JLP JOURNAL OF LAW AND POLITICS

DOI: https://doi.org//10.69648/VYRR6293

Journal of Law and Politics (JLP), 2025; 6(2): 21-32

jlp.ibupress.com

Online ISSN: 2671-3438

Application: 21.09.2025

Revision: 02.10.2025

Acceptance: 25.10.2025

Publication: 30.10.2025

Naumovski-Guli, B. (2025). AI as an early warning system: A framework for NGOs to predict failures in environmental policy implementation. Journal of Law and Politics, 6(2), 21-32. https://doi.org//10.69648/VYRR6293

Blagoja Naumoski Guli

University of Information Science and Technology St. Paul the Apostle, Ohrid, N. Macedonia

https://orcid.org/0009-0003-3969-1646

We declare no conflicts of interest

Correspondence concerning this article should be addressed to Blagoja Naumoski Guli

Email: blagoja.naumoski@mir.uist.edu.mk

BALKAN JOURNALS

AI as an Early Warning System: A Framework for NGOs to Predict Failures in Environmental Policy Implementation

Blagoja Naumoski Guli

Abstract

Environmental policies, though often ambitious in their conception, frequently falter during implementation due to a complex mix of political, economic, and social factors. This well-documented "implementation gap" creates a significant challenge for Non-Governmental Organizations (NGOs), which have traditionally operated in a reactive capacity, addressing environmental damage only after it has occurred. This paper puts forward a novel conceptual framework for an Artificial Intelligence-powered Early Warning System (EWS) designed to shift these organizations to a proactive and predictive model of advocacy. At its core, the system synthesizes varied, publicly available data streams to detect leading indicators of policy failure. It integrates computer vision to analyze satellite imagery for physical changes, Natural Language Processing (NLP) to monitor political and media discourse for shifts in sentiment and priority, and machine learning to analyze quantitative data and map complex networks. By systematically identifying subtle yet practical signals—such as resource deficits, negative discourse, and adverse lobbying activities—the EWS generates a predictive "failure risk score" for specific policies. This score is designed to equip NGOs to concentrate their limited resources on the most at-risk policies. As demonstrated through a detailed case study on corporate "No Deforestation" pledges in the Indonesian palm oil sector, the framework allows NGOs to engage policymakers with compelling, data-driven evidence before irreversible setbacks occur. This paper details the framework's architecture, outlines a methodology for its application, and considers its potential to redefine environmental advocacy by transforming data into pre-emptive action.

Keywords: Artificial intelligence (AI), non-governmental organizations (NGOs), environmental policy, early warning system (EWS), implementation gap

Introduction

The global environmental crisis has spurred the creation of landmark policies aimed at mitigating climate change, preserving biodiversity, and ensuring sustainable development. From the Paris Agreement to national-level carbon pricing, the legislative landscape is filled with ambitious commitments. Yet, a persistent and critical challenge remains: the gap between a policy's enactment and its effective implementation (Daly & May, 2020). A policy's success is determined not by its signing, but by its sustained execution on the ground. This execution is often undermined by a complex mix of social and political factors, with public perception frequently pointing toward governance failures as a primary cause (Aklin et al., 2014). This divergence between intent and outcome is known as the "implementation gap."

Figure 1:

Public perception of barriers to environmental law enforcement in Brazil (Aklin et al., 2014)

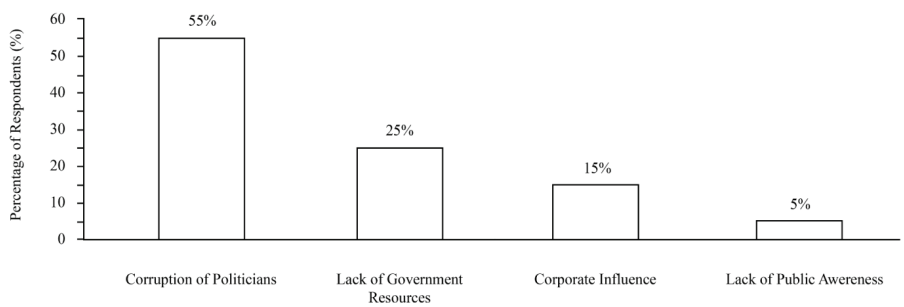

This chart displays the primary factors citizens identify as responsible for the poor enforcement of environmental laws. The data reveals that a majority of respondents blame political corruption, providing a clear, real-world example of the social and political challenges that contribute to the "implementation gap." This visual underscores the need for a system that can track such non-obvious indicators. Data is based on the survey findings from Aklin et al. (2014).

