

YOUR PRECOCIOUS INTERN

How to use
generative AI
responsibly
in corporate
reporting

Summary paper

Companies are increasingly using generative AI. The aim of this research with FTSE companies and investors was to understand how generative AI could, and should, be used in corporate reporting. This paper summarises our full report, which you can find [here](#).

Foreword by John Elkington	1
Introduction	2
About this research	3
Summary of research findings	4
How to use gen AI well – guiding principle: your precocious intern	5
How to use gen AI well – summary of practical recommendations	6
About Falcon Windsor	8
About Insig AI	8

INSIG
•AI

Falcon
Windsor

Foreword

I don't know whether you are using what this report labels the 'precocious interns' increasingly offered by AI, and if so, how, but I find that I am now calling on their services all the time. Firstly, as thought partners (e.g., ChatGPT) and, secondly, to develop visuals for articles and presentations (e.g., Artiphoria). I make no secret of the fact, but then again, I don't broadcast it because I don't want the wider world to think that I have been taken over by an alien intelligence.

Having spent almost 50 years trying to get business to open up to the wider world on issues around safety, health, environment, social impact and sustainability, all in the spirit of promoting greater transparency and accountability, I find this moment in the evolution of corporate disclosure and reporting particularly interesting.

It is interesting – in a troubling way – because new legal requirements like the EU Corporate Sustainability Reporting Directive (CSRD) signal an increasingly bureaucratic approach to the production of sustainability-related data and information. And it is also interesting because the rapid evolution and employment of AI suggests a vast array of potential applications both in producing data (the supply side) and processing, analysing and strategically using data (the demand side).

The ESG recession triggered by the Trump Administration, if we can call it that, will dent the sustainability information market for a while, but the underlying impetus is unstoppable. It will either be conscious, considered and increasingly well managed, or it will be chaotic and increasingly dangerous to major brands and the companies behind them.

Meanwhile, my impression is that we have made significant (if increasingly bureaucratic) progress on the information supply side – but, at the same time, worryingly little progress in terms of the use of AI and other techniques to address the demand side. In short, we are failing to produce the sort of market and wider system intelligence that leaders in the public and private sectors will increasingly need.

To understand where new technologies might fit in, nine years ago I began a series of visits to leading AI institutions, from DeepMind in London to Singularity University, X Development (formerly Google X) and HP in California, to see what they were up to in the sustainability space.

I discovered much younger, faster-thinking people who were generationally predisposed to think about how their tools and technologies might play into the market opportunity spaces sketched (even if often not significantly opened up) by the UN Sustainable Development Goals.

Back then, very few of these firms and institutions were dealing directly with the sustainability movement, though my sense is that that has changed significantly since then. In this context, *Your Precocious Intern* from the Falcon Windsor/Insig AI team is a timely, useful contribution to a critically important debate.

The survey sample size is impressive – some 60 people responded, including five institutional investors, one proxy agency and representatives from 40 companies of various sizes from a variety of sectors (including 20 FTSE 100s). The overall impression I get is that the use of AI is spreading fast, albeit from a very small base.

Its impact, for better and worse, needs to be addressed, sooner rather than later. And we need to tackle demand side issues, not just supply side ones.

I see this project as a dipstick test, a stepping stone toward much deeper engagement with the AI sector around these challenges and opportunities.

John Elkington

Co-founder of Environmental Data Services (ENDS), CounterCurrent, SustainAbility and Volans; author of *Green Swans: The Coming Boom in Regenerative Capitalism*, and *Tickling Sharks: How We Sold Business on Sustainability*.

Introduction

Generative AI will transform how we all do business – and for some that transformation has begun already. Even during this research project, the advances we’ve seen have been astonishing. Thus far, generative AI has not reached the creation and preparation of reporting in any meaningful way, but it will, and soon. Before this happens by accident, this paper proposes an approach and recommendations for bringing generative AI into reporting responsibly, by design.

The extra step needed to reap the benefits of investment in gen AI

As our research with 40 FTSE companies, and analysis of all FTSE 350 reports published 2020-24 showed, many companies are beginning to explore the use of generative AI (principally trials of Microsoft Copilot and the development of internal chatbots) in some form in the corporate information ecosystem. Yet so far, despite the considerable investment such systems require, where they are in use, few seem to be going the extra step and training people how to use these tools properly. And even fewer seem to be developing a robust approach for using them effectively in preparing and producing corporate reporting, an area ripe with potential if the risks are managed properly.

Investors want reassurance about accuracy, truthfulness and authenticity

Investors – reporting’s primary audience – are not blind to the risks, as our research with a number of institutional investors showed. However, they’re also keen to see companies taking advantage of the opportunities generative AI can bring in dealing with the vast quantities of information reporting must now contain. But, in that desire for efficiency, they want to be sure of two things:

first, that the information remains accurate; and second, that the opinions and decisions set out in reporting remain the preserve of the people responsible for them. And, at least in these early stages, they want companies to tell them how they’re using generative AI in their reporting.

According to other research,* companies are already investigating automation and the use of AI in the finance function in a measured, risk-focused way. It’s time for a similar approach in the critical area of narrative reporting.

