

Department for Science, Innovation and Technology

By email: evidence@officeforai.gov.uk

20 June 2023

In support of the Government's AI agenda: a short-lived opportunity to regulate the use of AI in corporate reporting

As noted in my online response submitted yesterday, I am writing with a proposal for how the UK could regulate the use of artificial intelligence (AI) systems, particularly large language model or similar AI systems (LLMs)¹, in corporate reporting. In this proposal, I first explain the case for regulation, in line with the principles set out in your policy paper [A pro-innovation approach to AI regulation](#); and second, I propose how existing regulation might simply and easily be adapted to support those principles. As far as I am aware from my discussions with my European and American colleagues, neither the EU nor the US has yet produced plans for regulation in the area of reporting, so we in the UK have the opportunity to take the lead if we act swiftly.

I write as one of the UK's leading experts in reporting, having worked in the field for more than 20 years; writing annual reports and advising UK plc, principally FTSE100 and 250 companies, on how to report well. I am the author of the standard work on the subject, [Trust me, I'm listed](#), published by the UK's Chartered Governance Institute. As such, I'm particularly well versed in the practical application of regulation, and how companies can best report to ensure they meet both the spirit and the letter of its many and varied requirements, and ultimately support a thriving, competitive and leading London equities market.

In my view, while AI systems that collect and manage data may have a place in reporting, the use of LLMs in their current form would undermine reporting's guiding principles of truthfulness, proportionality and accountability. This would in turn undermine confidence in our equities market, to the detriment of London as a financial centre.

Companies are generally not yet using LLMs in reporting, so we have a window of opportunity now to regulate its use, before it is even *in use*. I am therefore proposing some simple amendments to existing reporting regulation, set out below, that would put guardrails around the usage of AI in reporting before it becomes a problem. As already noted, my proposal sits firmly within the principles set out in *A pro-innovation approach to AI regulation*. I have sent it to the Financial Reporting Council (FRC), in its role as a reporting regulator, and to the Financial Conduct Authority. The FRC expressed interest but said AI regulation is a matter for Government, so it will be awaiting your guidance on the subject. I hope this proposal is useful to you in giving such guidance.

1. The case for regulation – safeguarding truth and accountability to support a thriving equities market

Truth and accountability are the bedrock of corporate reporting, and thus the bedrock of our system of capital markets. In this system, investors rely on the accuracy of the data and narrative information that companies produce, to help with assessing those companies as investment prospects. London's reputation as a global financial centre rests on the confidence that investors, regulators, governments and wider stakeholders have in this information.

¹ As defined in the UK Government's white paper, *A pro-innovation approach to AI regulation*, section 3.2.1.

Since LLMs went mainstream with ChatGPT only a few months ago, concerns around their truthfulness and accountability have exploded – with, for example, its ‘godfather’, Geoffrey Hinton, resigning from Google so that he could speak freely about its risks; and the chair of the US Federal Trade Commission, Lina Khan, speaking out on the necessity for AI to be regulated, to ensure markets remain open, fair and competitive. There may well be good uses for LLMs in some arenas, but currently their usage would be incompatible with the principles of reporting.

Government calls on regulators to take the lead, emphasising the danger of LLMs

Your white paper emphasises the values of transparency and accountability alongside innovation, if the UK is to lead the world in AI regulation. It states the importance of regulators leading in this field, ahead of any regulation – *empowering regulators to take a lead* – since *Initially we do not intend to introduce legislation. By rushing to legislate too early, we would risk placing undue burdens on businesses.*² Importantly, your white paper asks regulators to regulate the *use* of AI, rather than the technology itself: *Our framework is context-specific. We will not assign rules or risk levels to entire sectors or technologies. Instead, we will regulate based on the outcomes AI is likely to generate in particular applications.*³

It is in this context that I am proposing these amendments to reporting regulations, which I believe could easily be adopted by the FRC if supported by the Government. Indeed, I believe this proposal is a good example of the kind of innovative, fast action that you are asking of regulators, in light of the dangers of LLMs highlighted in the white paper: *LLMs also have limitations, for example, the models are not trained on a sense of truth, so they can reproduce inconsistent or false outputs that seem highly credible. Because they can be adapted to a wide variety of tasks downstream within an AI supply chain, any improvements or defects in a foundation model could quickly affect all adapted products.*⁴

We do not fully understand how LLMs work, so we cannot properly account for them

One of the key fundamentals of reporting is accountability, and the ability to analyse and verify sources of information, which ensures that companies provide decision-useful, accurate information to the market. With the advent of AI, companies need to be reminded of their responsibilities, no matter what sources they are using. In the case of LLMs, it is not possible at this stage to produce guidance on how best to use them in the reporting context, since no one, including in some cases their creators, fully understands how they work. And because their workings are not disclosed, these systems cannot be tested by anyone outside the company that created them either. Given LLMs’ known problems with accuracy, it would be irresponsible to allow their usage in reporting at this time.

Act now, while it’s easy to do so

It is essential that we act quickly to regulate the use of LLMs in reporting, for two practical reasons. First, AI is moving so quickly that any action must be swift; and second, we have a window of opportunity, since companies have not yet started to use it in this way. Acting now would make it very easy for companies to comply, since we would simply be requiring them not to change their existing practices.