Within this context, Environmental Non-Governmental Organizations (ENGOs) serve as crucial, if often overmatched, watchdogs. Historically, their role has been to monitor outcomes, document failures, and hold governments and corporations accountable through public campaigns and legal challenges. While vital, this model is fundamentally reactive. By the time an NGO can definitively prove a policy has

failed—by pointing to rising emission levels or observed deforestation—the window for effective intervention has often closed. The advocacy becomes a post-mortem analysis rather than a preventative action, a significant limitation for organizations constrained by funding and personnel.

The sheer volume of data surrounding policy implementation now presents both a challenge and an opportunity. The digital exhaust of governance—parliamentary debates, budget proposals, corporate sustainability reports, news media, and lobbying registers—contains subtle signals of future outcomes. For human analysts, synthesizing these vast and disparate datasets in real-time is an impossible task (Burt & Taylor, 2021). This informational overload can further disadvantage NGOs, which may lack the resources to compete with the analytical capacity of governments and corporate interests.

This paper argues that Artificial Intelligence (AI) can provide the tools to bridge this gap, transforming NGO advocacy from a reactive to a proactive and predictive endeavor. We propose the development of an AI-powered Early Warning System (EWS), a conceptual framework designed to continuously monitor a portfolio of environmental policies and identify leading indicators of implementation failure. By flagging policies that are trending towards failure long before their stated goals are missed, the EWS can serve as a strategic tool, enabling NGOs to allocate their limited resources with maximum impact (Onuche, 2025). This research details the system's rationale, its conceptual architecture, a methodology for its application, and its potential legal and political implications for the future of environmental advocacy.

Related Work

This research is situated at the intersection of three distinct but converging academic domains: the study of environmental policy, the role of NGOs in governance, and the application of AI for social good. A review of the literature reveals a clear opportunity for a framework that connects these fields.

The concept of an "implementation gap" is a foundational challenge in political science. Scholarship has established that a policy's journey from law into real-world effect is anything but straightforward. This divergence between a central government's aims and local outcomes is often a product of a country's unique governance structure. The result is a persistent gap where ambitious laws exist on paper but lack rigorous enforcement, leading to continued environmental degradation

despite legislative action (Daly & May, 2020). Much of this research, however, focuses on diagnosing these failures retrospectively, rather than on developing tools for prospective risk identification.

The conversation then shifts to the actors trying to close this gap, where NGOs have emerged as critical non-state players. As environmental governance has become less state-centric, NGOs have taken on roles as monitors, advocates, and even enforcers. Their ability to form transnational networks and apply pressure on states and corporations is a significant force in global politics. Yet, their effectiveness can be hampered by a "paradox of scale," where authority built at a global level does not always translate to effective change locally (Balboa, 2018). Furthermore, NGOs operate in a resource-constrained environment, facing challenges in maintaining accountability, especially when competing with well-funded corporate interests (Kim, 2020). This reality makes it nearly impossible to manually track the complex variables that signal policy failure, highlighting a need for tools that can amplify their analytical capabilities.

Here, the application of AI for social and environmental good offers a potential solution. The field of "AI for Social Good" is gaining traction, with an emphasis on tackling societal problems through AI-based solutions (Al-Maliki et al., 2023). Initiatives like Microsoft's "AI for Earth" program aim to deploy AI to help monitor, model, and manage the planet's natural systems (Joppa, 2019). However, many such projects focus on purely environmental phenomena, such as tracking deforestation from satellite imagery or modeling climate impacts. They do not typically engage with the messy, data-diverse world of political implementation. Within the policy sphere, recent advancements in Natural Language Processing (NLP) have demonstrated a remarkable ability to automate the analysis of complex documents (Cano, 2020; Sampedro & Larrabide, 2020). While AI has been used to forecast humanitarian crises, its application to prospectively identify the risk of environmental policy failure for advocacy purposes remains a novel and underexplored area. This paper seeks to bridge that gap by adapting these powerful computational techniques to the unique strategic context of environmental NGOs.

