSI_{*i,t*}.

Table 42: LSI by R&D intensity interaction on next-year patents, green vs non-green

Dependent Variables:	patents_green_lead1	patents_nongreen_lead1
Model:	(1)	(2)
<i>Variables</i>		
lsi_mean	8.867** (4.006)	-0.7949 (3.706)
rd_int	41.90*** (10.20)	19.22*** (7.121)
lsi_x_rd	-56.93*** (18.24)	-17.88 (12.64)
<i>Fixed-effects</i>		
naics2	Yes	Yes
year	Yes	Yes
<i>Fit statistics</i>		
Squared Correlation	0.73156	0.60775
Pseudo R ²	0.77165	0.73242
BIC	48,649.7	541,906.4
Observations	6,272	6,331
<i>Clustered (company_id) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Notes. PPML on next-year patent counts with NAICS-2-by-year fixed effects. Column 1 uses the green-patent count (CPC Y02+Y04S); column 2 uses the placebo non-green-patent count.

Table 43: (per-pillar robustness): Per-pillar LSI by R&D interactions on next-year patents

Dependent Variables:	patents_green_lead1	patents_nongreen_lead1
Model:	(1)	(2)
<i>Variables</i>		
E_mean	-6.478 (8.934)	-19.02*** (5.390)
S_mean	2.567 (8.710)	8.781 (8.163)
G_mean	-9.673 (6.682)	8.303 (6.444)
Sec_mean	23.68*** (6.642)	0.9423 (6.551)
rd_int	33.91*** (8.348)	16.97** (6.614)
E_x_rd	76.36* (45.39)	64.41** (30.68)
S_x_rd	-28.34 (46.43)	-35.21 (36.06)
G_x_rd	-66.39* (36.25)	-70.80** (28.57)
Sec_x_rd	-23.73 (35.25)	30.11 (25.33)
<i>Fixed-effects</i>		
naics2	Yes	Yes
year	Yes	Yes
<i>Fit statistics</i>		
Squared Correlation	0.72301	0.62326
Pseudo R ²	0.78188	0.74225
BIC	46,535.1	522,062.2
Observations	6,272	6,331

Clustered (company_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. PPML on next-year patent counts with NAICS-2-by-year fixed effects, parallel to Table 42 but with the LSI pillar means $E_{i,t}$, $S_{i,t}$, $G_{i,t}$, $Sec_{i,t}$ replacing the overall LSI, and each pillar interacted with R&D intensity. The per-pillar pattern reinforces the headline non-result for sustainability-specific R&D complementarity.

Table 44: Green-versus-non-green placebo on LSI loadings

Dependent Variables:	patents_green_lead1	patents_nongreen_lead1
Model:	(1)	(2)
<i>Variables</i>		
lsi_mean	5.546 (4.017)	-1.858 (3.438)
ln_at	1.102*** (0.0461)	0.8835*** (0.0348)
leverage	-0.1038 (0.4866)	-1.109* (0.6117)
roa	-0.1643 (0.7482)	1.967*** (0.6238)
rd_int	10.28*** (1.209)	9.302*** (0.8475)
capx_int	2.133 (2.019)	-0.6496 (1.922)
ln_age	0.4519 (0.7273)	0.0225 (0.7030)
<i>Fixed-effects</i>		
naics2	Yes	Yes
year	Yes	Yes
<i>Fit statistics</i>		
Squared Correlation	0.73179	0.60735
Pseudo R ²	0.77036	0.73216
BIC	48,913.1	542,414.2
Observations	6,272	6,331
<i>Clustered (company_id) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Notes. PPML on next-year patent counts with NAICS-2-by-year fixed effects and the standard baseline controls. Column 1 has the CPC Y02+Y04S green-patent count as the dependent variable; column 2 has the residual non-green-patent count.

