



# Large Language Model Application for Regulatory Horizon Scanning: Case Study on Anti-Greenwashing Regulations



## **Large Language Model Application for Regulatory Horizon Scanning: Case Study on Anti-Greenwashing Regulations**

Hao Zhang, Daniel Dao, James Bowden and Mark Cummins  
Strathclyde Business School, University of Strathclyde, and Financial Regulation  
Innovation Lab

31<sup>st</sup> March 2025

We acknowledge funding from Innovate UK, award number 10055559.

Corresponding authors: Hao Zhang, Daniel Dao, James Bowden and Mark Cummins

Strathclyde Business School, University of Strathclyde, 199 Cathedral Street, Glasgow,  
G4 0QU, UK

Email: [hao.zhang@strath.ac.uk](mailto:hao.zhang@strath.ac.uk)

Email: [daniel.dao@smithschool.ox.ac.uk](mailto:daniel.dao@smithschool.ox.ac.uk)

Email: [james.bowden@strath.ac.uk](mailto:james.bowden@strath.ac.uk)

Email: [mark.cummins@strath.ac.uk](mailto:mark.cummins@strath.ac.uk)

# Financial Regulation Innovation Lab

## Who are we?

The Financial Regulation Innovation Lab (FRIL) is an industry-led collaborative research and innovation programme focused on leveraging new technologies to respond to, shape, and help evolve the future regulatory landscape in the UK and globally, helping to create new employment and business opportunities, and enabling the future talent.

FRIL provides an environment for participants to engage and collaborate on the dynamic demands of financial regulation, explore, test and experiment with new technologies, build confidence in solutions and demonstrate their ability to meet regulatory standards worldwide.

## What is Actionable Research?

FRIL will integrate academic research with an industry relevant agenda, focused on enabling knowledge on cutting-edge topics such as generative and explainable AI, advanced analytics, advanced computing, and earth-intelligent data as applied to financial regulation. The approach fosters cross sector learning to produce a series of papers, actionable recommendations and strategic plans that can be tested in the innovation environment, in collaboration across industry and regulators.

**Locally-led Innovation Accelerators delivered in  
partnership with DSIT, Innovate UK and City Regions**

# Large Language Model Application for Regulatory Horizon Scanning: Case study on Anti-Greenwashing Regulations

Hao Zhang\* Daniel Dao\*\* James Bowden\* Mark Cummins\*

*\* University of Strathclyde*

*\*\* University of Oxford*

March 2025

**Abstract:** This white paper explores the application of Generative AI, specifically Large Language Models (LLMs), to enhance regulatory horizon scanning within financial services. Using the Financial Conduct Authority's (FCA) 2024 anti-greenwashing rule as a case study, we demonstrate how LLMs can be integrated into the strategic foresight process to detect early regulatory signals, analyse stakeholder feedback, and forecast future regulatory developments. Our framework builds upon the traditional horizon scanning model, comprising exploration, assessment, application, and continuation, and incorporates advanced text analysis techniques including semantic similarity testing with models such as BERT and RoBERTa. The study shows that LLMs can significantly improve the efficiency, accuracy, and scalability of horizon scanning by extracting meaningful insights from large, unstructured datasets. The results highlight the potential of LLM-driven foresight to enhance regulatory preparedness, guide compliance strategies, and inform policy design in an increasingly complex and dynamic regulatory environment.

# Table of Contents

I. PROBLEM STATEMENT.....	6
II. SOLUTION FRAMEWORK .....	8
2.1. <i>Horizon scanning process</i> .....	8
2.2. <i>Integrating Large Language Models into the horizon scanning process</i> .....	10
III. USE CASE DEMONSTATION.....	14
3.1. <i>Development timeline of the FCA’s anti-greenwashing rule</i> .....	14
3.2. <i>Data and filtering criteria</i> .....	15
3.3. <i>Analysis</i> .....	20
3.4 <i>Formal Similarity Testing of LLM-Based Anti-Greenwashing Rule Prediction</i> .....	30
IV. CONCLUSION .....	34
REFERENCE.....	36
ABOUT THE AUTHORS .....	38

## I. PROBLEM STATEMENT

The growing complexity of regulatory frameworks poses substantial challenges for corporations striving to ensure compliance within an evolving and dynamic institutional landscape. Given these significant challenges in keeping pace with the regulations, industry leaders (KPMG, PWC, Grant Thornton, among others<sup>1</sup>) have recognised the critical importance of regulatory foresight tools and applications. Regulatory horizon scanning has emerged as a valuable technique to address these challenges. It involves systematically anticipating future developments by detecting early signals around and identifying trends in

regulatory evolution. Horizon scanning enables businesses to capitalise on opportunities and manage risks effectively (Delaney, 2014; Cuhls et al., 2015). Traditional horizon scanning methods often rely on expert judgment and structured frameworks, which can be constrained by limited data processing capacity, and the inability to detect weak signals among vast unstructured information frameworks become ever more complex and dynamic, there is a growing need for innovative techniques that can enhance the scalability, objectivity, and predictive power of horizon scanning, enabling more effective strategic foresight.

---

<sup>1</sup> Please see some reports:

KPMG - <https://kpmg.com/us/en/capabilities-services/advisory-services/risk-and-compliance/financial-services-regulatory-compliance-risk/regulatory-change-management/horizon-scanning.html>

PWC -

<https://store.pwc.co.uk/en/products/horizon-scanning-portal>

Grant Thornton -

<https://www.grantthornton.co.uk/insights/regtech-automating-regulatory-change-compliance/>

The Financial Regulation Innovation Lab (FRIL) has further highlighted this industry need through its AI and Emerging Technologies innovation call, which includes specific requests for RegTech solutions to accelerate the identification and assessment of applicable regulations. This industry-driven challenge promotes the necessity for advancing research in regulatory risk management, particularly in leveraging artificial intelligence techniques. The aim of this white paper stems from this industry demand, contributing to the development of innovative methodologies - Generative AI and Large Language Models (LLMs) - in regulatory horizon scanning that enhance regulatory preparedness and strategic compliance management.

Our application context is the Financial Conduct Authority's (FCA) newly implemented anti-greenwashing rule. Effective in the UK since May 31, 2024, the regulation requires

FCA-authorized firms to ensure that all communications referencing environmental or social characteristics of their products and services are fair, clear, and not misleading. This rule, with objectives of protecting consumers from misleading sustainability-related claims, provides guidance to assist firms in achieving compliance. Firms that fail to align with these standards may face supervisory action, emphasising the importance of robust compliance frameworks. The rule's development involved an extensive consultation process, commencing in late 2023 and concluding with the final guidance in April 2024. We suggest that by tracking stakeholders' responses and feedback, and wider discourse, during the consultation process, regulatory horizon scanning could have helped organisations understand the trajectory of this regulatory evolution and predict regulatory changes in the finalised rule. Applying Generative AI – specifically LLMs -

within the regulatory horizon scanning process, we leverage early text signals (responses, feedback, comments) from a wide range of information sources in a consultation process to anticipate future regulations. Applying this model to the FCA's anti-greenwashing regulation offers insights into how businesses could have proactively aligned with the emerging compliance requirements, ensuring greater readiness and reducing risks associated with regulatory shifts.

