European Journal of Business, Economics & Management

Vol. 2 No. 1 2026

Article Open Access

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

Zaolin Zhang [1,*]

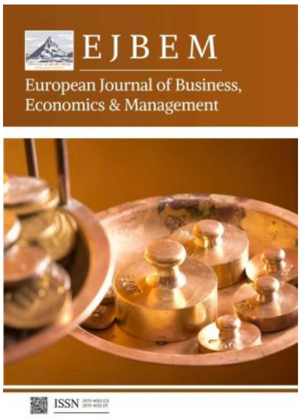

Received: 13 November 2025
Revised: 24 December 2025
Accepted: 06 January 2026
Published: 12 January 2026

[1] City University of New York, New York, New York, 10022, USA
* Correspondence: Zaolin Zhang, City University of New York, New York, New York, 10022, USA

Abstract: This study develops an integrated analytical framework that combines large language models with causal inference methods to strengthen sustainable investment decision-making in U.S. non-profit organizations. Drawing on advances in natural language processing, the proposed system applies an LLM-based architecture capable of interpreting regulatory texts, extracting domain-specific ESG signals, and synthesizing policy-relevant insights. Complementing this linguistic capability, the framework incorporates causal modeling-particularly Difference-in-Differences estimation-to identify the impact of policy changes on environmental, social, and governance performance factors. Together, these tools provide a structured foundation for supporting responsible investment strategies that align with mission-driven objectives. The model design also includes a multilayered feedback mechanism for continuous refinement, multilingual accessibility, and multi-format output generation, enabling diverse nonprofit stakeholders to access interpretable ESG results. The findings suggest that the integration of LLM-driven analytics with empirical causal evaluation enhances transparency, improves resilience in policy-sensitive contexts, and supports equitable governance practices. This research contributes to emerging scholarship on AI-enabled sustainability systems while offering practical implications for organizational strategy in the nonprofit sector.

Keywords: ESG analytics; sustainable investment; non-profit organizations; large language models; causal inference; policy analysis; resilience; artificial intelligence

1. Introduction

U.S. non-profit organizations increasingly rely on evidence-based frameworks to guide sustainable investment decisions, yet they often face structural barriers that limit access to advanced analytical tools. The rapid expansion of environmental, social, and governance policies and the growing complexity of regulatory environments heighten the need for accurate interpretation, timely insight generation, and transparent evaluation methods. At the same time, developments in artificial intelligence-particularly large language models-have introduced new possibilities for enhancing organizational capacity to synthesize unstructured information and identify emerging sustainability risks or opportunities. However, linguistic intelligence alone cannot fully address the empirical requirements for policy assessment [1]. Non-profits also need quantitative methods that clarify relationships between regulatory interventions and observed ESG outcomes. To meet this need, this study proposes an integrated framework combining LLM-based text interpretation with causal modeling approaches, with the objective of supporting sustainable investment decisions that are both mission-aligned and resilient to policy change. This integrated model seeks to offer a structured way for non-profits to navigate uncertainty, strengthen governance practices, and improve long-term strategic planning.

2. Conceptual Foundations of AI-Driven ESG Analytics

AI-driven ESG analytics is increasingly built on the convergence of natural language processing (NLP), structured data analysis, and policy evaluation, forming an integrated analytical framework capable of handling both qualitative and quantitative complexity. In the context of U.S. non-profit organizations, this convergence is especially significant because ESG performance is shaped by a highly fragmented landscape of policy environments, grant compliance requirements, tax regulations, and sector-specific reporting standards [2]. Unlike for-profit entities, non-profits often operate under multiple regulatory regimes simultaneously, including federal, state, and local oversight, as well as private foundation and donor-driven accountability mechanisms. As a result, the ability to interpret and synthesize diverse policy materials efficiently becomes a critical operational capability.