Table 45: Pillar decomposition of LSI on next-year patents, green vs non-green

Dependent Variables:	patents_green_lead1	patents_nongreen_lead1
Model:	(1)	(2)
<i>Variables</i>		
E_mean	-2.402 (7.528)	-15.94*** (4.553)
S_mean	-0.2082 (7.817)	5.482 (6.577)
G_mean	-13.24** (6.353)	4.113 (5.548)
Sec_mean	23.52*** (6.208)	4.401 (6.056)
ln_at	1.132*** (0.0509)	0.8431*** (0.0338)
leverage	-0.4226 (0.5878)	-0.9750* (0.4981)
roa	-0.2462 (0.6412)	1.800*** (0.5302)
rd_int	10.83*** (1.053)	8.471*** (0.8245)
capx_int	1.279 (2.056)	-0.1640 (2.080)
ln_age	0.4123 (0.6526)	-0.0205 (0.6877)
<i>Fixed-effects</i>		
naics2	Yes	Yes
year	Yes	Yes
<i>Fit statistics</i>		
Squared Correlation	0.71808	0.61591
Pseudo R ²	0.77992	0.73896
BIC	46,915.1	528,681.8
Observations	6,272	6,331
<i>Clustered (company_id) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Notes. PPML on next-year patent counts with NAICS-2-by-year fixed effects and the standard baseline controls. The four LSI pillar means $E_{i,t}$, $S_{i,t}$, $G_{i,t}$, $Sec_{i,t}$ (Section 3.6) replace the overall $LSI_{i,t}$.

Future profitability: sales, labour productivity, and capex interaction

Table 46: LSI on future $\ln(\text{sales})$ at horizons $k \in \{1, 3, 5\}$

Dependent Variables: Model:	$\ln_sale_lead1$ (1)	$\ln_sale_lead3$ (2)	$\ln_sale_lead5$ (3)
<i>Variables</i>			
lsi_mean	0.0517 (0.0744)	-0.1121 (0.1922)	-0.0917 (0.3333)
$\ln_sale$	0.9058*** (0.0109)	0.8185*** (0.0214)	0.8054*** (0.0326)
$\ln_at$	0.0800*** (0.0099)	0.1398*** (0.0194)	0.1314*** (0.0293)
leverage	-0.0395*** (0.0152)	-0.0846** (0.0377)	-0.1117* (0.0618)
roa	0.0285 (0.0431)	-0.1105 (0.0940)	-0.2866* (0.1697)
rd_int	0.1076 (0.0815)	0.7936*** (0.2053)	1.558*** (0.3057)
capx_int	0.2011** (0.0933)	0.5445** (0.2508)	0.9202** (0.4136)
$\ln_age$	-0.1610*** (0.0163)	-0.3518*** (0.0585)	-0.6616*** (0.1590)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Within R^2	0.97338	0.94074	0.90868
R^2	0.97509	0.94446	0.91435
Observations	18,342	9,751	4,188
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. Dependent variable is firm-year $\ln(\text{sales})$ k years ahead, with $k \in \{1, 3, 5\}$ across columns. The right-hand-side variable of interest is the contemporaneous headline aggregator $LSI_{i,t}$.

Table 47: LSI on future $\ln(\text{labour productivity})$ at horizons $k \in \{1, 3, 5\}$

Dependent Variables: Model:	$\ln_labour_prod_lead1$ (1)	$\ln_labour_prod_lead3$ (2)	$\ln_labour_prod_lead5$ (3)
<i>Variables</i>			
lsi_mean	0.0192 (0.0536)	-0.0244 (0.1219)	-0.2127 (0.2062)
$\ln_labour_prod$	0.9447*** (0.0040)	0.9061*** (0.0082)	0.8826*** (0.0116)
$\ln_at$	0.0080*** (0.0010)	0.0110*** (0.0023)	0.0158*** (0.0037)
leverage	0.0239** (0.0104)	0.0290 (0.0233)	0.0214 (0.0357)
roa	-0.2773*** (0.0298)	-0.6044*** (0.0613)	-0.6818*** (0.0948)
rd_int	-0.0066 (0.0531)	0.1797 (0.1225)	0.6633*** (0.1805)
capx_int	-0.1788*** (0.0599)	-0.2248 (0.1397)	-0.2067 (0.2250)
$\ln_age$	-0.0353* (0.0184)	-0.0451 (0.0500)	-0.1133 (0.2000)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Within R ²	0.89933	0.84439	0.81212
R ²	0.92334	0.88562	0.86673
Observations	17,741	9,462	4,075

Clustered (company_id) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Dependent variable is firm-year $\ln(\text{sales/employment})$ k years ahead, with $k \in \{1, 3, 5\}$ across columns. The right-hand-side variable of interest is the contemporaneous headline aggregator $LSI_{i,t}$.