## **II. SOLUTION FRAMEWORK**

### **2.1. *Horizon scanning process***

Horizon scanning is a vital strategic foresight tool that directly informs policy and regulatory processes by systematically identifying and analysing early signals and trends (Kim et al., 2019; Idoko and Mackay, 2021). By leveraging patterns of surrounding information, it identifies indicators of change to uncover

potential risks and opportunities, providing a structured foundation for proactive decision-making (Delaney, 2014; Cuhls et al., 2015).

A key strength of horizon scanning lies in its ability to detect future-oriented issues at an early stage, enabling organisations, industries, governments, and societies to implement proactive measures or adapt to evolving trends (Wintle et al., 2020). This process facilitates the anticipation of emerging threats and opportunities, fostering the development of informed strategies to mitigate risks and capitalise on positive changes (Sutherland et al., 2019; Palomino et al., 2012). Importantly, horizon scanning does not aim to produce deterministic projections; instead, it offers a dynamic and objective approach to continuously exploring, monitoring, and assessing present developments and their long-term implications (Miles and Saritas, 2012).



At its core, horizon scanning is a data-driven process that involves systematic collection, interpretation, and validation of relevant information to enable evidence-based decision-making and policy development (Marsh et al., 2014). Upon reviewing various horizon scanning approaches, Rowe et al. (2017) identify a shared framework (process) underpinning these approaches, which typically involves four key stages: Exploration, Assessment, Application, and Continuation (Figure 1).

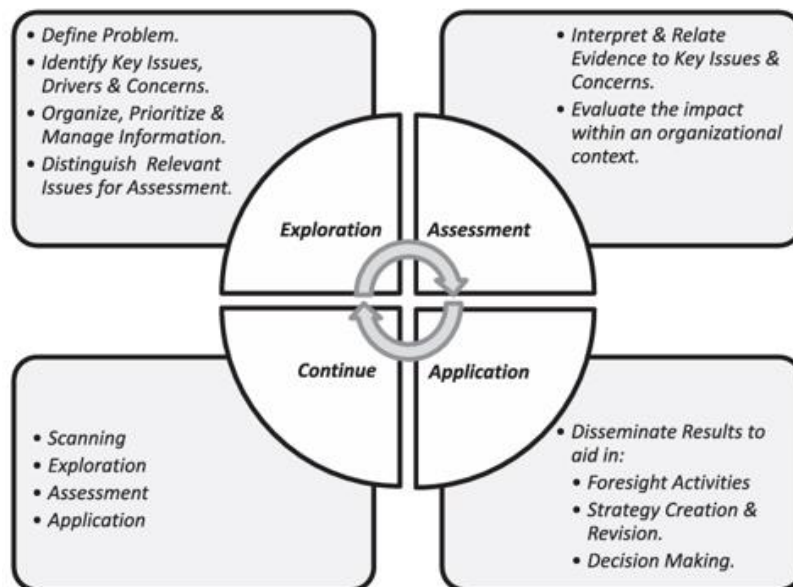
**Exploration:** The initial step in horizon scanning involves exploring the external environment by systematically scanning and gathering all relevant information. Traditionally, this has been achieved through manual methods such as expert consultations during workshops or brainstorming sessions. However, with the advancement of the internet and technological

tools, this exploration phase is increasingly conducted using automated and web-based collection methods.

**Assessment:** Once the information is gathered, the next step is to extract evidence from the data to assess critical issues and address public concerns effectively.

**Application:** This stage focuses on communicating the results of the assessment phase to support activities such as foresight development, strategic planning, policy formulation or revision, risk evaluation, and informed decision-making.

**Continue:** The final stage acknowledges the importance of continuing horizon scanning activities to update changes in the surrounding environment and organizational knowledge, thus enhancing the decision-making process.



**Figure 1: Horizon scanning process - Source: Rowe et al. (2017)**

**2.2. Integrating Large Language Models into the horizon scanning process**

As outlined, horizon scanning is an essential strategic process aimed at identifying emerging trends, risks, and opportunities that may significantly impact regulatory frameworks and market dynamics. Within the financial regulatory landscape, characterised increasingly by attention to sustainability and anti-greenwashing measures, effective horizon

scanning enables both regulatory bodies and financial institutions to anticipate changes proactively and adjust their strategies accordingly. However, the proliferation of vast amounts of unstructured data from diverse sources such as news articles, industry reports, social media, and multimedia content poses significant challenges for systematic and efficient analysis.

Recent advancements in Large Language Models (LLMs) offer promising solutions to these challenges, notably through their sophisticated capabilities in natural language understanding, information extraction, sentiment analysis, topic modeling, summarisation, and text generation. These features make LLMs particularly valuable in enhancing the efficiency and effectiveness of the horizon scanning process, especially during the assessment and application phases. Specifically, the natural language understanding capability of LLMs enables accurate interpretation and contextualisation of complex textual data, which is fundamental for processing the heterogeneous and voluminous information gathered during the initial exploratory stages of horizon scanning. Further, through robust information extraction techniques, LLMs can efficiently pinpoint and retrieve critical insights such as emerging regulatory developments, shifts in stakeholder

priorities, and significant market signals, effectively distilling meaningful patterns from extensive datasets.

Building on the application of Generative AI for simplified ESG reporting in financial services as set out by Hao et al. (2025), which demonstrates efficacy of LLMs in processing complex sustainability data, we now investigate the extension of these capabilities to the regulatory foresight domain. LLMs offer the capability to analyse temporal data across diverse sources to monitor and interpret the evolution of regulatory topics and market trends. This capability not only enhances foresight but also provides regulatory bodies and financial institutions with a more agile and responsive approach to managing regulatory compliance and strategic risk.

In this study, conducted during the consultation period for the United Kingdom's new anti-greenwashing rule, we aim to enhance the

horizon scanning process by leveraging LLMs. This approach allows us to gather a more diverse corpus of textual data from the internet, while rigorously applying predefined screening criteria to ensure precise and relevant sample selection. Our data collection strategy draws on content from a range of online forums and discussion platforms, capturing public discourse and stakeholder commentary. We specifically highlight the perspectives of companies and financial institutions regarding the anti-greenwashing rule, as these viewpoints are directly aligned with the rule's objectives. The anti-greenwashing rule is a regulatory measure that prioritises combating greenwashing in the form of misleading or overstated sustainability claims and empowering consumers to make informed decisions aligned with their sustainability preferences.

The rule also aims to create a level playing field for firms in an evolving market by encouraging greater transparency and by promoting products and services that genuinely represent sustainable choices. By strengthening stakeholder trust in companies' sustainability-related claims about their products and services, these measures are expected to bolster overall market confidence. Integrating advanced tools such as LLMs into the horizon scanning process offers an efficient way to manage the rapidly growing volume of online information. By automating and scaling up the identification of relevant discussions and perspectives, our approach helps capture emerging issues and stakeholder views more comprehensively.

To clarify our approach, we map the details to the high-level stages covered in Section 2.1.