As highlighted in the uploaded document, advanced large language model (LLM) systems provide several core functionalities that directly support ESG analysis in this context. These include automated reasoning across complex texts, multi-format response generation, highlight-to-query interaction, and educational scaffolding that adapts explanations to different user expertise levels. Together, these features transform dense and unstructured policy documents-such as IRS guidance, environmental compliance rules, grant agreements, and governance bylaws-into concise, actionable summaries. This transformation significantly lowers the cognitive and technical barriers that traditionally limit non-profits' engagement with sustainability and compliance frameworks, enabling staff without specialized legal or data science training to participate meaningfully in ESG-related decision-making [3].

Beyond linguistic processing, effective ESG analytics also depends on the structured interpretation of relationships among environmental risk exposure, social impact outcomes, and governance compliance mechanisms. Traditional ESG evaluation frameworks typically rely on standardized rating systems or fixed scoring models designed primarily for corporate entities [4]. While these models offer comparability, they often lack the flexibility required by non-profit organizations that operate across diverse mission areas such as public health, education, environmental conservation, housing, or social justice. These organizations may prioritize impact dimensions that are underrepresented or poorly captured by conventional ESG metrics.

By incorporating AI-driven analytics, the ESG evaluation process becomes more adaptive and mission-aligned. Rather than applying a one-size-fits-all scoring methodology, AI systems allow users to define thematic priorities-such as climate resilience, community equity, workforce inclusion, ethical fundraising, or board governance-and generate tailored analytical outputs. LLMs can automatically extract relevant indicators from policy texts, reports, and internal documents, linking them to organizational goals and regulatory expectations. This capability supports scenario analysis, gap identification, and strategic planning, enabling non-profits to proactively address ESG risks and opportunities rather than reacting to compliance failures after the fact [5].

Furthermore, AI-driven ESG systems enhance accessibility and inclusivity through multilingual support and contextual explanation. Many U.S. non-profits serve linguistically diverse communities or collaborate with international partners, making it essential to interpret regulatory materials and sustainability guidelines across languages and cultural contexts. Multilingual AI tools allow organizations to translate, contextualize, and explain ESG-related requirements without losing semantic nuance, thereby improving transparency and stakeholder engagement. Educational scaffolding further ensures that complex sustainability concepts are communicated in a way that aligns with the knowledge level of board members, program staff, volunteers, and community stakeholders.

In combination, flexible linguistic interpretation and customizable indicator extraction create a more inclusive, responsive, and strategically useful ESG analytics framework. For non-profit organizations, this approach not only improves regulatory compliance and reporting accuracy but also strengthens mission alignment, organizational learning, and long-term sustainability. By leveraging AI-driven ESG analytics, non-profits can better navigate complex policy environments, demonstrate accountability to funders and communities, and integrate sustainability considerations into core organizational decision-making processes.

3. Integration of Large Language Models for ESG Interpretation

Large language models (LLMs) function as the cognitive layer of the proposed ESG analytic system, serving as the central mechanism through which complex information is interpreted, synthesized, and transformed into actionable knowledge. Their primary strength lies in the ability to process high-volume and heterogeneous textual data, including policy documents, investment disclosures, federal and state guidelines, sustainability frameworks, grant agreements, and non-profit organizational reports. These materials are typically unstructured, fragmented, and written in highly technical language, making manual interpretation both time-consuming and error-prone. By leveraging advanced natural language processing capabilities, LLMs can efficiently analyze these sources at scale and generate structured insights that support informed ESG-related decision-making.

As described in the uploaded document, the LLM system is designed with a multi-tiered output architecture that dynamically adapts to user needs. This architecture allows the system to deliver simplified explanations for general users, academically grounded discussions for researchers and policy analysts, and technically detailed responses for compliance officers or sustainability professionals. Such adaptability is particularly valuable in the non-profit sector, where stakeholders often possess diverse educational backgrounds and levels of technical expertise. By tailoring outputs to the user's cognitive and informational requirements, the system ensures inclusive engagement with ESG content while maintaining analytical rigor and accuracy.

In the context of sustainable investment and organizational governance, LLMs perform three primary interpretive functions that directly enhance ESG analytics. First, they extract ESG-relevant keywords and concepts from large bodies of text and classify them into predefined thematic categories, such as environmental stewardship, climate risk management, labor practices, community impact, ethical leadership, and governance integrity. This automated classification enables organizations to map complex documents onto recognized ESG dimensions without relying on rigid or overly simplistic scoring frameworks.