Table 48: LSI $\times$ capex-intensity interaction on ROA($t + 3$)

Dependent Variable:	roa_lead3
Model:	(1)
<i>Variables</i>	
lsi_mean	0.0106 (0.0429)
capx_int	0.8683** (0.3999)
lsi_x_capx	-1.488** (0.7558)
ln_at	0.0053*** (0.0007)
leverage	0.0196** (0.0087)
roa	0.4404*** (0.0248)
rd_int	-0.1353*** (0.0448)
ln_age	0.0118 (0.0112)
<i>Fixed-effects</i>	
naics2	Yes
year	Yes
<i>Fit statistics</i>	
Within R ²	0.30582
R ²	0.34080
Observations	9,790

Clustered (company_id) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Dependent variable is firm-year ROA three years ahead, ROA_{*i,t+3*}. The right-hand-side variables of interest are the contemporaneous headline aggregator LSI_{*i,t*}, capex intensity (capex divided by total assets), and their product LSI_{*i,t*} $\times$ capx_int_{*i,t*}.

Intangible-capital integration: robustness

Table 49: Investment-q regression with wage-weighted dollar LSI-K stock

Dependent Variables: Model:	IK_PT (1)	IK_PT_ext (2)	IK_PT (3)
<i>Variables</i>			
q_tot	-0.0001 (0.0007)	0.0009 (0.0014)	-0.0340** (0.0148)
lsik_share			47.35*** (17.97)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Within R ²	1.75×10^{-8}	1.13×10^{-7}	0.04537
R ²	0.01444	0.01187	0.05918
Observations	19,993	19,959	19,737
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. Dependent variables: column (1) is $IK_{i,t}^{PT} = CAPX_{i,t}/k_{int_{i,t}}$, the Peters-Taylor investment-rate (CAPX scaled by total Peters-Taylor intangible capital, $k_{int} = k_{int_know} + k_{int_org}$); column (2) replaces the denominator with $k_{int_{i,t}} + LSI-K_{i,t}$, augmenting the intangible-capital base with the wage-weighted LSI-K stock; column (3) uses the baseline denominator from column (1) and adds the LSI-K share covariate $lsik_share_{i,t} = LSI-K_{i,t}/(k_{int_{i,t}} + LSI-K_{i,t})$. The right-hand-side variable of interest is the Peters-Taylor augmented Tobin's q (q_tot); see Peters and Taylor (2017).

Table 50: robustness: Investment-q regression with unweighted LSI-K stock

Dependent Variable: Model:	(1)	IK_PT (2)
<i>Variables</i>		
q_tot	-0.0001 (0.0007)	-0.0002 (0.0007)
lsi_k_unw_std		-0.6483 (0.4501)
<i>Fixed-effects</i>		
naics2	Yes	Yes
year	Yes	Yes
<i>Fit statistics</i>		
Within R ²	1.75×10^{-8}	0.00050
R ²	0.01444	0.01493
Observations	19,993	19,993
<i>Clustered (company_id) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Notes. Robustness counterpart to Table 49. Dependent variable in both columns is the Peters-Taylor investment-rate $IK_{i,t}^{PT} = CAPX_{i,t}/k_{int_{i,t}}$.

Table 51: (orthogonalised): investment-q with LSI-K residual to organisation capital

Dependent Variable:	IK_PT
Model:	(1)
<i>Variables</i>	
q_tot	-0.0001 (0.0007)
lsi_k_orth_std	0.0362 (0.0820)
<i>Fixed-effects</i>	
naics2	Yes
year	Yes
<i>Fit statistics</i>	
Within R ²	1.83×10^{-6}
R ²	0.01446
Observations	19,737
<i>Clustered (company_id) standard-errors in parentheses</i>	
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>	

Notes. The covariate `lsi_k_orth_std` is the standardised residual from a first-stage NAICS-2-by-year fixed-effect regression of LSI-K_{*i,t*} on Peters-Taylor organisation capital (`k_int_org`), $\ln(AT)$, and leverage. The residual is the component of the LSI-K level that is orthogonal to organisation capital and to firm size, and is in principle the version of LSI-K that should isolate the sustainability-specific contribution beyond general organisation capital.

