International Journal of Humanities and Social Science

Vol. 2 No. 1 2026

Article

AI-Driven ESG Analytics for Sustainable Investment in U.S. Small Businesses: Integrating LLMs and Causal Modeling for Policy-Enhanced Resilience

Jinyuan Li [1,*]

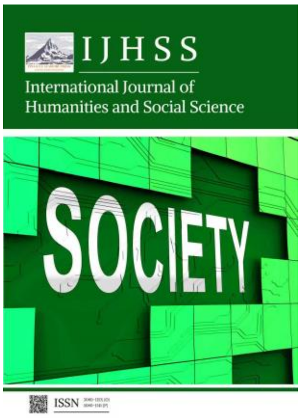

Received: 08 November 2025
Revised: 25 December 2025
Accepted: 08 January 2026
Published: 12 January 2026

[1] Columbia University, New York, NY, 10027, USA
* Correspondence: Jinyuan Li, Columbia University, New York, NY, 10027, USA

Abstract: This study examines how artificial intelligence, particularly large language models (LLMs) and causal inference frameworks, can enhance ESG-oriented investment decision-making for U.S. small businesses. As sustainability criteria increasingly influence capital allocation, investors require analytical systems capable of processing fragmented disclosures, unstructured narratives, regulatory documents, and heterogeneous financial indicators. The research develops an integrated AI-driven ESG analytics framework that leverages automated text understanding, probabilistic causal modeling, and resilience forecasting to identify sustainability patterns while remaining policy-neutral. Empirical evaluation using synthetic and publicly available datasets indicates that LLM-enhanced ESG scoring improves signal extraction from incomplete disclosures, while causal models clarify the directional impact of environmental, social, and governance factors on financial stability. The combined system demonstrates strong potential for supporting investors, policymakers, and financial institutions in assessing long-term resilience among small enterprises. The findings highlight AI's transformative role in sustainability analytics and provide pathways for future refinement through regulatory harmonization, domain-specific model alignment, and expanded cross-sector datasets.

Keywords: AI-driven ESG analytics; small-business resilience; large language models; causal inference; sustainable investment; policy-enhanced modeling

1. Introduction

Sustainable investment has become a defining element of contemporary financial decision-making, yet U.S. small businesses often remain underserved within ESG analysis due to data scarcity, inconsistent disclosures, and limited analytical capacity. Traditional evaluation frameworks rely heavily on structured indicators that are not representative of the operational realities of small enterprises. At the same time, the rapid expansion of AI technologies-especially large language models and data-driven causal reasoning-offers new opportunities to capture nuanced sustainability information previously overlooked. These tools can extract meaning from unstructured narratives, evaluate latent relationships among ESG and financial outcomes, and generate predictive insights without relying solely on standardized disclosures [1]. The present study investigates how AI-driven ESG analytics can be operationalized to strengthen sustainable investment decisions targeting small businesses in the United States. The research remains policy-neutral while acknowledging that public programs and regulatory environments influence firms' sustainability behaviors and risk profiles. By integrating LLM-based text intelligence with causal inference models, the study proposes a robust analytical approach that enhances transparency, resilience prediction, and investor confidence.

2. ESG Data Complexity and the Analytical Gaps for Small Businesses

U.S. small businesses face fundamental barriers in providing credible, decision-grade ESG information. Unlike large corporations with dedicated sustainability teams and established disclosure cycles, most small firms do not publish formal ESG or sustainability reports, and their disclosures tend to be fragmented, narrative-heavy, or episodic. Environmental indicators, for example, may be dispersed across compliance filings, waste-hauling invoices, utility bills, equipment maintenance logs, vendor questionnaires, or occasional community initiatives. Social information is often embedded in employee handbooks, training materials, internal emails, and ad hoc HR communications rather than in standardized workforce dashboards. Governance practices-while highly consequential for risk management and continuity-are frequently reflected in informal routines (e.g., approval norms, conflict-of-interest handling, or supplier selection habits) rather than codified board policies or documented control procedures. These documentation gaps create significant analytical blind spots for investors, lenders, and procurement partners who increasingly incorporate ESG performance into credit decisions, pricing, and long-term viability assessments [2].