Second, LLMs identify and contextualize references to regulatory and compliance requirements embedded within legislative, administrative, and policy texts. These references may include obligations related to environmental reporting, labor standards, financial transparency, or board oversight. By highlighting relevant compliance elements and explaining their implications in accessible language, the system assists non-profit organizations in understanding their legal and ethical responsibilities. This capability is especially important for organizations operating across multiple jurisdictions or funding environments, where regulatory requirements may overlap or conflict.

Third, LLMs generate comparative and cross-document analyses that allow users to evaluate alignment between internal ESG objectives and external policy expectations. For example, the system can assess whether an organization's sustainability goals are consistent with federal climate guidelines, donor requirements, or recognized ESG frameworks. This comparative function supports strategic planning, risk mitigation, and performance evaluation by revealing gaps, redundancies, or areas of strong alignment. As a result, non-profit staff are relieved of much of the cognitive burden associated with navigating complex and evolving regulatory landscapes.

Beyond interpretive functionality, model alignment plays a critical role in ensuring that AI-generated outputs remain consistent with recognized ESG standards and ethical principles. By integrating human validation cycles, expert review protocols, and iterative feedback mechanisms, the system continuously refines its reasoning processes and reduces interpretive ambiguity. These alignment strategies enhance transparency, accountability, and trust in AI-assisted decision-making. In the nonprofit sector-where public credibility, donor confidence, and ethical governance are paramount-such safeguards are essential.

Through the combination of advanced interpretive capabilities and robust alignment mechanisms, LLM-based ESG analytics systems provide a reliable and responsible foundation for sustainability assessment. They empower non-profit organizations to engage more effectively with ESG frameworks, improve governance practices, and integrate sustainability considerations into long-term strategic decisions.

4. Causal Modeling for Policy Impact Assessment

While large language models (LLMs) provide significant interpretive depth and contextual understanding, they do not inherently establish causal relationships or quantify the magnitude of policy effects on organizational outcomes. LLMs excel at explaining what policies say, how they are framed, and how they relate to ESG concepts, but they cannot, on their own, determine whether observed changes in ESG performance are directly attributable to specific regulatory interventions. To address this limitation and ensure analytical rigor, causal inference methods-most notably Difference-in-Differences (DiD)-are incorporated into the ESG analytical framework. DiD offers a structured econometric approach for estimating causal effects by comparing changes in ESG outcomes for organizations exposed to a policy intervention with those for similar organizations that are not exposed.

Within the non-profit sector, this comparative architecture is particularly valuable. Non-profits operate in environments characterized by shifting funding conditions, economic cycles, demographic changes, and mission-driven strategic adjustments. These external and internal dynamics can produce fluctuations in ESG-related indicators that are unrelated to policy interventions. DiD analysis helps isolate the net effect of regulatory changes by filtering out common trends and time-invariant organizational characteristics. As a result, non-profits are better able to distinguish between ESG performance shifts driven by genuine policy impact and those caused by broader market forces, donor behavior, or organizational restructuring. This enhanced clarity enables decision-makers to identify which dimensions of environmental stewardship, social responsibility, or governance quality are most sensitive to regulatory change.

The causal modeling framework follows a series of systematic steps designed to improve the reliability, validity, and interpretability of estimated effects. The first step involves constructing appropriate treatment and control groups. Because non-profits are not randomly assigned to regulatory interventions, selection bias poses a major analytical challenge. To mitigate this issue, propensity score matching (PSM) is applied prior to DiD estimation. PSM matches organizations subject to a policy intervention with similar organizations that are not, based on observable characteristics such as organizational size, mission focus, geographic location, funding structure, and baseline ESG performance. This matching process reduces baseline imbalances and improves the comparability of the two groups, strengthening the plausibility of causal inference.