Table 52: robustness: firm characteristics by LSI-K / `k_int_org` quintile

LSI-K / <code>k_int_org</code> quintile	<i>N</i>	Mean ratio	Mean ROA	Mean Tobin's q	Mean <code>q_tot</code>	Mean R&D intensity	Mean CAPX intensity
Q1 (lowest)	4,303	0.008	+0.031	1.64	1.38	0.022	0.036
Q2	4,300	0.040	+0.006	1.59	1.42	0.033	0.035
Q3	4,298	0.083	−0.004	1.74	1.62	0.043	0.031
Q4	4,300	0.161	−0.027	1.98	2.13	0.054	0.031
Q5 (highest)	4,301	–	−0.043	4.97	6.29	0.053	0.037

Notes. Robustness counterpart to Table 19 using the within-organisation-capital ratio $LSI-K_{i,t}/k_int_org_{i,t}$ as the sorting variable. The ratio is computed only on firm-years with strictly positive `k_int_org` ($N = 21,502$ firm-years across the five quintiles).

Table 53: Appendix A.8: LSI-K depreciation-rate sensitivity for the headline G2 PT decomposition

Dependent Variable: Model:	IK_PT		
	(1)	(2)	(3)
<i>Variables</i>			
q_tot	-0.0300** (0.0129)	-0.0335** (0.0144)	-0.0391** (0.0168)
I(k_int_org/at)	-1.376*** (0.5309)	-1.385*** (0.5320)	-1.376*** (0.5303)
I(k_int_know/at)	0.3416** (0.1639)	0.3598** (0.1711)	0.3763** (0.1780)
lsik_share_d15	42.64*** (16.11)		
lsik_share_d20		47.80*** (18.12)	
lsik_share_d30			56.20*** (21.41)
<i>Fixed-effects</i>			
naics2	Yes	Yes	Yes
year	Yes	Yes	Yes
<i>Fit statistics</i>			
Within R ²	0.04131	0.04599	0.05333
R ²	0.05517	0.05978	0.06702
Observations	19,737	19,737	19,737
<i>Clustered (company_id) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Notes. Sensitivity of the headline G2 PT-decomposition regression (Table 18 in Section 5.7) to the perpetual-inventory depreciation rate δ used to construct LSI-K. The dependent variable in every column is I/k_{int} , the investment-to-augmented-capital ratio from Peters and Taylor (2017).

Validation battery and posting language distribution

Table 54: Posting language distribution (langdetect on a stratified 80,000-posting sample)

Language (ISO code)	Share
English (en)	97.56%
German (de)	0.75%
French (fr)	0.58%
Spanish (es)	0.40%
Chinese, simplified (zh-cn)	0.34%
Portuguese (pt)	0.16%
Japanese (ja)	0.08%
Polish (pl)	0.08%
Dutch (nl)	0.08%
Other / unknown	0.14%
Total non-English	2.44%

Notes. Language detection performed with **langdetect** (Shuyo's port of Google's language-detection library) on the first 1,000 characters of each posting's cleaned description text. Sample stratified by year (10,000 postings per calendar year for 2013-2020).

Table 55: T1.1: AUC of LSI scores against the O*NET binary green-occupation flag, by pillar

LSI score	Curated (primary)			Curated minus alt-corpus	
	AUC	95% CI (lower)	95% CI (upper)	vs alt-A	vs alt-B
total_max	0.579	[0.576	0.582]	−0.009	−0.014
E_score	0.626	[0.623	0.629]	+0.024	+0.019
S_score	0.568	[0.565	0.572]	−0.011	−0.005
G_score	0.552	[0.548	0.555]	−0.001	−0.009
Sec_score	0.614	[0.610	0.617]	+0.016	+0.013

Notes. Computed on a stratified 250,000-posting validation sample (balanced across calendar years 2013-2023). The curated 95% confidence interval is from a 1,000-replication bootstrap.

Table 56: T1.1.b: Walk-versus-Talk discrimination gap within the 204 O*NET green occupations

LSI score	Curated	Alt-A	Alt-B	Curated / best alt
total_max	+0.0212	+0.0099	+0.0081	2.1×
E_score	+0.0349	+0.0145	+0.0107	2.4×
S_score	+0.0212	+0.0102	+0.0080	2.1×
G_score	−0.0041	−0.0063	−0.0070	—
Sec_score	+0.0225	+0.0114	+0.0037	2.0×

Notes. Walk-versus-Talk classification of the 204 O*NET green occupations from [Chen \(2024\)](#) (186 Walk, 18 Talk). Reported values are the mean LSI gap between Walk and Talk postings within each pillar, computed on the validation sample.