The Implementation Gap in Practice: From Signals to Failures

While the implementation gap is a well-documented academic concept, its practical indicators are often buried in vast streams of public data. The challenge for NGOs is not a lack of information, but the inability to connect disparate signals to predict

a negative outcome. The failure of an environmental policy does not happen overnight; it is preceded by a series of measurable, observable events that serve as early warnings. Our framework categorizes these practical signals into three main types. The first, Resource Signals, includes quantitative data related to funding and staffing. For example, a national law protecting a rainforest may exist on paper, but a significant year-over-year budget cut to the agency responsible for its enforcement is a concrete signal of waning political commitment (Aklin et al., 2014). The second type, Discourse Signals, involves the qualitative analysis of language used by policymakers and corporations. A government that once championed a specific environmental initiative might subtly shift its language to prioritize competing interests, a change that can be detected by tracking the frequency of key terms in official documents and media reports (Sampedro & Larrabide, 2020). The final category, Activity Signals, tracks the actions of actors who may oppose the policy. A sharp, documented increase in lobbying meetings between a corporate consortium and the environmental ministry, cross-referenced with satellite data showing new development near a protected area (Hansen et al., 2013), provides a powerful, multi-layered signal of impending policy failure. By systematically tracking these practical signals, an EWS can transform the abstract concept of an "implementation gap" into a concrete, data-driven forecast of policy failure.

The Rationale for an AI-Powered EWS

Transitioning from a reactive to a proactive advocacy model is not merely an incremental improvement for NGOs; it represents a fundamental strategic shift. The rationale for developing an AI-powered Early Warning System emerges from the clear need to overcome the practical limitations these organizations face. The benefits can be understood as an integrated set of capabilities.

The system first confronts the overwhelming issue of Scale and Speed. The sheer volume of relevant information—legislative amendments, media reports in multiple languages, corporate filings, scientific data—far exceeds the capacity of any human research team. An AI system, however, can ingest, process, and connect these disparate data points in near real-time, uncovering subtle patterns and correlations that would otherwise remain invisible (Burt & Taylor, 2021). This provides a level of monitoring that is both broader and faster than anything currently possible. This speed is critical for the system's primary function: Early Detection. The goal is not to document failures that have already happened, but to detect the leading indicators that predict them. A budget cut to an environmental agency, a subtle change in

ministerial rhetoric, or a documented increase in lobbying by opposing industries are all signals that can precede measurable environmental decline by months or even years. Detecting these signals early provides a crucial window for intervention.

These early, data-driven insights, in turn, provide a foundation for Data-Driven Objectivity. NGO advocacy can sometimes be dismissed by policymakers as being purely ideological. An EWS changes this dynamic by arming organizations with quantitative, data-driven evidence. Approaching a government agency with a report that states, "This policy has a 75% risk of failure based on budgetary, political, and media signals," presents a far more powerful and credible argument than a general statement of concern. It shifts the basis of the conversation from opinion to evidence (Lawal et al., 2023). Ultimately, this more powerful approach allows for Strategic Resource Allocation. NGOs almost universally operate with finite resources of time, money, and personnel (Zavodna et al., 2025). The system essentially functions as a triage tool. Instead of spreading their resources thinly across the entire legislative landscape, organizations can use the EWS to focus their intensive advocacy efforts on the specific policies identified as most at-risk, ensuring maximum impact from their limited means (Onuche, 2025).

Case Study: AI-Driven Monitoring of Indonesian Palm Oil Deforestation

To illustrate the framework's practical application, we analyze the widespread corporate pledges for "No Deforestation, No Peat, No Exploitation" (NDPE) within the Indonesian palm oil sector. Since 2013, many of the world's largest agribusinesses have made public commitments to eliminate deforestation from their supply chains. However, significant deforestation continues, making this a prime example of an implementation gap where corporate policy fails in practice due to opaque supply chains and symbolic compliance. An AI-powered EWS is uniquely equipped to tackle this challenge by processing vast, varied data sets far beyond human capacity. Its role is not just to collect data, but to synthesize it into predictive intelligence.

The Role of AI in Monitoring NDPE Pledges

Computer Vision for Satellite Analysis (Quantitative Module): The EWS's core is an AI-powered computer vision model that continuously analyzes thousands of satellite images from public sources like the Global Land Analysis and Discovery

(GLAD) lab (Hansen et al., 2013). Trained on millions of images, the model automatically detects the visual signatures of deforestation—changes in forest texture and canopy cover—and flags new areas of cleared land in near-real time. This allows it to quantify the weekly deforestation rate (in hectares) within a specific supplier's concession area.