The second step focuses on validating the parallel trends assumption, which is central to the credibility of DiD estimates. Pre-trend analysis examines ESG outcome trajectories for treatment and control groups prior to policy implementation to ensure that they evolved in a similar manner before the intervention occurred. If pre-policy trends diverge significantly, post-policy differences may reflect pre-existing dynamics rather than the policy effect itself. By formally testing for parallel trends and visualizing outcome paths over time, the analytical framework enhances transparency and guards against spurious conclusions.

The third step incorporates a comprehensive set of robustness checks to assess the stability of estimated effects. Placebo tests introduce fictitious policy implementation dates to confirm that significant effects do not appear in periods where no intervention occurred. Sensitivity analyses evaluate how results change when model assumptions or sample definitions are modified, such as altering matching criteria or excluding influential observations. Alternative model specifications, including fixed-effects variations and dynamic DiD models, further test whether findings are robust to methodological choices. These validation layers collectively reduce vulnerability to unobserved confounders, measurement error, or random shocks that could otherwise distort inference.

Through this multi-stage causal modeling process, complex and often opaque policy environments are translated into quantifiable and interpretable ESG impacts. Rather than relying on anecdotal evidence or qualitative judgment alone, non-profit organizations gain empirical estimates of how specific regulations affect measurable outcomes, such as emissions reporting practices, workforce diversity metrics, financial transparency, or board governance structures. These estimates provide a data-driven foundation for strategic responses, including revising sustainability plans, reallocating resources toward high-impact initiatives, or designing targeted interventions to address policy-sensitive risk areas.

Importantly, causal modeling does not operate in isolation but complements the interpretive strengths of LLM-driven analysis. While LLMs decode the language, intent, and contextual significance of policy texts, causal inference methods measure the real-world effects of those policies on organizational behavior and performance. For example, an LLM may identify that a new environmental regulation emphasizes climate risk disclosure, while DiD analysis can determine whether organizations subject to the regulation actually improve disclosure quality or environmental performance relative to unaffected peers. This integration ensures that ESG evaluation is both conceptually grounded and empirically validated.

Together, LLM-based interpretation and causal inference form a comprehensive analytical system capable of supporting rigorous ESG assessment in the non-profit sector. This hybrid framework strengthens strategic decision-making by aligning qualitative understanding with quantitative evidence, enhances organizational resilience in policy-dependent environments, and promotes accountability to stakeholders. By combining narrative intelligence with causal measurement, non-profits are better equipped to navigate regulatory complexity, prioritize impactful actions, and sustain long-term mission effectiveness under evolving policy conditions.

5. System Architecture and Functional Design

The proposed ESG analytic system adopts a layered and modular architecture that integrates advanced language modeling, causal inference algorithms, and interactive user-facing tools into a unified analytical workflow. The uploaded document outlines a deployment model built on a serverless infrastructure capable of managing high-volume model requests, producing structured PDF outputs, logging user interactions, and storing feedback for iterative refinement. This structural design provides high scalability, supports parallel processing of analytical tasks, and ensures stable system performance under varying workloads. The serverless environment also allows real-time access to policy interpretations and ESG assessments, enabling non-profit organizations to obtain timely insights without maintaining complex technical infrastructure.

Functionally, the system incorporates several key design elements intended to accommodate diverse user needs. Multilingual processing expands accessibility for non-profits that engage multilingual communities or operate across culturally diverse regions. Through flexible output structuring, the system can generate executive summaries, detailed technical analyses, compliance-oriented documentation, and simplified explanatory notes, ensuring that both experienced policy analysts and general staff members can effectively use the platform. Automated documentation functionalities support the creation and archiving of compliance reports, organizational ESG summaries, and regulatory audit materials, significantly reducing administrative burden. Interactive text-highlighting features allow users to select specific policy segments and receive targeted, context-aware interpretations generated by the language model. This enables deeper engagement with policy content and supports efficient knowledge extraction. Moreover, a built-in feedback mechanism captures user evaluations, identifies misinterpretations or ambiguities, and incorporates these signals into ongoing model optimization, gradually enhancing system precision and practical relevance.