**Exploration:** The exploration phase, spanning November 28, 2023, to January 1, 2024, involved automated web-based data collection, focusing exclusively on text-based sources. This included news articles, reports, stakeholder feedback, and relevant publications<sup>2</sup> on the FCA's anti-greenwashing rule. The aim of the process is to gather a wide array of information to support subsequent analysis, ensuring a foundation for identifying regulatory trends.

**Assessment:** In the assessment phase, we utilised LLMs to analyse the collected data. This analysis involved identifying high-frequency keywords, capturing concerns raised by firms, investment institutions, and assessing public opinions on key issues.

**Application:** The application phase focused on disseminating the insights gained during the

assessment. Leveraging LLMs, we attempted to predict potential responses and key information to be included in the finalised anti-greenwashing guidance. Additionally, we developed similarity tests to formally compare our predictive outcomes with the finalised guidance, allowing for an evaluation of the alignment and accuracy of the predictions against actual outputs.

**Continue:** The continuation phase of the horizon scanning process emphasised iterative refinement to maintain both reliability and adaptability. Specifically, this involved clearly outlining the models, scenarios, and observations employed, and directing readers to relevant sections of the document for further detail. Regarding the models, the study integrates traditional methods, such as Jaccard

---

<sup>2</sup> News articles were sourced from major financial outlets like Financial Times (Financial Times), Bloomberg (Bloomberg), and Reuters (Reuters), covering the FCA's policy statement PS23/16 on sustainability disclosures.

Reports included analyses from industry bodies like the Investment Association (Investment Association) and consulting firms such as KPMG (KPMG) and PwC (PwC), discussing the implications of the anti-greenwashing rule.

similarity and N-gram analysis, alongside advanced approaches like BERT, and customised LLMs, allowing for comprehensive comparisons and optimal selection tailored to specific analytical needs.

In the following section, we will present a comprehensive overview of the detailed results obtained from each step of the horizon scanning process. This demonstration will highlight key insights and findings related to the horizon scanning conducted for the FCA's anti-greenwashing rule, offering valuable perspectives on the application of this methodology in addressing emerging regulatory challenges.

### **III. USE CASE DEMONSTATION**

#### ***3.1. Development timeline of the FCA's anti-greenwashing rule***

Development of the FCA's anti-greenwashing rule involved several distinct phases, from the

initiation of the feedback solicitation process to the eventual publication of the finalised guidance. The development process included a structured two-month consultation stage, during which stakeholders (consumers, corporations, and investment institutions) were invited to provide input on the proposed regulatory framework. This consultation phase commenced on November 28, 2023, and concluded on January 26, 2024 (Figure 2). Following the consultation period, the FCA dedicated three months to analysing the formal feedback received and incorporating necessary revisions. The finalised guidance was subsequently published in April 2024 (Figure 2), marking the culmination of the rule's development process.

This phased approach highlights the iterative nature of regulatory development, wherein stakeholder input plays a pivotal role in shaping the final outcomes. The timeline underscores

the importance of the consultation process as a critical window for anticipating regulatory changes. By analysing the information gathered during this stage, organisations can

proactively predict shifts in regulatory frameworks, assess potential impacts, and align their responses accordingly.



**Figure 2: Development timeline of the Financial Conduct Authority's (FCA) anti-greenwashing rule**  
*- Source: Financial Conduct Authority (FCA)*

### **3.2. Data and filtering criteria**

We employed a horizon scanning approach, which involved systematically collecting and analysing data from Google search results of stakeholders' feedback and concerns pertaining to the FCA's anti-greenwashing rule. Specifically, our data contains three components: (i) text data sourced from mainstream platforms such as Google and LinkedIn, comprising the formal consultation responses and institutional commentaries,

along with news articles, relevant blogs and forum discussions; (ii) video content (from YouTube), and (iii) audio content (i.e. audio discussion programs around the anti-greenwashing rule). For the purposes of our study, we use only the text data type. Future research will consider the audio and video information sources. To ensure the relevance and validity of the data collected, we apply filtering criteria based on the following

standards set around the FCA's consultation process for the anti-greenwashing rule:<sup>3</sup>

1. Feedback must be correct and capable of being substantiated;
2. Content must be clear and presented in a way that can be readily understood;
3. Feedback should be complete, avoiding omissions or the concealment of critical information while considering the full lifecycle of the product or service in question;
4. Claims must be fair and meaningful, especially when comparisons are drawn between products or services.

Using these criteria as a foundation, stakeholders were prompted to address three core questions during the consultation:

- a) Does the proposed guidance clarify the anti-greenwashing rule? If not, what additional measures could enhance clarity?
- b) Do you have any specific comments on the proposed guidance, including the examples provided?
- c) Do you agree with the proposed enforcement date of 31 May 2024 for the guidance?

As noted, the consultation period was limited to November 28, 2023, to January 26, 2024. The dataset constructed for this study was compiled to capture public and stakeholder commentary regarding the UK's proposed anti-greenwashing rule over this timeframe. Using automated web-based methods, such as targeted searches via Google and LinkedIn, we initially collected 10,471 feedback artefacts. These artefacts encompassed multiple content

[guidance/fg24-3.pdf](#)

---

<sup>3</sup> For a summary of the standard used see <https://www.fca.org.uk/publication/finalised->



formats: approximately 75% were native text-based content, including news articles, reports, blogs, and online forum discussions; about 20% originated from podcasts, converted into text using automated speech-to-text transcription; and roughly 5% were derived similarly from video sources, including webinars and interviews, similarly converted into text using automated speech-to-text transcription. Again, for the purposes of this study, we only use the text data.

To ensure the relevance and reliability of our dataset, we implemented strict filtering criteria aligning with the principles outlined above. To scale the filtering across the extensive text data collected, the implementation is automated through a natural language processing (NLP) based textual analysis. Firstly, artefacts were required to be directly relevant to the anti-greenwashing rule and the context of its official consultation. We included only content that

explicitly discussed the rule's provisions, implications, or feedback to the FCA's proposals. Comments that were tangential or unrelated (for example, general statements about "greenwashing" without reference to the UK rule or its consultation) were excluded. Secondly, we selected substantive feedback that offered structured arguments or targeted suggestions regarding the rule. To be included, an entry needed to articulate clear points – such as critiques of the rule's scope, support or concerns with its requirements, or recommendations for implementation. Simple, unelaborated opinions (e.g., a one-line remark of support or opposition with no detail) or off-topic musings were filtered out. This criterion ensured that the dataset focused on meaningful, content-rich responses that could inform analysis (each entry typically contained reasoned commentary or concrete proposals rather than vague sentiments). Thirdly, we

eliminated duplicate content to ensure unique contributions. If the same textual entry appeared multiple times (for instance, a comment reposted across different websites or a news article quoted verbatim on a forum), we retained only one representative instance. We also removed entries that were essentially identical paraphrases or copies of other feedback. This de-duplication step ensured that each data point in the final set was a distinct piece of feedback, preventing any single individual's comments from being over-represented and avoiding bias from repeated content.