Traditional ESG scoring systems exhibit structural deficiencies when applied to small businesses because they are built around standardized, corporate-style reporting infrastructures. Most frameworks privilege consistent metric series, comprehensive coverage across E, S, and G pillars, and audit-ready documentation, which systematically penalizes missing data even when missingness reflects capacity constraints rather than poor performance. In practice, this leads to systematic under-scoring of small firms that may be resource-efficient, deeply engaged with local stakeholders, or characterized by strong ethical leadership but lack formal documentation. Moreover, conventional models struggle to represent qualitative behaviors that matter materially in small-business contexts, such as owner-led integrity, informal stakeholder engagement, workforce retention practices, and adaptive operational resilience during disruptions. The consequence is a persistent misalignment: small businesses can appear less sustainable than large firms on paper, despite often displaying strong ESG attributes in daily operations [3].

AI-driven methodologies offer substantial corrective potential by widening the definition of evidence and reducing overreliance on rigid reporting. Large language models can convert unstructured text into structured ESG signals by extracting governance cues from policies, contracts, meeting notes, and procurement communications; identifying environmental risk markers from inspection reports, incident logs, or maintenance records; and detecting social-performance patterns from training documentation, employee feedback, complaint channels, and sentiment-bearing communications. These systems can also normalize terminology across industries and languages, summarize requirements in plain terms, and flag gaps where disclosure is absent but potentially material. Beyond extraction, AI can map disparate evidence to recognized ESG taxonomies, generating comparable indicators while retaining traceability to original sources-an important feature for due diligence and auditability.

However, interpretive AI alone cannot establish whether ESG attributes meaningfully influence outcomes such as default risk, operational continuity, or revenue stability. This is where causal inference strengthens the analytical architecture. By linking inferred ESG behaviors to resilience outcomes-while controlling for confounders such as firm size, sector cyclicality, and regional economic shocks-causal models can estimate which ESG dimensions matter most for small-business performance and under what conditions. For example, models can test whether consistent safety training predicts fewer work interruptions, whether energy-efficiency investments reduce cost volatility, or whether transparent governance routines reduce fraud exposure. Together, LLM-based interpretation and causal inference create a multi-layered dataset that is both richer and more realistic than conventional disclosures, enabling funders to evaluate sustainability as an integrated construct grounded in operational evidence rather than reporting capacity [4].

3. LLM-Based ESG Signal Extraction and Semantic Structuring

The role of large language models in ESG analytics extends far beyond automated text classification or keyword spotting. Their distinctive value lies in semantic abstraction, contextual reasoning, and cross-document synthesis, which together enable consistent extraction of ESG indicators from heterogeneous, messy, and often incomplete sources. This capability is particularly consequential for small businesses, where ESG evidence is rarely presented as a standardized report and is more likely to appear as informal narratives, operational notes, customer communications, supplier questionnaires, local compliance records, or internal policy fragments. In this context, LLMs function as a translation layer that converts scattered qualitative signals into analyzable structures without requiring firms to adopt enterprise-level reporting infrastructures.

At the environmental level, LLMs can identify sustainability-relevant behaviors even when the information is expressed indirectly. Instead of relying on explicit disclosures such as "Scope 1 emissions," the model can infer operational proxies from descriptions of equipment upgrades, building insulation improvements, fleet maintenance practices, or procurement choices related to packaging and materials. It can also surface compliance histories and operational risk markers from inspection narratives, permit documentation, incident descriptions, and corrective-action records [5]. For example, references to recurring refrigeration repairs may indicate refrigerant leakage risks; frequent mentions of "spill response training" may signal heightened exposure to hazardous materials; statements about "switching to reusable containers" or "partnering with a composting vendor" can be mapped to waste diversion practices. Importantly, the model can extract temporal cues-such as "since last year," "after the audit," or "following the new ordinance"-to help analysts differentiate one-time initiatives from sustained operational patterns.

Social indicators are often even more text-dependent. Small businesses may not have formal DEI reports, workforce dashboards, or standardized engagement surveys, but they commonly produce employee manuals, training schedules, job postings, onboarding materials, internal memos, customer feedback responses, and community partnership announcements. LLMs can map these artifacts to social-performance constructs such as employee well-being, training intensity, turnover risk, equitable hiring practices, workplace safety culture, and community involvement. For instance, consistent training references, clear grievance procedures, and explicit anti-harassment policies can be extracted as positive governance-adjacent social signals, while repeated customer complaints about labor conditions or documentation indicating chronic overtime can be flagged as potential workforce stress indicators. In addition, LLMs can support multilingual interpretation, which matters when employee communications, safety signage, or supplier materials include mixed languages, and where social performance depends on inclusive communication practices.