To promote responsible and trustworthy use, the system includes alignment-oriented safeguards designed to mitigate bias, reinforce factual accuracy, and ensure transparency in interpretive processes. These safeguards incorporate rule-guided reasoning templates, factual validation layers, and constraint-based response filtering. Together, they reduce the likelihood of inaccurate policy interpretations and ensure that model outputs remain consistent with recognized regulatory frameworks. Additionally, the system aligns with federal AI ethical guidelines by emphasizing fairness, explainability, and accountability, and reflects sector-specific standards for data governance and responsible analytics. Through this multi-layered approach, the system provides not only technical robustness but also an ethically grounded foundation for ESG evaluation in the U.S. non-profit sector, strengthening organizational capacity to navigate regulatory complexity and support sustainable, policy-responsive investment strategies.

6. Implications for Sustainable Investment and Organizational Resilience

Integrating LLM-driven ESG interpretation with causal modeling strengthens the strategic capacity of U.S. non-profits in several ways. First, it enhances regulatory comprehension by translating complex policy documents into consistent, actionable insights. Second, it improves decision quality by grounding sustainability planning in empirical evidence drawn from causal evaluation. Third, it supports equitable governance by making advanced analytical tools accessible to organizations that traditionally lack technical resources.

The combined framework also improves resilience against policy fluctuations. As federal and state-level ESG regulations continue to evolve, non-profits must adjust investment strategies and reporting practices accordingly. The system's ability to identify emerging regulatory trends, quantify their potential effects, and generate actionable recommendations supports adaptive capacity. Furthermore, by integrating structured outputs with user feedback, the model evolves alongside the policy environment, maintaining long-term relevance.

From a broader perspective, AI-enabled ESG analytics contributes to the national objective of promoting trustworthy, transparent, and equitable AI systems. By supporting non-profits-key actors in social welfare and community development-the model indirectly strengthens community resilience, supports sustainable economic development, and fosters inclusive participation in the sustainability economy.

7. Conclusion

This study presents an integrated framework that combines large language models with causal modeling to support sustainable investment strategies in U.S. non-profit organizations. The linguistic capabilities of LLMs facilitate comprehensive interpretation of regulatory texts and ESG datasets, while causal inference methods clarify the empirical impact of policy shifts on sustainability outcomes. Together, these tools promote transparency, improve decision quality, and strengthen resilience within the nonprofit sector. The system architecture incorporates multilingual support, structured outputs, automated documentation, and an iterative feedback loop to enhance accessibility and alignment across user groups. The findings underscore the potential for AI-driven ESG analytics to advance responsible governance practices and enable mission-driven organizations to navigate complex policy environments with confidence. Future research may explore integration with forecasting models, reinforcement learning simulations, and sector-specific ESG benchmarking to further refine the analytical capacity of the system.

References

1. L. Liang, T. Ding, R. Tan, and M. Song, "International ESG Standards and Frameworks," In *The ESG Systems in Chinese Enterprises: Theory and Practice*, 2025, pp. 25-61. doi: 10.1007/978-981-96-5459-8_2
2. A. Chambers, "Re: Artificial Intelligence Risk Management Framework," 2021.
3. D. Fougère, and N. Jacquemet, "Causal inference and impact evaluation," *Economie et Statistique/Economics and Statistics*, no. 510-511-512, pp. 181-200, 2019. doi: 10.24187/ecostat.2019.510t.1996
4. M. R. Anto, A. Shameem, and R. D. Ranjani, "SUSTAINABLE FINANCE AND ESG INVESTING: EMERGING TRENDS AND TECHNIQUES IN AI ERA," *INTERNATIONAL RESEARCH ON SUSTAINABILITY AND INNOVATION: A MULTIDISCIPLINARY APPROACH TO GLOBAL DEVELOPMENT*, vol. 296, 2025.
5. A. Romolini, S. Fissi, E. Gori, and M. Sances, "Integrating ESG into a Business Model: A Case Study from a Non-profit Organization," In *Environmental, Social, Governance (ESG) Risk, Performance, Monitoring*, 2025, pp. 705-718. doi: 10.1007/978-3-031-76618-3_33