Finally, we discarded contributions that were poorly written to the point of being unintelligible or that contained largely disjointed text (e.g., heavy use of broken sentences or irrelevant strings of characters from web scraping noise). For transcribed audio/video content, we imposed an

additional quality control: we required a high transcription confidence level from the speech-to-text system. Transcribed artefacts were included only if the automated transcription's confidence score exceeded a predetermined threshold (ensuring the words were recognised with high accuracy). This helped filter out mis-transcribed or ambiguous passages (common in low-quality audio or with heavy background noise) and guaranteed that the textual data reliably reflected the speakers' intended statements.

With the filtering complete, the remaining feedback from customers, firms, and financial institutions was input into a chosen LLM (OpenAI's GPT-4o) to forecast the content of the finalised guidance. To validate this forecasting approach, we then conduct a similarity analysis comparing our forecasted document with the actual finalised guidance. Specifically, the performance of various

similarity testing methods is evaluated using customised LLMs, employing metrics such as Accuracy, Precision, F1 Score, Pearson Correlation, and Spearman Correlation to provide a comprehensive assessment of the effectiveness of each method in determining textual similarity.


The comparison methods used in our analysis include both traditional statistical approaches, such as Jaccard similarity, and advanced deep learning transformer models (BERT and RoBERTa). These methods have been benchmarked against each other to clearly demonstrate their respective performance. This benchmarking highlights the strengths of different analytical frameworks in capturing nuanced textual patterns. Section 3.5 provides a detailed demonstration of how early-stage signals - such as stakeholder consultation feedback, draft regulatory documents, and indicators of societal interest - are

systematically compared to the content of the final regulatory guidance. This approach allows us to identify which early signals correspond most closely to eventual regulatory content; for instance, RoBERTa's strong performance enables the detection of semantic alignments between preliminary signals and the final guidance. In our results, the model recognised an early emphasis on the principle of proportionality in consultation documents, mirroring how the FCA incorporates the principle of proportionality in the final guidance. Such findings suggest that RoBERTa can effectively anticipate the inclusion of key concepts from consultations in the ultimate regulatory outcomes.

To clarify our horizon scanning process, we structure our analysis into two stages: an exploratory stage, during which we gather diverse early-warning signals such as stakeholder consultation feedback, preliminary

regulatory drafts, and related news articles or reports that reflect broader societal interests; and an assessment stage, where these collected signals are analysed using RoBERTa and other language models to determine their semantic similarity to the final regulatory text.

High semantic alignments identified in the assessment stage indicate that specific concepts or terms from earlier discussions strongly predict the content of the finalised regulation, thus enabling effective forecasting of regulatory developments.



**Figure 3: Dynamic visualisation of high-frequency keywords in consultation stage**

### 3.3. Analysis

Figure 3 presents a ranking of public concerns derived from keyword frequency analysis within our feedback dataset collected during the FCA’s consultation phase on anti-greenwashing guidance. The most frequently

mentioned terms were associated with broad sustainability concepts, including “sustainability,” “eco,” “claim,” “investment,” and “product,” each appearing over 1,000 times. These frequently referenced terms likely reflect widespread public interest in ensuring

integrity and transparency in environmental claims and in the authenticity of sustainable investment offerings.

Following these prevalent themes, the analysis also identifies frequently mentioned keywords more explicitly connected to regulatory aspects, such as “disclosure,” “risk,” “policy,” “regulation,” and “law.” While the frequent occurrence of these terms initially suggests stakeholder attention towards regulatory mechanisms aimed at governing environmental claims and mitigating greenwashing, it is important to acknowledge that keyword frequency alone provides preliminary insights rather than definitive conclusions about stakeholder intentions. Nevertheless, the prominence of the keyword “disclosure” may preliminarily indicate stakeholder recognition of transparency as a potentially important mechanism for accountability and trust-building in sustainable investment products.

To explore specific stakeholder concerns in greater depth beyond initial keyword analysis, Table 1 summarises key concerns raised by firms and investment institutions regarding the FCA’s Anti-Greenwashing Rule, as generated from our use of OpenAI’s GPT-4o large language model. Note that we confine our scope to firms and investment institutions in order to conserve on space for this study. Table 1 explicitly outlines the main thematic areas of stakeholder commentary, clarifying both the nature and context of the identified concerns from firms. The main concerns highlighted include scope ambiguity, proportionality, disclosure requirements, and alignment with international standards that stakeholders believe need clearer definition and practical adjustment to facilitate effective compliance and market confidence.

**Table 1: LLM Analysis of Firm and Investment Institutions' Concerns on the Scope and Implementation of the FCA's Anti-Greenwashing Rule**

No.	Key Concern	Description	Analysis
1	Scope of Application	ICMA highlighted concerns regarding the ambiguous scope of the Anti-Greenwashing Rule (AGR) and its implications for the bond markets. They stated that "The scope of the AGR is not immediately clear and therefore any practical implications for the bond context are still to be fully determined." Additionally, ISDA pointed out that guidance primarily focused on retail could lead to uncertainties for wholesale applications, stating the need for clarity on the application of the rule to complex products such as derivatives.	Investment institutions are worried that the AGR's ambiguity could lead to misinterpretation, causing unintended compliance requirements or restrictions. Without a clear scope, firms may need to over-prepare or delay certain activities to avoid potential non-compliance. This results in inefficiencies and heightened caution within the market. Institutions are thus recommending more detailed definitions to prevent disruptions and ensure smooth operational continuity.
2	Application to Underwriting	ICMA expressed concerns that AGR could inadvertently encompass underwriting activities, which would complicate compliance with existing financial promotion regulations. They noted, "The wording of the AGR could bring within its scope a firm's underwriting and bookrunning activities." This risk was further underlined by the potential misalignment with current conduct rules for prospectuses and third-party communications.	The potential extension of AGR to underwriting might conflict with existing rules governing financial promotions and could disrupt established practices. Including underwriting under AGR might lead to conflicting regulations with the current financial promotion exemptions outlined in the FCA Handbook. This raises operational complexities and may reduce the willingness of firms to participate in primary market activities. Clarifying the AGR's non-application to underwriting is crucial to preserve current market structures.
3	Proportionality	EICMA and ISDA emphasised the importance of a proportionate approach in the application of AGR. ICMA stated, "A communication addressed to a professional client... may not need to include the same information." This reflects the belief that professional clients possess a higher level of expertise compared to retail clients and therefore require less detailed communication.	Applying the same level of detailed disclosures to both retail and professional clients could be redundant and ineffective. Proportionality ensures that communications are tailored to the audience's expertise. Investment institutions highlight this to avoid overloading professional clients with unnecessary details, which could slow down decision-making processes. This approach supports efficient and streamlined communication in financial markets.
4	Alignment with Existing Regulations	ISDA warned of potential legal and reputational risks stemming from unclear AGR guidelines, stating, "Otherwise it could expose firms... to significant risk of litigation and reputational risk.". This	The AGR should complement existing regulations rather than create overlapping or contradictory requirements. Investment institutions advocate for alignment to integrate the AGR seamlessly into current