Table 57: T1.4: Walk-versus-Talk AUC within green SOC, text-level LSI vs occupation-level imputation

Measure	Text-level AUC	Occupation-level AUC	Text minus occupation
Curated LSI total_max	0.611	0.496	+0.115
Alt-A LSI total_max	0.560	0.496	+0.064
Alt-B LSI total_max	0.555	0.496	+0.059

Notes. The occupation-level benchmark is the Vona-style **green_task_count** from [Vona et al. \(2018\)](#), imputed to each posting by its SOC code. All three corpora beat occupation-level imputation by 0.06 to 0.12 AUC, demonstrating that text-level scoring delivers within-occupation discrimination that pure occupation-level measures cannot.

Table 58: T1.3: Head-to-head AUC of LSI scores vs published keyword dictionaries

Corpus	Wins (of 35 pairings)	Notes
Curated	34 / 35	Sole tie: G_score vs Chen top-120 governance
Alt-A	34 / 35	Same governance-domain tie
Alt-B	35 / 35	Sweeps via max-over-chunks coverage

Notes. Pairings combine three keyword dictionaries (Sautner et al. 2023; Baier-Berninger-Kiesel 2020; Chen 2024 top-120) with the five LSI scores and the seven dictionary subsets, evaluated by paired-bootstrap AUC difference against the O*NET binary green flag. All three corpus-based measures beat published keyword dictionaries decisively on the 35 pairings; no dictionary statistically dominates any sustainability-corpus measure anywhere.

Table 59: T6.1: Aggregator robustness (Spearman correlations across firm-year aggregator choices)

Corpus	Min pairwise Spearman across 6 aggregators	Cluster-axis vs pillar-axis Spearman
Curated	0.963	0.986
Alt-A	0.991	0.997
Alt-B	0.975	0.990

Notes. Six aggregators tested: **total_max**, **total_pillar_mean**, **total_top3_mean**, **total_top5_mean**, **total_top10_mean**, and the cluster-axis equivalent of **total_pillar_mean**. All pairwise Spearman correlations are well above the 0.85 robustness threshold established in the validation protocol.

T3.2 truncation sensitivity. On the validation sample, fewer than 0.01% of postings exceed the 4,096-token truncation cap I apply during embedding (14 postings out of 250,000). The empirical 99th percentile of posting

Table 60: T5: Construct-internal geometry of the topic-anchor space

Metric	Curated	Alt-A
Within-pillar mean cosine	0.71	0.94
Between-pillar mean cosine	0.68	0.91
Within minus between	+0.029	+0.032
Within / between ratio	1.04	1.04

Notes. The curated corpus’s anchors are conceptually distinct (within-pillar cosine of 0.71), so cosine similarity to the curated anchors captures fine-grained topic discrimination. The alt-corpus anchors are nearly identical to each other (within-pillar cosine of 0.94) because they are mean-pooled from largely overlapping source PDFs, so cosine similarity to the alt-corpus anchors captures a coarse green-vocabulary signal.

length on the broader 80,000-posting language-and-tokens sample is 1,964 tokens, well below the cap. Truncation loss is therefore negligible for the empirical results.

T0 face validity. The top-decile postings on `total_max` carry job titles such as "Chief Sustainability Officer", "Director of ESG Reporting", "Senior Climate Strategist", "EHS Engineering Manager", and "Sustainable Finance Director". The bottom-decile postings are routine retail, sales, and operational roles with no sustainability content. The face-validity check is a necessary but not sufficient validation; the T1.x external benchmarks above carry the substantive load.