² Ibid, page 3.

³ Ibid, page 25.

⁴ Ibid, page 58.

To highlight the key points:

- **LLMs frequently create false yet highly plausible narratives.** There are two types of information in reporting – data, and narrative. Data alone is not truth – it is simply a fact that is (or is not) true. The truth (or otherwise) about a company, as told through reporting and results announcements, is revealed in the interpretation of that data – the story that is told around it. We know that large language model AI systems invent ‘truth’, creating highly plausible but false narratives, with false attributions and sources. It’s essential, therefore, that we keep such systems out of creating reporting narratives.
- **AI itself is unregulated and changing rapidly – its use in reporting must have guardrails.** It is irresponsible to allow any system that is in a rapid state of flux, with consequences as yet unknown, to be involved in the important business of creating information on which people base their investment decisions.
- **In practice, there is no way for a director to be properly accountable for the information produced by LLMs.** Directors’ responsibility for the accuracy of reporting is a fundamental principle, which would not change were the reporting to be created by LLMs. However, in practice, the notion of directors’ accountability works because directors trust the people and processes within the company that help them produce accurate reporting. These people and processes are not infallible, but individuals can be held accountable for their actions, whereas a machine cannot. This is especially worrying when the machine in question is known to create plausible yet false narrative, with false attributions. Knowingly bringing such a system into reporting is irresponsible, and compromises the chain of accountability.
- **Everyone – including its creators – is calling for AI to be regulated; the UK has a unique opportunity to act in the case of corporate reporting.** All over the world, people are clamouring for AI to be regulated – even its architects. While governments grapple with the bigger question of its general use, the UK Government has taken the bold step, through the white paper, to propose that those who can act to regulate the use of AI in particular domains should do so. The proposal for regulation outlined below is, in my view, entirely in line with this approach, and would give UK regulators the power to take the lead among reporting regulators worldwide to restrict AI’s ability to put false but plausible information into the public domain.
- **We have a window of opportunity in which regulation will be easy to implement – if we act quickly.** New reporting regulation always needs to consider the practicalities, and whether additional burdens will be placed on companies. Right now, the practicalities are easily surmountable, because in general, companies aren’t using LLMs to produce reporting and results announcements. We therefore have a window of opportunity, before work begins on the next round of December year-end reporting, to introduce such safeguards in a way that places no extra burden on companies.

2. What might regulation look like? Some simple amendments to our existing framework

In line with the principles of the ‘Better Regulation Framework’, regulation needs to be practical, proportionate and enforceable. Prohibiting the use of AI systems in the creation of reporting altogether would be both disproportionate and unenforceable: besides which, some types of AI data-gathering systems may well prove useful to reporting professionals. However, while the implications of LLMs are still so little understood, prohibiting their use in creating narrative, and requiring the disclosure of any use of AI in any part of the information upon which that narrative relies, is something that could be both

tested and enforced. Furthermore, it would support the important reporting principles of transparency, consistency and accountability.

With that in mind, existing regulation, encapsulated in the UK Corporate Governance Code's requirement for directors to *present a fair, balanced and understandable assessment of the company's position and prospects*⁵, could be updated as follows:

- **Narrative reporting – whether in annual reporting or other statements to the market – must be written by human beings.** Large language model or similar AI systems (LLMs) must not be used in any form to create narrative reporting, including first drafts.
- **Any use of AI systems** in the gathering or preparation of any form of source material that is used in annual reporting or other statements to the market **must be disclosed in full.**
- **Accountability – the existing directors' 'fair, balanced and understandable' requirement in the annual report should be expanded,** under the fairness principle, to include a statement that:
 - Confirms that the narrative has been written by a human being, not by an LLM system
 - Discloses either:
 - that AI has not been used at all in the creation of source material, or
 - a list of information that has been generated by an AI system, and which AI system or systems have been used in each case
 - Confirms that the director is confident of the veracity of all the information that has been included, from whatever source.

It's important that this issue is the responsibility of the Executive Directors, particularly the Chief Executive, and not just the non-executives; since the practicalities of creating reports sit within the company, and any use of AI systems will be a management issue that permeates the organisation.

In conclusion – act now to support directors in carrying out their reporting obligations responsibly

The large-scale adoption of AI is coming fast, and it is untested and unregulated. The proposal for regulating the use of AI in corporate reporting, outlined above, could make a meaningful contribution to ensuring that new such technologies are a benefit not a threat, and thus support the UK Government's agenda of making this country a leader in AI regulation. If such regulation were introduced before companies start adopting AI systems wholesale, it will be easy for them to adhere to it. In this way, the UK would take a lead in safeguarding the truthfulness and reliability of information that companies provide to the market, and their accountability for it.

Thank you for considering my proposal, which I would be glad to discuss with you further. In the meantime, I am convening a series of focus groups on this issue with around 20 FTSE companies plus a few other interested parties, and would be happy to share the outputs with you in due course.

Claire Bodanis

Founder and Director, Falcon Windsor; Author of *Trust me, I'm listed*

LinkedIn: <https://www.linkedin.com/in/claire-bodanis-00a7831/>

W: www.falconwindsor.com; www.trustmeimlisted.com

⁵ [The UK Corporate Governance Code 2018](#), section 4, Audit, risk and internal controls, pages 10-12.