No.	Key Concern	Description	Analysis
		sentiment was shared by ICMA, which highlighted concerns over unforeseen regulatory interpretations that could lead to litigation.	compliance frameworks. This prevents regulatory duplication, reducing administrative burdens and allowing firms to focus on actual sustainable practices instead of procedural compliance.
5	Litigation Risks	ISDA warned of potential legal and reputational risks stemming from unclear AGR guidelines, stating, "Otherwise it could expose firms... to significant risk of litigation and reputational risk." This sentiment was shared by ICMA, which highlighted concerns over unforeseen regulatory interpretations that could lead to litigation. Furthermore, ISDA emphasised potential liabilities linked to sustainability claims supported by third-party data, should these claims later prove inaccurate.	Ambiguity in guidelines can lead to disputes over sustainability claims, exposing firms to legal challenges. Clear regulations mitigate the risk of lawsuits or accusations of greenwashing. Firms need definitive standards to protect their reputation and ensure they can promote their products confidently without fear of misrepresentation claims. This provides a stable environment for sustainable investment.
6	Definitions and Consistency	ISDA and ICMA both stressed the importance of clear, consistent definitions of key terms such as "net zero" and "carbon neutral." ISDA stated, "Technical terms... should be clearly and consistently defined." This ensures that sustainability claims are uniformly understood and applied across the industry.	Inconsistent terminology creates confusion and complicates adherence to regulations. Defining key terms uniformly ensures all firms interpret and apply the rules in the same way. This consistency helps avoid misinterpretation, making it easier for firms to align their sustainability practices with regulatory expectations and ensuring fair market competition.
7	Retail vs. Wholesale Focus	ISDA noted that the AGR guidance seemed overly focused on retail markets, stating, "It would be helpful... for future guidance to include a wholesale focus." This highlights the need for tailored rules that reflect the differences in market participants' scale and operations.	Wholesale clients have different needs and levels of understanding compared to retail clients, and regulations should reflect this distinction. Guidance that is heavily retail-focused may not be applicable to larger, institutional investors. A dual focus ensures all participants can effectively comply without unnecessary adaptations or confusion, fostering a more inclusive and adaptable regulatory landscape.
8	Granularity of Disclosure	BlackRock and ICMA pointed out that overly detailed disclosure requirements could burden firms. BlackRock stated, "Given the granularity of the guidance in regard to fund disclosures, the FCA may want to consider..." Such granular requirements might overwhelm firms, complicating their communication efforts.	Excessively detailed disclosure requirements can burden firms and obscure key messages. Simplifying disclosure requirements without sacrificing transparency can help firms maintain clear communication with investors while avoiding operational inefficiencies. Overly complex requirements can deter clear understanding and reduce the effectiveness of disclosures.
9	Timeframe Compliance	BlackRock raised concerns about the tight implementation timeline, suggesting, "The FCA may want to consider... aligning the date of the guidance coming into	A tight implementation schedule can place undue stress on firms as they adapt their processes and training. Investment institutions often need sufficient time to revise their systems and ensure compliance.

No.	Key Concern	Description	Analysis
		force". ICMA also highlighted the need for reasonable timelines.	A more flexible timeline enables thorough preparation, mitigating rushed implementation that could result in errors or partial compliance.
10	International Consistency	ISDA emphasised that the AGR should align with global sustainability standards for consistency. They noted, "The importance for firms of international coherence and interoperability of rules." BlackRock echoed the need for harmonisation to prevent firms from facing conflicting regulations.	Divergent regulations between jurisdictions create a patchwork of compliance obligations that can be difficult for multinational firms. Harmonising rules with global standards allows firms to maintain streamlined operations across borders. This approach reduces costs associated with complying with multiple sets of rules and supports the global expansion of sustainable finance practices.

Through the systematic scraping and analysis of relevant online documents, our horizon scanning exercise reveals that institutional investors are particularly concerned about the implications of anti-greenwashing laws, with a notable emphasis on their impact on the bond market. As shown in Table 1, the International Capital Market Association (ICMA) expressed concerns regarding the ambiguous scope of the Anti-Greenwashing Rule and its potential implications for bond markets. Additionally, institutions such as BlackRock highlighted litigation and regulatory risks, while several entities noted that detailed disclosure

requirements could increase underwriting and compliance costs. Furthermore, given the proliferation of regulatory frameworks and guidance issued by various authorities, investment institutions highlight challenges related to the alignment of new anti-greenwashing laws with existing regulations, which may create additional compliance complexities. *Institutional* investors were also concerned about the tight timelines associated with consultation periods and the official implementation of anti-greenwashing laws and guidance. They emphasised the importance of having adequate time to prepare for



compliance with the new regulatory requirements. The accelerated pace at which guidance is coming into force poses challenges for these institutions, as they need sufficient time to adjust their policies, processes, and communication strategies to ensure alignment with the regulations.

Another key concern pertains to the granularity and level of detail in the guidance provided. Institutions were wary that overly detailed requirements may impose excessive burdens on firms, potentially overwhelming them and complicating their efforts to communicate effectively with stakeholders. This concern underscores the need for a balanced approach in regulatory design to ensure clarity and alignment while minimising unintended consequences that could hinder transparency and market efficiency.

---

<sup>4</sup> See [https://www.fca.org.uk/publication/finalised-](https://www.fca.org.uk/publication/finalised-guidance/fg24-3.pdf)

Building on the insights thus far, Table 2 presents a comparative analysis of actual feedback and recommendations from investment institutions, collected during the FCA's anti-greenwashing rule consultation, against the FCA's responses in its final rule.<sup>4</sup> This table draws on real consultation responses from firms and institutions, systematically compared with the final regulatory text using LLMs as a facilitating tool. The analysis, positioned in the application phase, is critical for evaluating how effectively the FCA addressed institutional concerns, identifying gaps that may influence future regulatory refinements. Through this LLM-facilitated comparison, several key observations emerge regarding the alignment (or lack thereof) between institutional priorities and the FCA's final anti-greenwashing framework. For brevity, we focus on a selection of these [guidance/fg24-3.pdf](https://www.fca.org.uk/publication/finalised-guidance/fg24-3.pdf).

observations, leaving the others for the reader to peruse.

First, the analysis points to the partial resolution of clarity and scope of application, a concern raised by investment institutions such as ICMA, who stressed the need for clearer boundaries, especially regarding underwriting activities. While the FCA's final guidance provides some general clarity, it does not fully address overlaps in underwriting regulations, leaving firms uncertain about how to proceed. This lack of specificity could lead to inefficiencies and hesitation, particularly in underwriting green or sustainability-linked products. To build confidence and ensure smooth operations, more detailed guidance is still needed.

Second, institutions like ICMA and ISDA highlighted the importance of tailoring communication requirements to suit different audiences, such as professional versus retail clients. The final guidance reflects this by

aligning with Conduct of Business Sourcebook (COBS) standards, which allow firms to provide detailed information to professional clients without burdening them with unnecessary simplicity. This balanced approach supports transparency while reducing unnecessary complexity, helping both firms and clients make informed decisions.

Third, litigation and legal risks remain a concern, particularly around the use of third-party data for sustainability claims. ISDA had recommended adding "safe harbor" provisions to protect firms from legal challenges if third-party data turned out to be inaccurate, but the FCA did not include this in the final guidance. The FCA stated only that if firms depend on third parties for information, they should assess whether it is appropriate to rely on the data, research, analysis, and other materials provided by those third parties to support the claims they make. Without legal protections, however, firms might hesitate to use external

data, limiting collaboration and transparency in sustainability reporting. Addressing these risks with clearer safeguards could encourage firms to engage more openly in sustainable finance efforts without fear of litigation. Lastly, institutions like BlackRock had emphasised the need for gradual timelines to avoid overwhelming firms. The FCA’s staggered

deadlines, particularly around the Sustainability Disclosure Requirements (SDR) and investment labels regime, has allowed for phased adaptation. This approach gives firms the time they need to update processes and systems, ensuring compliance without causing unnecessary disruptions.

