The Good, the Bad and the (tremendously) Ugly:
AI AND THE QUEST FOR BETTER REPORTING

Last week, one of my US reporting pals alerted me to a survey by Workiva, which found that "83% of [US] executives… say their company uses generative AI to augment business reporting staff" and "89% say their company will likely do so in the next five years". I find these statistics fascinating, because either the US is miles ahead of the UK in using AI in its reporting, or I've been phrasing my own questions to UK reporters somewhat differently. My research in the late summer of 2023 found that, in general, AI was little-used in reporting in this country. And I doubt this will have changed that much in the intervening months, given where we are in the reporting cycle. Are we then, as often noted, two countries divided by a common language? I don't have an answer yet, but I am investigating what lies behind those survey numbers.

By Claire Bodanis

But what's more interesting, and chimes with my own research, is the statement from the same survey that, despite this enthusiasm, "72%… say their company is limiting the use of generative AI", due to concerns over data security. This gives me hope that it's not too late to promote the responsible use of AI in reporting, not just in my patch in the UK and Europe but in the US too.

Before I describe my research, here's what prompted it. Last year, when generative AI and ChatGPT in particular exploded into our consciousness, I realised that it would raise serious issues if it were used to create corporate reporting. What appeared to be a genius tool that could write at a prompt would be like crack cocaine for corporate reporters who struggle to write, or were just in a hurry. And so I felt that regulation, or at least guidance, would be essential if large language model systems (LLMs) of the ChatGPT type were not to compromise the ultimate purpose of reporting – which is to build a relationship of trust with investors and other stakeholders through truthful, accurate, clear reporting that people believe because it tells an honest, engaging story.*

Essentially, the two most important concepts in the purpose of reporting – communicating truth and building relationships of trust – could be compromised by the indiscriminate use of LLMs to create narrative. Truth is at risk because LLMs are well known to create false yet highly plausible narratives. Relationships of trust are at risk because, in the words of a tech company executive: "If reporting is about giving insights into the minds of management and the board, how is that achieved by narrative being automated? It makes reporting pointless."

AI was not, therefore, something I felt should be allowed to just "happen" to reporting. Careful analysis was needed to ensure that any usage would support rather than compromise reporting's purpose. In the absence of action from UK regulators, who told me they were waiting for government to act, last summer I ran a research project in which 40 + UK corporates, investors and advisors, including 10% of the FTSE 100, shared their thoughts about the potential impact of AI on reporting. From that I developed guidance, launched in November, to get the ball rolling.

While the research focused on LLMs as the AI type most likely to be used in reporting, an important point raised was that it's essential to differentiate between types of AI, since they have different benefits and risks. Which brings us to the good, the bad and the (tremendously) ugly.

So far I've focused on the use of AI in *creating* reporting. Here, LLMs have clear potential to qualify as Bad AI – and without proper guardrails, Tremendously Ugly AI. Something no doubt the executives in that US survey had in mind when urging caution in its use.

But we mustn't forget the potential benefits of using AI in analysing reporting. And that's where I see great potential for "good" AI, particularly when it comes to analysing the growing datasets required for sustainability reporting. Such tools are already in use by some investors: my latest research project, in partnership with Imperial College London and ESG data specialist Insig AI, is investigating the opportunities for companies too.

Boiled down to its essentials, my view is this. Good AI does things that human beings cannot do – like the massive job of interrogating thousands of reports and data points in a matter of seconds. Bad AI does things that human beings ought to be doing for the long-term good of their brains and expertise. And when it's used for editing text that should only ever be written by a human – such as the analysis and opinion central to corporate reporting – that's when AI gets decidedly Ugly.

My latest research project came about because, to me, the regulatory response to these issues in both the UK and EU has been woefully inadequate so far. When the UK government called for evidence from regulators on AI recently, the Financial Reporting Council was not even asked to respond. I am reliably informed that the issue is not yet properly on EFRAG's agenda either (although I'm pleased to say that I've since been asked to talk to EFRAG about my work in AI and reporting!). So our research aims to give regulators, as well as companies, insights that will help them develop the guardrails we need to ensure that Good AI prevails.

Want to take part?
Contact me at claire@falconwindsor.com.

Stay up to date
AI campaign at Falcon Windsor:
https://www.falconwindsor.com/aicampaign

*My definition, supported by government representatives and, in my first research project, a multidisciplinary group of UK reporting professionals.

CLAIRE BODANIS
is a UK authority on reporting. In 2004, she founded the reporting advisory company Falcon Windsor. In 2021, the UK CGI published her guide to reporting, *Trust me, I'm listed – why the annual report matters and how to do it well.* Claire's goal today is to ensure AI supports, rather than undermines, the purpose of reporting.