References

- Simona Abis and Laura Veldkamp. The changing economics of knowledge production. *Review of Financial Studies*, 37(1):89–118, 2024. doi: 10.1093/rfs/hhad059.
- Daron Acemoglu, David Autor, Jonathon Hazell, and Pascual Restrepo. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340, 2022.
- Philippe Aghion, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen. Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51, 2016.
- Dogu Araci. FinBERT: Financial sentiment analysis with pre-trained language models. arXiv:1908.10063, 2019.
- Doron Avramov, Si Cheng, Abraham Lioui, and Andrea Tarelli. Sustainable investing with esg rating uncertainty. *Journal of Financial Economics*, 145(2):642–664, 2022.
- Tania Babina, Anastassia Fedyk, Alex Xi He, and James Hodson. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745, 2024.
- Florian Berg, Julian F. Kölbel, and Roberto Rigobon. Aggregate confusion: The divergence of esg ratings. *Review of Finance*, 26(6):1315–1344, 2022.
- Julia Anna Bingler, Mathias Kraus, Markus Leippold, and Nicolas Webersinke. Cheap talk and cherry-picking: What ClimateBert has to say on corporate climate risk disclosures. *Finance Research Letters*, 47:102776, 2022. doi: 10.1016/j.frl.2022.102776.
- Julia Anna Bingler, Mathias Kraus, Markus Leippold, and Nicolas Webersinke. How cheap talk in climate disclosures relates to climate initiatives, corporate emissions, and reputation risk. *Journal of Banking and Finance*, 164, 2024.
- Patrick Bolton and Marcin Kacperczyk. Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549, 2021.
- Brantly Callaway and Pedro H. C. Sant’Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021.
- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. M3-Embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. arXiv:2402.03216, 2024.
- Shuang Chen. Green investors and green transition efforts: Talk the talk or walk the walk? Working paper, University of Melbourne; SSRN 4254894; IWFSAS 2024 version, 2024.
- Dane M. Christensen, George Serafeim, and Anywhere Sikochi. Why is corporate virtue in the eye of the beholder? the case of esg ratings. *The Accounting Review*, 97(1):147–175, 2022.

- Lauren Cohen, Umit G. Gurun, and Quoc H. Nguyen. The ESG-innovation disconnect: Evidence from green patenting. *Management Science*, 2026. Forthcoming; published online ahead of print 2024. DOI 10.1287/mnsc.2024.09124.
- Davide Consoli, Giovanni Marin, Alberto Marzucchi, and Francesco Vona. Do green jobs differ from non-green jobs in terms of skills and human capital? *Research Policy*, 45(5):1046–1060, 2016.
- Carol A. Corrado, Charles R. Hulten, and Daniel E. Sichel. Measuring capital and technology: An expanded framework. In Carol Corrado, John Haltiwanger, and Daniel Sichel, editors, *Measuring Capital in the New Economy*, volume 65 of *NBER Studies in Income and Wealth*, pages 11–46. University of Chicago Press, 2005.
- Carol A. Corrado, Charles R. Hulten, and Daniel E. Sichel. Intangible capital and u.s. economic growth. *Review of Income and Wealth*, 55(3):661–685, 2009.
- Nicolas Crouzet, Janice C. Eberly, Andrea L. Eisfeldt, and Dimitris Papanikolaou. The economics of intangible capital. *Journal of Economic Perspectives*, 36(3):29–52, 2022.
- Alper Darendeli, Kelvin K. F. Law, and Michael Shen. Green new hiring. *Review of Accounting Studies*, 27(3):986–1037, 2022.
- David J. Deming and Lisa B. Kahn. Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1):S337–S369, 2018.
- Erich C. Dierdorff, Jennifer J. Norton, Donald W. Drewes, Christina M. Kroustalis, David Rivkin, and Phil Lewis. Greening of the world of work: Implications for O*NET-SOC and new and emerging occupations. Technical report, National Center for O*NET Development, 2009.
- Alex Edmans. Does the stock market fully value intangibles? employee satisfaction and equity prices. *Journal of Financial Economics*, 101(3):621–640, 2011.
- Andrea L. Eisfeldt and Dimitris Papanikolaou. Organization capital and the cross-section of expected returns. *Journal of Finance*, 68(4):1365–1406, 2013.
- Andrea L. Eisfeldt, Edward Kim, and Dimitris Papanikolaou. Intangible value. *Critical Finance Review*, 11(2):299–332, 2022.
- Robert F. Engle, Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel. Hedging climate change news. *Review of Financial Studies*, 33(3):1184–1216, 2020.
- Matthew Gentzkow, Bryan T. Kelly, and Matt Taddy. Text as data. *Journal of Economic Literature*, 57(3):535–574, 2019.
- Simon Glossner. Repeat offenders: ESG incident recidivism and investor underreaction. Working paper, 2021.
- Bronwyn H. Hall. The manufacturing sector master file: 1959–1987. NBER Working Paper 3366, National Bureau of Economic Research, 1990.