Natural Language Processing for Discourse Analysis (Discourse Module): The system uses NLP to scan and understand tens of thousands of news articles, NGO reports, and social media posts in both English and Indonesian. It performs sentiment analysis to gauge if the tone surrounding a supplier is becoming more negative and uses topic modeling to identify emerging risk factors, such as discussions about "peatland fires" or "lobbying efforts" connected to a specific company.

Machine Learning for Network and Risk Synthesis: Finally, a central machine learning model synthesizes these inputs. It analyzes public supplier lists and shipping manifests to map hidden relationships between high-risk mills and major consumer brands. It then learns from historical data which combination of signals—for example, a small deforestation event followed by negative media sentiment—most reliably predicts a major future violation, allowing it to calculate a final Failure Risk Score.

Application of the EWS

An environmental NGO using the EWS could proactively identify companies at high risk of violating their NDPE pledges. By integrating the AI-driven signals described above, the system generates a specific, data-rich alert.

Figure 2:

EWS alert dashboard for corporate NDPE pledge (Hansen et al., 2013)

Policy Monitored	Company	Failure Risk Score	Triggering Signals and AI-Generated Statistics (Past 30 Days)
NDPE Pledge (2015)	Global Food Corp.	78% (High Risk)	1. Quantitative Anomaly (Computer Vision): AI detected 150 hectares of primary forest cleared within Supplier Mill X's concession (Source: GLAD Alerts). 2. Discourse Shift (NLP): AI flagged a 45% increase in negative media sentiment in Indonesian news outlets related to Supplier Mill X and peatland cultivation. 3. Network Link (Machine Learning): AI analysis of shipping data confirmed 3 shipments from Supplier Mill X to Global Food Corp.'s main refinery.

How It Works: The Conceptual Framework

The proposed Early Warning System is best understood not as a single piece of software, but as a modular framework that processes varied data streams and synthesizes them into an actionable output. Its architecture integrates several AI-driven components to create a cohesive analytical pipeline.

Figure 3:

Conceptual framework of the AI-powered early warning system

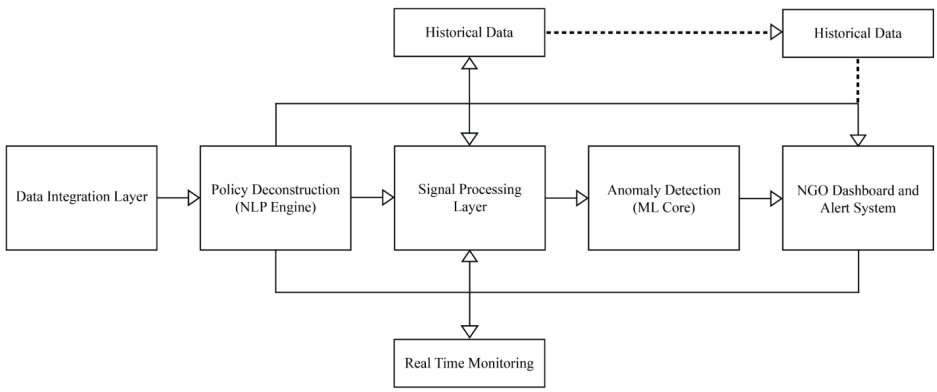

The system's foundation is its Data Ingestion Layer (Module 1), which is responsible for automatically collecting unstructured and structured data from a wide array of public sources, such as parliamentary websites, news databases, and budget portals. This information is then processed by the Policy Deconstruction (NLP Engine) (Module 2), where natural language processing is used to analyze legislative text and create a structured "implementation blueprint" of key commitments and metrics (Sampedro & Larrabide, 2020). This blueprint serves as the baseline against which all future data is compared.

The raw data flows to the Signal Processing Layer (Module 3), which synthesizes the deconstructed policy information with external inputs from real-time monitoring and historical data to identify trends. This processed information is then fed into the Anomaly Detection (Machine Learning Core) (Module 4), the analytical heart of the EWS. It uses the implementation blueprint as a baseline and employs machine learning models to flag statistically significant deviations—such as a budget cut or a missed deadline—as "failure signals." Finally, these signals are synthesized and displayed on the NGO Dashboard & Alert System (Module 5). This user interface provides a clear, dynamic "Failure Risk Score" and allows NGOs to see the specific data points that contributed to it, transforming complex data into actionable intelligence.