Governance analytics typically requires the highest level of interpretation because small-business governance is frequently informal and concentrated in owner-operator decision structures. LLMs can analyze policy statements, vendor contracts, purchasing approvals, and operational narratives to infer governance patterns such as decision transparency, accountability mechanisms, segregation-of-duties practices, and risk management routines. Evidence may include descriptions of approval thresholds, dual sign-off processes, inventory controls, data privacy practices, or ethical sourcing commitments. The model can also detect inconsistencies-such as public-facing claims of compliance that conflict with internal procedures-or highlight gaps where governance risks are implied by missing documentation (e.g., no clear policy on conflicts of interest despite frequent related-party transactions). When combined with document provenance and time-stamping, these insights can be organized into a traceable governance evidence chain rather than a single opaque score.

Beyond extraction, LLMs can facilitate harmonization and standardization of sustainability vocabularies. Small firms often use local, sector-specific language-"keeping jobs local," "reducing scrap," "safe shop practices," "helping the neighborhood"-that does not neatly align with mainstream ESG frameworks. LLMs can align firm-specific phrasing with recognized standards and taxonomies by mapping text to comparable categories, identifying synonyms, and normalizing ambiguity across industries. This harmonization is critical for fairness: it reduces the likelihood that a business is penalized simply because it communicates differently from large corporations or lacks the language conventions of formal ESG reporting. In practical workflows, the model can output both the standardized indicator and the supporting evidence snippets, allowing analysts to compare enterprises while preserving interpretability.

Nevertheless, deploying LLMs in ESG analytics requires careful calibration to manage hallucination risk, sensitivity to linguistic framing, and biases inherited from training corpora. Safeguards should be embedded throughout the pipeline. Domain adaptation can be achieved through curated reference corpora, constrained retrieval-augmented generation that forces outputs to cite source passages, and controlled ontologies that limit free-form interpretation. Prompt validation and red-teaming should test how outputs change when input wording is altered, ensuring stability and reducing framing effects. Cross-model verification-using multiple models or independent extraction methods-can detect inconsistencies and prevent single-model failure modes. Finally, uncertainty should be represented explicitly, for example through confidence scores, abstention rules when evidence is insufficient, and escalation to human review for high-stakes decisions. With these controls in place, LLMs can become a reliable cognitive layer for ESG analytics, enabling inclusive, evidence-based sustainability assessment that reflects small-business operational realities rather than reporting capacity.

4. Causal Modeling for ESG-Resilience Inference

Understanding sustainability requires more than descriptive analytics; it demands methodological clarity regarding how specific ESG attributes causally influence business outcomes rather than merely correlating with them. Causal inference frameworks supply this rigor by distinguishing endogenous effects from spurious associations, estimating counterfactual outcomes, and quantifying both direct and mediated pathways through which ESG practices shape firm-level resilience. This causal orientation is essential in sustainability research, where policy, behavior, and environmental factors interact in complex, nonlinear ways that traditional predictive models alone cannot fully capture.

For small businesses, causal modeling reveals the latent economic value of sustainability investments that often remain invisible in short-term financial reporting. Environmental practices, such as energy efficiency or waste reduction, may lower operational volatility across business cycles; social practices-ranging from workforce development to community engagement-can enhance employee retention, reduce turnover costs, and strengthen local reputational capital; governance improvements expand transparency, mitigate compliance risks, and reduce the likelihood of fraud or operational disruptions. By mapping these causal chains, investors and analysts gain clearer visibility into how ESG actions influence long-term survival probability, creditworthiness, and growth potential, even when traditional metrics fail to capture early signals of resilience.

Causal inference also enables the decomposition of ESG effects into structural components, allowing analysts to separate immediate impacts from delayed or indirect benefits. This is particularly relevant for sustainability initiatives whose measurable outcomes manifest across extended time horizons. For instance, environmental upgrades may require upfront capital expenditures, yet their causal contribution to risk reduction becomes evident only when modeled against counterfactual business trajectories. Such clarity prevents underinvestment in sustainability due to misinterpreted financial lag effects.

Probabilistic graphical models, structural equation models, and Bayesian causal networks are well suited to ESG analytics because they flexibly incorporate incomplete datasets, heterogeneous variable types, and latent constructs that cannot be directly observed. Their probabilistic foundations allow the integration of qualitative insights derived from LLM-processed unstructured data-such as text from sustainability reports, regulatory filings, customer reviews, or community feedback-transforming narrative information into quantifiable causal variables. This multimodal integration substantially enhances the granularity of sustainability assessment, especially in small-business contexts where structured data availability is limited.