**Table 2: LLM Analysis of Investment Institutions’ Observations and the FCA’s Responses in its Final Anti-Greenwashing Guidance**

Concern Category	Investment Institutions’ Key Concerns	Presence in Final FCA Guidance	Analysis
1 Clarity and Scope of Application	ICMA expressed concerns about ambiguous rule application, particularly regarding underwriting activities and the potential overlap with existing regulations.	Partially addressed. The FCA's final guidance maintains general clarity but does not explicitly outline underwriting distinctions.	ICMA highlighted the need for clearly defined rule boundaries to prevent unintended regulatory burdens, especially in underwriting activities. This aspect remains partially addressed, with general clarity in the final FCA guidance but lacking explicit distinctions for underwriting. ICMA’s concerns stem from the multi-layered nature of underwriting, involving facilitating, promoting, and executing transactions on behalf of third-party entities. If anti-greenwashing rules are not clearly defined, firms could be forced into comprehensive due diligence on external claims, leading to extensive resource allocation. This creates a compliance framework where responsibilities for verifying third-party issuer claims become ambiguous. Without precise exclusions, the potential for operational uncertainty and legal exposure rises. Firms could face

Concern Category	Investment Institutions' Key Concerns	Presence in Final FCA Guidance	Analysis	
2	Proportional Approach	ICMA requested a proportionate approach, allowing for differentiation between communications targeting professional vs. retail clients.	Addressed. The guidance references the need for proportionate communication, aligned with existing COBS guidance.	<p>duplicative checks and delays, extending transaction timelines and creating inefficiencies. This ambiguity may discourage participation in underwriting, especially for green or sustainability-linked products, affecting sustainable finance growth. Regulatory overlap can deter firms due to the perceived increased exposure to liability and procedural complexity, leading to reduced market activity and innovation.</p> <p>ICMA advocated for communication tailored to client types, ensuring proportionality when addressing professional versus retail clients. The final guidance references proportionate communication aligned with COBS standards. The guidance's acknowledgment of varied communication strategies allows institutions to present detailed, nuanced information to professional clients, avoiding oversimplification. This helps maintain effective dialogue while fulfilling compliance. Such targeted communication balances the need for transparency with practical disclosure requirements, benefiting institutions by reducing complexity. Professional clients gain access to in-depth data, enhancing informed decision-making, while simplified versions ensure broader retail compliance.</p>
3	Litigation and Legal Risks	ISDA highlighted potential liabilities if sustainability claims based on third-party data were found inaccurate.	Mentioned but not deeply expanded. The final document indicates firms must substantiate claims but does not add new liability protections.	ISDA emphasised potential liabilities linked to sustainability claims supported by third-party data, should these claims later prove inaccurate. The final guidance requires substantiation but lacks new liability protections. Firms must substantiate claims independently, raising the stakes if third-party data is later discredited. The absence of protective measures increases potential litigation, making firms cautious in using such data. While maintaining strict evidentiary standards ensures robust sustainability claims, the lack of liability safeguards can elevate legal risk. Firms may limit reliance on external data,

Concern Category	Investment Institutions' Key Concerns	Presence in Final FCA Guidance	Analysis	
4	Granularity of Fund Disclosures	Investment institutions like BlackRock argued that detailed fund disclosures needed phased implementation.	Addressed with modifications. The FCA revised disclosure timelines to phase in gradually, supporting operational readiness.	<p>reducing transparency and comprehensive disclosures.</p> <p>BlackRock and others noted that detailed disclosures need phased timelines to prevent operational strain. The FCA's revised timelines include phased disclosure requirements. A phased approach recognises the significant time and resources needed for data collection and system upgrades. This adaptation period supports structured integration, ensuring accurate and timely reporting. Phased implementation aligns with feedback, allowing firms to progressively adapt, mitigating disruptions. This structured approach facilitates compliance without overloading resources and supports a gradual transition.</p>
5	Global Alignment and Standardization	ISDA and ICMA emphasised aligning UK rules with global standards to prevent jurisdictional fragmentation.	Partially incorporated. The FCA's focus remains on UK-specific standards with ongoing consideration for international developments.	<p>ISDA and ICMA called for global alignment to prevent fragmentation across jurisdictions. The FCA's guidance mentions international coherence but does not fully align with global standards. The recognition of global frameworks indicates awareness, but the absence of full harmonization could create disparities for UK-based firms needing dual compliance. Without uniform alignment, UK firms may face challenges adapting to differing standards, potentially impacting competitiveness. Navigating divergent regulations adds to administrative burdens and strategic complexity.</p>
6	Use of Third-Party Verification	Concerns were raised about liability when using third-party verified data.	Not comprehensively addressed. The final guidance reiterates that firms must verify claims independently.	<p>Institutions called for "safe harbor" clauses, which the FCA did not include. The absence of such provisions may increase litigation risks related to external data dependencies. Concerns were raised about liability when relying on third-party verified data. Institutions sought 'safe harbor' clauses to mitigate risks. The FCA maintained independent verification requirements without such clauses. Firms using third-party data remain liable if the data is later challenged, which could deter reliance on external verification due to legal risks. The lack of 'safe harbor' provisions may discourage firms from engaging external verification, impacting</p>

Concern Category	Investment Institutions' Key Concerns	Presence in Final FCA Guidance	Analysis
			data reliability and collaborative transparency. The resulting hesitancy may undermine comprehensive sustainability disclosures.
7 Implementation Timelines	Calls for phased or delayed timelines to allow for operational adaptation.	Addressed. The FCA incorporated phased dates for implementing new requirements.	Phased timelines help institutions manage transitions. Aligns with institutional feedback, easing the shift to new practices and reducing compliance pressure. The institution call for phased timelines were aimed at accommodating compliance preparation. The FCA incorporated phased dates for implementing the new requirements. Phased timelines provide firms with the needed period to update processes, implement system changes, and ensure staff readiness, minimising disruption. The phased approach supports efficient adaptation and aligns with institutional needs, fostering compliance readiness while reducing transitional difficulties.
8 Greenwashing Definition and Examples	Institutions requested more practical examples to clarify what would qualify as misleading.	Partially addressed. The final guidance provides examples but not exhaustive lists.	Institutions requested detailed examples to clarify misleading claims. The FCA offered examples but not an exhaustive list. Examples help in interpreting the guidance but may not cover every scenario, requiring firms to exercise caution to avoid non-compliance. Limited examples can lead to varied interpretations and implementation inconsistencies. Comprehensive guidance would provide firms with clearer parameters, reducing ambiguity and enhancing compliance efforts.