Methodology

The operationalization of the proposed EWS framework follows a five-stage methodological process, designed to systematically collect data, train predictive models, and validate their output.

1. Corpus Definition and Data Sourcing: The initial stage involves the selection of a representative corpus of 10-15 significant environmental policies with varying implementation outcomes. For each policy, a multimodal dataset is compiled, comprising a) official legislative and policy texts, b) quantitative time-series data, including agency budgets, enforcement statistics, and relevant environmental metrics (e.g., emissions data), and c) a qualitative corpus of news articles, parliamentary records, and corporate reports.
2. Data Pre-processing: The raw, heterogeneous data is then cleaned, structured, and normalized. Textual data is standardized for Natural Language Processing (NLP), and numerical data is formatted into consistent time-series to enable trend analysis and anomaly detection.
3. Model Training and Feature Engineering: This stage utilizes a supervised learning approach. The historical dataset, with known policy outcomes (i.e., "success" or "failure"), is used to train the machine learning models. Key predictive "features" are engineered from the data, such as the rate of change in an agency's budget, the sentiment polarity of media coverage, or the frequency of specific keywords in political discourse. The model learns the statistical correlations between these features and the eventual implementation outcome.
4. Risk Scoring and Validation: To ensure predictive accuracy, the trained model is validated against a holdout set of historical policies it has not previously analyzed. The model's predictions are compared to the known outcomes, and its performance is evaluated using standard statistical metrics. This process allows for the fine-tuning of the "Failure Risk Score" algorithm to ensure its reliability and robustness.
5. Output and Visualization: The final stage involves the design of a user-facing dashboard. This interface is engineered to translate complex analytical outputs into an intuitive visualization, such as a "report card" for each monitored policy, often using a traffic light system (green, yellow, red) to provide an at-a-glance assessment of risk and direct an NGO's attention to the most critical areas.

Results

The intended output of the EWS is a clear, data-driven alert that synthesizes multiple signals into an actionable narrative. The system moves beyond simply documenting past environmental damage by providing the predictive intelligence necessary for pre-emptive intervention. For example, in monitoring the NDPE pledge of a company like "Global Food Corp.," the EWS would generate a critical alert upon the convergence of several leading indicators.

The NGO dashboard would show that the company's policy has reached a high failure risk score of 78%. This score would be substantiated by tangible evidence from the system's modules. It would pinpoint a quantitative anomaly detected by the computer vision model: 150 hectares of recent primary forest clearing within the concession of a key supplier, Mill X (Hansen et al., 2013). This alert would be automatically cross-referenced with a significant discourse shift identified by the NLP model, such as a 45% increase in negative media sentiment surrounding Mill X's operations in Indonesian news outlets. Finally, the machine learning core would confirm a direct network link, verifying recent shipments from the at-risk mill to Global Food Corp.'s main refinery.

This synthesized alert provides the NGO with a specific and undeniable evidence package. It empowers them to move beyond general accusations and engage Global Food Corp. with a targeted, data-backed demand for action regarding a specific supplier, effectively closing the gap between the company's pledge and its real-world impact.

Discussion and Limitations

The implementation of an AI-powered EWS could represent a paradigm shift for environmental advocacy, moving the sector from a reactive to a proactive stance. This paper's primary contribution is a framework designed to level the informational playing field. Corporations and governments already utilize sophisticated data analytics; the proposed EWS provides a powerful counterpart for civil society, fostering a more transparent and accountable governance process. The political implications are significant. As demonstrated in the Indonesian palm oil case study, pre-emptive campaigns armed with satellite imagery and supply chain data are far more potent than retroactive condemnations.

However, several limitations must be acknowledged. The most immediate challenge is the system's reliance on Data Accessibility and Bias. The framework assumes access to public data, but in political systems with low transparency, sourcing unbiased budgetary or lobbying information would be difficult. Furthermore, all public data carries inherent biases, and an AI model could amplify these distortions if not carefully managed (Torrance & Tomlinson, 2025). Another significant hurdle is the "Black Box" Problem. For an NGO to credibly use a 78% failure risk score, it must be able to explain the contributing factors. Incorporating principles of Explainable AI (XAI) is therefore not just beneficial but essential for the tool's legitimacy (Fundraising.AI, 2025).