- Samuel M. Hartzmark and Abigail B. Sussman. Do investors value sustainability? a natural experiment examining ranking and fund flows. *Journal of Finance*, 74(6):2789–2837, 2019.
- Tarek A. Hassan, Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics*, 134(4):2135–2202, 2019.
- Brad Hershbein and Lisa B. Kahn. Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–1772, 2018.
- Gerard Hoberg and Gordon Phillips. Product market synergies and competition in mergers and acquisitions: A text-based analysis. *Review of Financial Studies*, 23(10):3773–3811, 2010.
- Gerard Hoberg and Gordon Phillips. Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465, 2016.
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, 2020.
- Mozaffar Khan, George Serafeim, and Aaron Yoon. Corporate sustainability: First evidence on materiality. *The Accounting Review*, 91(6):1697–1724, 2016.
- Leonid Kogan, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. Technological innovation, resource allocation, and growth. *Quarterly Journal of Economics*, 132(2):665–712, 2017.
- Philipp Krueger, Zacharias Sautner, and Laura T. Starks. The importance of climate risks for institutional investors. *Review of Financial Studies*, 33(3):1067–1111, 2020.
- Philipp Krüger. Corporate goodness and shareholder wealth. *Journal of Financial Economics*, 115(2):304–329, 2015.
- Tim Loughran and Bill McDonald. When is a liability not a liability? textual analysis, dictionaries, and 10-ks. *Journal of Finance*, 66(1):35–65, 2011.
- Thomas P. Lyon and John W. Maxwell. Greenwash: Corporate environmental disclosure under threat of audit. *Journal of Economics and Management Strategy*, 20(1):3–41, 2011.
- Asaf Manela and Alan Moreira. News implied volatility and disaster concerns. *Journal of Financial Economics*, 123(1):137–162, 2017.
- Christopher Marquis, Michael W. Toffel, and Yanhua Zhou. Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, 27(2):483–504, 2016.
- Nora M. C. Pankratz, Rob Bauer, and Jeroen Derwall. Climate change, firm performance, and investor surprises. *Management Science*, 69(12):7352–7398, 2023.
- Luboš Pástor, Robert F. Stambaugh, and Lucian A. Taylor. Sustainable investing in equilibrium. *Journal of Financial Economics*, 142(2):550–571, 2021.

- Luboš Pástor, Robert F. Stambaugh, and Lucian A. Taylor. Dissecting green returns. *Journal of Financial Economics*, 146(2):403–424, 2022.
- Luboš Pástor, Robert F. Stambaugh, and Lucian A. Taylor. Sustainable investing. *Annual Review of Financial Economics*, 17:321–341, 2025.
- Ryan H. Peters and Lucian A. Taylor. Intangible capital and the investment-q relation. *Journal of Financial Economics*, 123(2):251–272, 2017.
- Aneesh Raghunandan and Shivaram Rajgopal. Do socially responsible firms walk the talk? *Journal of Law and Economics*, 67(4), 2024.
- Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, 2019.
- Aurelien Saussay, Misato Sato, Francesco Vona, and Layla O’Kane. Who’s fit for the low-carbon transition? emerging skills and wage gaps in job ad data. FEEM Working Paper No. 31 / Grantham Research Institute working paper, 2022.
- Zacharias Sautner, Laurence van Lent, Grigory Vilkov, and Ruishen Zhang. Firm-level climate change exposure. *Journal of Finance*, 78(3):1449–1498, 2023.
- Tobias Schimanski, Andrin Reding, Nico Reding, Julia Bingler, Mathias Kraus, and Markus Leippold. Bridging the gap in ESG measurement: Using NLP to quantify environmental, social, and governance communication. *Finance Research Letters*, 61:104979, 2024. doi: 10.1016/j.frl.2024.104979.
- Daniel E. Sichel. Intangibles: How far we have come and the road ahead. *Review of Income and Wealth*, 71(1), 2025. DOI 10.1111/roiw.12706.
- J. M. C. Santos Silva and Silvana Tenreyro. The log of gravity. *Review of Economics and Statistics*, 88(4):641–658, 2006.
- Liyang Sun and Sarah Abraham. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199, 2021.
- Francesco Vona, Giovanni Marin, Davide Consoli, and David Popp. Environmental regulation and green skills: An empirical exploration. *Journal of the Association of Environmental and Resource Economists*, 5(4):713–753, 2018.
- Nicolas Webersinke, Mathias Kraus, Julia Anna Bingler, and Markus Leippold. ClimateBERT: A pretrained language model for climate-related text. In *Proceedings of the AAAI 2022 Fall Symposium: The Role of AI in Responding to Climate Challenges*, 2022.