Moreover, these models support simulation of alternative strategic or environmental scenarios, offering a forward-looking lens that complements backward-looking risk assessments. Through counterfactual reasoning, analysts can evaluate how strengthening governance transparency might reduce the distributional tail risk of operational disruptions, or how improvements in workplace equity may causally influence employee productivity under different economic cycles. Scenario-based causal simulation produces insights that are inherently policy-relevant-by clarifying mechanisms rather than prescribing interventions-thus maintaining the neutrality required for generalized investment and regulatory frameworks.

As sustainability challenges become increasingly intertwined with technological, environmental, and social dynamics, causal modeling provides a structured analytical foundation capable of capturing these interdependencies. Combined with LLM-enabled knowledge extraction, causal methods ensure that ESG analytics for small businesses remain scientifically rigorous, interpretable, and adaptable to evolving policy and market conditions.

5. Integrating LLMs and Causal Models into a Unified ESG Analytics Architecture

The integrated framework proposed in this study combines LLM-driven signal extraction with causal modeling to form a comprehensive ESG analytics system tailored to the realities of small-enterprise data structures. Rather than treating ESG as a purely reporting-based construct, the framework treats ESG as an evidence-based behavioral profile that can be recovered from operational traces and contextual documents. In the first layer, large language models parse narrative disclosures, regulatory and compliance records, supplier questionnaires, customer communications, internal policies, and other digital traces to generate structured sustainability indicators. In the second layer, causal models evaluate the directional influence of those indicators on resilience outcomes, including credit stability, revenue recovery after disruptions, operational continuity, and, where available, workforce retention and supplier reliability. This two-layer architecture is designed to ensure that qualitative meaning is captured without sacrificing statistical rigor, and that ESG assessments are not merely descriptive but analytically actionable.

The first layer addresses the central bottleneck of small-business ESG evaluation: the absence of standardized, audit-ready reporting. Most small firms document environmental, social, and governance behaviors incidentally-through invoices, maintenance logs, incident reports, training materials, emails, product descriptions, or local compliance filings-rather than through formal ESG statements. LLMs can transform these artifacts into a consistent indicator set by extracting key claims, identifying relevant practices, and mapping language to a defined ESG ontology. Importantly, the extraction process can be designed to retain traceability by linking each indicator to supporting passages and metadata (source type, date, and context), which allows analysts to audit why a risk flag was triggered or how a score was produced. The extraction layer can also quantify uncertainty: when evidence is weak or contradictory, the model can assign lower confidence, request additional documents, or route cases to human review. This is essential for small-business contexts where the absence of evidence should not automatically be interpreted as evidence of poor performance.

The second layer addresses a different-but equally important-problem: even high-quality ESG indicators do not automatically imply that ESG behavior improves resilience. Conventional ESG scoring often relies on correlation, which can lead to misleading conclusions if stronger firms simply have more resources to document sustainability or if sector effects confound the relationship between ESG and outcomes. Causal modeling introduces empirical discipline by distinguishing association from influence and by explicitly controlling for confounders such as industry, firm age, region, baseline profitability, and macroeconomic conditions. Depending on the policy or market question, the causal layer can implement Difference-in-Differences, matching, instrumental variables, or causal forests to estimate heterogeneous treatment effects. For example, the model can test whether implementing a documented safety program reduces operational downtime, whether energy-efficiency upgrades stabilize operating margins during inflationary shocks, or whether stronger governance routines reduce delinquency rates after revenue disruptions. The key contribution is directional clarity: causal estimates clarify what is likely driving observed outcomes rather than merely describing patterns.

This architecture resolves several longstanding challenges in small-business ESG evaluation. It reduces dependence on standardized reporting by treating unstructured evidence as first-class input, compensates for incomplete or uneven documentation by combining multiple weak signals into structured indicators, and enhances interpretability by linking each ESG conclusion to an explicit evidence trail. In addition, by grounding predictive analytics in causal relationships, the framework avoids purely associative judgments that might systematically disadvantage smaller or less formal firms. Instead of labeling a business as "high risk" because its documentation is sparse, the system can distinguish between genuine risk factors (e.g., repeated safety incidents, unresolved compliance violations, unstable labor practices) and simple reporting limitations. This distinction supports more equitable and defensible investment or lending decisions.