Furthermore, two strategic challenges exist. The first is Adversarial Adaptation. As NGOs begin using such systems, governments and corporations may adapt their behavior, using more sophisticated language to mask intentions or burying critical data in obscure reports. This would create an "adversarial game" requiring the EWS models to be continuously updated. The second is the risk of a Technological Divide. The resources required to build and maintain such a system could be prohibitive for smaller, grassroots NGOs, potentially deepening the resource gap between large international organizations and local ones (Zavodna et al., 2025). Addressing this will be critical for the equitable deployment of this technology.

Conclusion

The gap between what environmental policies promise and what they deliver remains one of the most significant barriers to addressing the global climate crisis. This paper has proposed a potential pathway to closing that gap. The AI-powered Early Warning System framework offers a way to turn vast, complex datasets into clear, predictive, and actionable intelligence, empowering NGOs to evolve. It allows them to move beyond documenting past failures and towards actively preventing future ones. This proactive stance, grounded in empirical evidence, has the potential to make environmental advocacy more efficient, strategic, and ultimately, more effective.

While significant technical and ethical challenges remain, the development of such systems represents a critical and necessary step in ensuring that the environmental policies of today become the ecological realities of tomorrow. Future work should focus on developing a pilot of the EWS, testing it on historical policies to validate its predictive accuracy, and exploring the creation of open-source modules to make this technology more accessible to the entire environmental community, thereby mitigating the risk of a technological divide within the sector.

References

Aklin, M., Bayer, P., Harish, S. P., & Urpelainen, J. (2014). Who blames corruption for the poor enforcement of environmental laws? Environmental Economics and Policy Studies, 16(4), 365–384.

Al-Maliki, S., Qayyum, A., Ali, H., Abdallah, M., Qadir, J., Hoang, D. T., Niyato, D., & Al-Fuqaha, A. (2023). Adversarial machine learning for social good: Reframing the adversary as an ally. arXiv. https://doi.org/10.48550/arXiv.2310.03614

Balboa, C. M. (2018). The paradox of scale: How NGOs build, maintain, and lose authority in environmental governance. MIT Press.

Burt, E., & Taylor, J. (2021). AI unlocking the potential of NGOs. International Journal of Research and Innovation in Social Science, 9(7), 5184–5197.

Cano, A. (Ed.). (2020). Social media and machine learning. IntechOpen.

Daly, E., & May, J. (2020). Environmental constitutionalism and the implementation gap. International Journal of Constitutional Law, 18(1), 308–312.

Fundraising.AI. (2025). *A framework toward responsible and beneficial AI for fundraising*. Fundraising.AI. https://fundraising.ai/

Hansen, M. C., et al. (2013). High-resolution global maps of 21st-century forest cover change. Science, 342(6160), 850–853.

Joppa, L. (2019, April 17). AI for Earth [Audio podcast]. In Microsoft Research Podcast. Microsoft.

Kim, E. (2020). The limits of NGO-led environmental governance: The case of forest certification in Indonesia. Third World Quarterly, 41(3), 497–516.

Lawal, A. M., Vincent-Akpu, I. F., & Woke, G. N. (2023). Evaluation of stakeholder perceptions and applicability of artificial intelligence integration into environmental impact assessment processes in Nigeria. Journal of Environmental Science and Sustainable Development, 6(2), 1–15.

Onuche, J. (2025). Enhancing non-profit efficiency and impact through data-driven strategies: Addressing challenges and leveraging emerging technologies – A literature review. Texila International Journal of Management, 11(1), 1–12.

Sampedro, J., & Larrabide, I. (2020). A review on the application of text mining to the analysis of scientific and technical information in the energy sector. IEEE Access, 8, 211833–211847.

Torrance, A. W., & Tomlinson, B. (2025). Organic websites: Certification of AI-generated or human-written content on the internet. Chicago-Kent Journal of Intellectual Property, 24(1), 72–100.

Zavodna, L. S., Überwimmer, M., & Frankus, E. (2025). Barriers to the implementation of artificial intelligence in small and medium sized enterprises: Pilot study. Journal of Competitiveness, 17(1), 332–345.