### 3.4 Formal Similarity Testing of LLM-Based Anti-Greenwashing Rule Prediction

Understanding how early-stage regulatory discussions correspond to finalised guidance is crucial for anticipating policy shifts and helping

organisations prepare effectively. During the consultation phase, a broad range of stakeholders – particularly from the financial sector – submitted feedback outlining their expectations, concerns, and suggestions for

improvement. Yet, there is often limited clarity on the extent to which this input ultimately shapes the outcome. To bridge this gap, we apply similarity testing as a core methodological step, measuring how closely stakeholder feedback aligns with the content of the final guidance.

Similarity testing allows us to systematically assess the degree of overlap between consultation inputs and regulatory outcomes, whether in terms of specific themes, language, or proposals. By combining traditional methods such as Jaccard similarity and TF-IDF (Term Frequency-Inverse Document Frequency) with more advanced models like BERT and RoBERTa, we capture both surface-level patterns and deeper semantic connections. The results help reveal which stakeholder contributions were considered and highlight areas where their concerns may not have been fully addressed.

Building on this analysis, we again use OpenAI's GPT-4o to forecast the likely content of the final guidance. This step demonstrates how horizon-scanning techniques, when combined with LLMs and insights from consultation processes, can be used not only to assess alignment retrospectively but also to anticipate regulatory direction. In doing so, this dual approach offers organisations practical foresight, enabling them to prepare for changes in regulation, strengthen compliance systems, and reduce associated risks well in advance. To understand how closely consultation documents align with the final guidance, we use a variety of methods to compare the similarity between texts. These include both traditional techniques – such as comparing word overlap or word frequency - and more advanced approaches that analyse the meaning of whole sentences or documents. For example, we use methods like Jaccard similarity and TF-IDF, which focus on

surface-level word usage, alongside modern tools such as Doc2Vec and the Universal Sentence Encoder, which can capture deeper semantic meaning. We also include powerful transformer language models like BERT and RoBERTa, which are trained to understand the context of words in a sentence. To evaluate how well each method works, we apply several performance measures – such as accuracy and correlation scores – that help us determine how effectively each approach identifies meaningful similarities. This comparison allows us to find out which techniques are most reliable for analysing how early-stage consultation documents may relate to the final regulatory text.

The results are highly promising, with several advanced techniques demonstrating very high

similarity scores. For instance, transformer-based models, such as BERT and RoBERTa revealed similarity scores exceeding 0.8. Other techniques also delivered solid results, with similarity scores ranging between 0.6 and 0.8 (Table 3). These findings indicate that LLMs, when combined with horizon-scanning processes, can reliably forecast the likely content of finalised guidance. Specifically, the high scores mean these models successfully matched patterns and ideas from consultation feedback with what appeared in the FCA’s final rules. By systematically scanning early inputs like consultation documents and using LLMs to analyse them, we could predict key elements of the guidance before being finalised



**Table 3: Similarity Test Performance Comparison Metrics Based on LLMs**

	Accuracy	Precision	F1 Score	Pearson	Spearman
Jaccard	0.6234	0.5987	0.6102	0.5876	0.5923
BERT	0.8567	0.8489	0.8523	0.8678	0.8589
N-gram	0.6543	0.6387	0.6463	0.6234	0.6312
Doc2Vec	0.7234	0.7123	0.7178	0.7045	0.7123
RoBERTa	0.8678	0.8567	0.8621	0.8745	0.8678

**Notes:** The table compares five similarity testing methods applied to the provided text, highlighting their purposes and how to interpret their results. Jaccard similarity measures overlap between word sets, focusing on exact matches, as reflected in its lower scores, N-gram similarity extends this by considering word sequences, offering moderate results (e.g., Precision = 0.6387). Doc2Vec generates vectorised document representations, reflecting thematic alignment with mid-range scores (e.g., Pearson = 0.7045). Deep learning models like BERT and RoBERTa excel in capturing contextual and semantic relationships (e.g., RoBERTa Spearman = 0.8678). Metrics such as Accuracy, Precision, F1 Score, and Pearson/Spearman correlations assess different aspects of similarity, from overall correctness to rank-order alignment.

Expanding on Table 2 further, the Jaccard similarity test measures how much overlap exists between the exact words used in the two texts. The recorded score of 0.6234 for accuracy is a relatively low score and highlights that the FCA’s guidance did not use identical language to the recommendations made by institutions. For example, ICMA's call for clarity on underwriting was not directly mirrored in the guidance, leaving room for interpretation. Jaccard is useful for identifying literal overlaps

but does not account for broader context or meaning, which is why its results alone may not fully capture whether the guidance addressed key issues effectively.

On the other hand, BERT and RoBERTa, which achieved high F1 scores of 0.8523 and 0.8621 respectively, go beyond surface-level word matching by understanding the context and relationships between words. This explains why RoBERTa, with its refined training method,

slightly outperformed BERT across all metrics, such as Pearson (0.8745) and Spearman (0.8678). These results show that the FCA's final guidance reflects many of the deeper concerns raised by institutions, even if the exact wording differs. For instance, the FCA's adoption of a proportional approach aligns with ICMA's recommendations, which ROBERTa is better at capturing. These similarity tests demonstrate that while the guidance might not address every point explicitly, it aligns semantically with much of the feedback.

#### **IV. CONCLUSION**

This white paper aims to introduce innovative techniques, specifically Generative AI and LLMs, to enhance regulatory horizon scanning. We use the 2024 Financial Conduct Authority (FCA) anti-greenwashing regulations as a case study.

To achieve this objective, we first conduct a literature review and outline the traditional horizon scanning process, which consists of four stages: Exploration, Assessment, Application, and Continue. We then develop a framework that integrates Generative AI and LLMs into the horizon scanning process, with a particular focus on the Assessment and Application stages, where these advanced techniques enhance the collection, analysis, and interpretation of diverse information sources and unstructured text data. Specifically, in the context of the FCA's anti-greenwashing regulation, we sought to capture a broad base of relevant news, articles, feedback and comments during the formal consultation process. By leveraging the predictive capabilities of Generative AI and LLMs, we demonstrate how these techniques can anticipate future regulatory changes with high accuracy, offering a more proactive and

data-driven approach to regulatory preparedness.

Future research will focus on expanding the scope of this approach by extending the timeframe and incorporating the full extent of the unstructured data collected (i.e. text, audio, video). By incorporating a broader dataset over longer time horizons, we aim to check the robustness and strategic utility of the LLM-based approach horizon scanning for regulatory compliance and risk management. In particular, we are keen to discern whether

more informal sources of information (e.g. blogs) can add to the predictive capability we coupled with formal sources of information (e.g. formal consultation submissions) or whether such informal information merely creates noise. The results thus far suggest that the use of LLMs may allow firms to anticipate and adapt to regulatory developments with greater precision and confidence, while and policymakers may benefit through enhanced insights that would allow for more proactive regulatory design practice .