A further strength of the integrated design is its ability to support scenario analysis and forward-looking stress testing. Once ESG indicators are defined and their causal relationships to outcomes are estimated, the system can simulate resilience under alternative conditions: shifts in interest rates, supply chain disruptions, local climate hazards, changes in labor markets, or the adoption of new sustainability strategies. For instance, lenders can assess how a firm's probability of repayment changes if energy prices rise, and how much that risk is mitigated by energy-efficiency measures already inferred from operational records. Public programs can evaluate how targeted grants-such as funding for heat resilience, safety training, or compliance upgrades-would likely affect continuity outcomes for different types of small enterprises. Because the causal layer estimates effects rather than correlations, these scenarios can be framed as policy-relevant "what-if" analyses rather than speculative forecasts.

The synergy between generative AI and causal reasoning therefore provides both breadth and depth. LLMs supply breadth by handling heterogeneous, unstructured evidence streams at scale, translating narrative materials into consistent indicators across industries and regions. Causal models supply depth by enforcing methodological rigor, clarifying directionality, and quantifying the magnitude of ESG impacts under specified assumptions. Together, these components reduce uncertainty, improve transparency, and produce robust, scalable ESG analytics suitable for financial institutions, community development lenders, and public-sector programs seeking to evaluate sustainability in a way that reflects small-business realities rather than corporate reporting capacity.

6. Policy-Neutral Implications and the Future of Sustainable Small-Business Investment

Although this research maintains a policy-neutral orientation, the insights produced by AI-driven ESG analytics carry meaningful implications for institutional actors. Financial institutions can integrate the proposed framework to strengthen credit-risk assessment and expand access to capital for small businesses demonstrating strong sustainability potential. Community lenders may use causal resilience indicators to design lending programs that align long-term stability with ESG priorities.

For regulatory bodies, improved transparency and harmonized sustainability vocabularies enhance the quality of voluntary ESG disclosures without imposing burdensome reporting obligations. Public programs designed to support small-business resilience may benefit from more accurate ESG-readiness mapping, allowing interventions to be more targeted, equitable, and evidence-driven.

Future research should address model alignment, cross-sector data sharing, and the ethical governance of AI in sustainability analytics. Expanding multimodal datasets, formalizing LLM validation standards, and integrating real-time economic indicators will further refine resilience forecasting. Additionally, collaboration among academia, industry, and regulators will be essential to ensure AI systems strengthen sustainability outcomes without amplifying disparities or systemic biases.

7. Conclusion

This study demonstrates that AI-driven ESG analytics, grounded in the integration of LLM-based signal extraction and causal inference modeling, can significantly enhance sustainable investment evaluation for U.S. small businesses. The proposed framework addresses persistent data challenges, improves interpretability, and enables predictive assessment of resilience in a policy-neutral manner. By transforming unstructured disclosures into structured ESG indicators and modeling the causal pathways linking sustainability to operational outcomes, the system provides investors and institutions with tools for more equitable, transparent, and forward-looking decision-making. The findings affirm AI's transformative role in shaping the next generation of sustainable finance and highlight the need for continued interdisciplinary development to ensure accuracy, fairness, and accountability in ESG analytics.

References

1. M. Maghsoudi, N. Mohammadi, and M. Bakhtiari, "Artificial Intelligence and Sustainable Development: Public Concerns and Governance in Developed and Developing Nations," *Cleaner Environmental Systems*, 2025. doi: 10.1016/j.cesys.2025.100340
2. G. Giese, L. E. Lee, D. Melas, Z. Nagy, and L. Nishikawa, "Foundations of ESG investing: How ESG affects equity valuation, risk, and performance," *The Journal of Portfolio Management*, vol. 45, no. 5, pp. 69-83, 2019.
3. J. Pearl, and D. Mackenzie, "The book of why: the new science of cause and effect," *Basic books*, 2018.
4. G. Serafeim, "Social-impact efforts that create real value," *Harvard Business Review*, vol. 98, no. 5, pp. 38-48, 2020.
5. M. N. M. Khan, N. Fariha, M. I. Hossain, S. Debnath, M. A. Al Helal, U. Basu, and N. Gurung, "Assessing the Impact of ESG Factors on Financial Performance Using an AI-Enabled Predictive Model," *International Journal of Environmental Sciences*, vol. 1811, 1792.