## REFERENCE

- Cuhls, K., van der Giessen, A., & Toivanen, H. (2015). *Models of Horizon Scanning: How to integrate Horizon Scanning into European Research and Innovation Policies*. European Commission. Retrieved from <https://d1wqtxts1xzle7.cloudfront.net/113228227/Models-of-Horizon-Scanning-libre.pdf>
- Delaney, K. (2014). Innovation Tool Kit. A practical guide: introduction to horizon scanning in the public sector. APO. [http://apo.org.au/files/Resource\\_publicsectorinnovationtoolkithorizonsscanningmodule2014.pdf](http://apo.org.au/files/Resource_publicsectorinnovationtoolkithorizonsscanningmodule2014.pdf), 25735152(0).
- Hao, Z., Bowden, J., and Cummins, M. (2025). Generative AI for simplified ESG reporting in financial services. *Financial Regulation Innovation Lab White Paper Series*. Available at <https://www.fintechscotland.com/wp-content/uploads/2025/04/Generative-AI-for-Simplified-ESG-Reporting-in-Financial-Services.pdf>.
- Idoko, O., & MacKay, R. B. (2021). The performativity of strategic foresight tools: Horizon scanning as an activation device in strategy formation within a UK financial institution. *Technological Forecasting and Social Change*, 162, 120389.
- Kim, H., Ahn, S. J., & Jung, W. S. (2019). Horizon scanning in policy research database with a probabilistic topic model. *Technological Forecasting and Social Change*, 146, 588-594.
- Marsh, A., Bhagat, A., Muller, A., Clarkson, H., Knowles, H., Park, J., & Goodman, J. (2014). Horizon scanning to improve social impact. *The Rockefeller Foundation*.
- Miles, I., & Saritas, O. (2012). The depth of the horizon: searching, scanning and widening horizons. *Foresight*, 14(6), 530-545.
- Palomino, M. A., Bardsley, S., Bown, K., De Lurio, J., Ellwood, P., Holland-Smith, D., ... & Owen, R. (2012). Web-based horizon scanning: concepts and practice. *Foresight*, 14(5), 355-373.
- Rowe, E., Wright, G., & Derbyshire, J. (2017). Enhancing horizon scanning by utilizing pre-developed scenarios: Analysis of current practice and specification of a process improvement to aid the identification of important 'weak signals'. *Technological Forecasting and Social Change*, 125, 224-235.

Sutherland, W. J., Broad, S., Butchart, S. H., Clarke, S. J., Collins, A. M., Dicks, L. V., ... & Ockendon, N. (2019). A horizon scan of emerging issues for global conservation in 2019. *Trends in ecology & evolution*, *34*(1), 83-94.

Wintle, B. C., Kennicutt, M. C., Sutherland, W. J., NM, B. P., Nancy, O., & Nathalie, P. (2020). Scanning horizons in research, policy and practice. *Conservation research, policy and practice*, 29-47.

## ABOUT THE AUTHORS



**Dr. Hao Zhang** is a Research Associate at the Financial Regulation Innovation Lab (FRIL), University of Strathclyde. He holds a PhD in Finance from the University of Glasgow, Adam Smith Business School. Hao held the position of Senior Project Manager at the Information Center of the Ministry of Industry and Information

Technology (MIIT) of the People's Republic of China. His recent research has focused on asset pricing, risk management, financial derivatives, intersection of technology and data science.

Email: [hao.zhang@strath.ac.uk](mailto:hao.zhang@strath.ac.uk)



**Daniel Dao** is a Research Associate at the Oxford Sustainable Finance Group, University of Oxford. He previously worked at the Financial Regulation Innovation Lab (FRIL), University of Strathclyde; and a Research Economist (Consultant) at the International Bank for Reconstruction and Development

(IBRD), The World Bank, Washington DC Headquarters. He is a CFA Charterholder and an active member of the CFA UK. He has earned his PhD (2024) in Finance at Coventry University, and MBA (2017) in Finance at Bangor University, UK. His expertise lies in the field of Fintech; Sustainable Finance; and Entrepreneurial Finance. He has published in high-ranking academic journals, including Risk Analysis (ABS-4); British Journal of Management (ABS-4), Information and Management (ABS-3), International Journal of Finance and Economics (ABS-3) and various policy and industry research reports affiliated with The World Bank (Dominican Economic Memorandum, 2023; World Development Report, 2024; Labour and Policy Reform, 2024) and Fintech Scotland (White papers and Blog Posts in AI, Fintech, ESG, and Financial Regulation).

Email: [daniel.dao@strath.ac.uk](mailto:daniel.dao@strath.ac.uk)



**Dr. James Bowden** is Lecturer in Financial Technology at the Strathclyde Business School, University of Strathclyde, where he is the programme director of the MSc Financial Technology. Prior to this, he gained experience as a Knowledge Transfer Partnership (KTP) Associate at Bangor Business School, and he has previous industry experience within the global financial index team at FTSE Russell. Dr

Bowden's research focusses on different areas of financial technology (FinTech), and his published work involves the application of text analysis algorithms to financial disclosures, news reporting, and social media. More recently he has been working on projects incorporating audio analysis into existing financial text analysis models and investigating the use cases of satellite imagery for the purpose of corporate environmental monitoring. Dr Bowden has published in respected international journals, such as the European Journal of Finance, the Journal of Comparative Economics, and the Journal of International Financial Markets, Institutions and Money. He has also contributed chapters to books including "Disruptive Technology in Banking and Finance", published by Palgrave Macmillan. His commentary on financial events has previously been published in The Conversation UK, the World Economic Forum, MarketWatch and Business Insider, and he has appeared on international TV stations to discuss financial innovations such as non-fungible tokens (NFTs).

Email: [james.bowden@strath.ac.uk](mailto:james.bowden@strath.ac.uk)



**Professor Mark Cummins** is Professor of Financial Technology at the Strathclyde Business School, University of Strathclyde, where he leads the FinTech Cluster as part of the university's Technology and Innovation Zone leadership and connection into the Glasgow City Innovation District. As part

of this role, he is driving collaboration between the FinTech Cluster and the other strategic clusters

identified by the University of Strathclyde, in particular the Space, Quantum and Industrial Informatics Clusters. Professor Cummins is the lead investigator at the University of Strathclyde on the newly funded (via UK Government and Glasgow City Council) Financial Regulation Innovation Lab initiative, a novel industry project under the leadership of FinTech Scotland and in collaboration with the University of Glasgow. He previously held the posts of Professor of Finance at the Dublin City University (DCU) Business School and Director of the Irish Institute of Digital Business. Professor Cummins has research interests in the following areas: financial technology (FinTech), with particular interest in Explainable AI and Generative AI; quantitative finance; energy and commodity finance; sustainable finance; model risk management. Professor Cummins has over 50 publication outputs. He has published in leading international discipline journals such as: European Journal of Operational Research; Journal of Money, Credit and Banking; Journal of Banking and Finance; Journal of Financial Markets; Journal of Empirical Finance; and International Review of Financial Analysis. Professor Cummins is co-editor of the open access Palgrave title *Disrupting Finance: Fintech and Strategy in the 21st Century*. He is also co-author of the Wiley Finance title *Handbook of Multi-Commodity Markets and Products: Structuring, Trading and Risk Management*.

Email: [mark.cummins@strath.ac.uk](mailto:mark.cummins@strath.ac.uk)



Get in touch

FRIL@FinTechScotland.com

This is subject to the terms of the  
Creative Commons license.

A full copy of the license can be found at

<https://creativecommons.org/licenses/by/4.0/>