Because, as these tools start gaining real traction in the workplace, with Copilot apps and company chatbots a standard feature of the business desktop, generative AI will come to reporting, planned or not, whether it’s in the sources on which reporting is based, or in the way people write and produce the documents.

A window of opportunity to get gen AI right and support the purpose of reporting

But, as our research showed, we’re not there yet, which means we have an opportunity to develop an approach that will satisfy investors, taking advantage of the benefits while managing the risks. An approach that supports rather than detracts from reporting’s critical purpose, 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.

Get the approach right, with appropriate guardrails and guidelines, and it has the potential to offer real benefits to the overstretched, over-stressed preparer, saving significant time and money. Get it wrong, and companies risk publishing inaccurate, untrustworthy reports for which ‘AI did it’ is no excuse when it comes to the (unchanged) responsibilities of the directors who signed them off.

After all, generative AI does not change reporting itself – it’s a tool to help us do it better.

Our proposed approach: ‘your precocious intern’

This paper offers one such approach. A way of bringing generative AI into the reporting process in a thoughtful, responsible manner that should help companies do reporting better, while protecting themselves, their investors and other stakeholders against the risks. It’s also one that should remain relevant as the technology evolves, and, while based on UK research, should nonetheless be relevant in other jurisdictions too.

We call our approach ‘your precocious intern’. Bright, capable, efficient, diligent – yet limited in experience and prone to overconfidence in their own abilities, so must never work unsupervised. Everything about how you use generative AI, from what you ask it to do for you, and how you ask it, to what you allow it to take part in, to how you view its output, should be considered with that role in mind.

A starting point for debate

Companies, investors, and everyone concerned with the truthfulness and accuracy of reporting: we hope you find this approach, and the practical recommendations set out in this paper, a useful starting point for developing a shared view of how generative AI should be used in reporting. And we’d welcome questions, challenge and debate so we can keep this agenda alive as generative AI and reporting evolve.

Claire Bodanis

Founder and Director, Falcon Windsor

On behalf of the research team from Falcon Windsor and Insig AI

*KPMG: AI and automation in financial reporting, December 2024

About this research

Background

This research project was born out of an initial study with focus groups of FTSE companies and investors about the potential use of generative AI in reporting, carried out in 2023 by Claire Bodanis, Founder and Director of Falcon Windsor.

With generative AI evolving so fast, and starting to be adopted by companies in the form of tools like Copilot and chatbots, it was clear it would soon have an impact on corporate reporting. And yet, despite the potential benefits and risks posed by generative AI to report preparers, it seemed that little was being said or done about it.

So in 2024, together with Diana Rose of tech company Insig AI, Claire launched a new, more in-depth study, based on a research plan developed with Imperial College London. Its aim was to find out how companies are using or planning to use generative AI, including in reporting, with a view to starting the debate about its usage, and to prompt action by proposing practical recommendations.

With thanks to research respondents, Imperial College London and the UK's Chartered Governance Institute

Our thanks to all the companies, investors and others who responded to and took part in this research. Special thanks to the team at ICL who created an initial research plan for us, and to the CGI, who were carrying out their own research with FTSE companies on AI governance at the same time as this project, and also took part in ours.

Quantitative analysis – all FTSE 350 publications since 2021

The quantitative analysis was carried out by Diana and the team at Insig AI in early January 2025 using their proprietary AI-enabled research platform, which contains a searchable database of corporate documentation. The research covered all FTSE 350 publications from calendar years 2021-2024 inclusive. They looked at two things:

- = General trends in discussion of AI over time, across types of documents and between sectors, alongside detailed contextual searches
- = Mentions of the use of generative AI in reporting itself, or in areas that could imply a potential impact on reporting.

Qualitative analysis – focus groups with FTSE 350 companies and institutional investors

From June to November 2024, Claire chaired a series of focus groups with FTSE companies and investors. In total, 60 people responded, including five institutional investors, one proxy agency and representatives from 40 companies of various sizes from a variety of sectors (including 20 FTSE 100s) that produce reporting. An observer from the UK's CGI sat in on some of the sessions. The aim of the company focus groups was to find out:

- = How generative AI is being used inside companies and its impact on reporting
- = Views on future use in reporting, the benefits and risks
- = Views on regulation and disclosure.

The aim of the investor groups was to sound out their views of the use of generative AI in reporting, given their importance as reporting's primary audience. A representative of a UK investor body joined one of the investor groups.

Quantitative research sample

- = 21,350 corporate documents
- = All FTSE 350 companies' annual reports
- = A range of other FTSE 350 publications
- = All documents published in PDF form on corporate websites in calendar years 2021-2024

Qualitative research sample

- = 60 people responded
- = 40 companies of varying size, from a variety of sectors, including 20 FTSE 100s
- = 5 institutional investors
- = 1 proxy agency

Meaning of 'generative AI'

Our research focused exclusively on generative AI, the 'black box' of AI – as popularised by ChatGPT, Microsoft Copilot, Perplexity, and others. Such tools are probabilistic, statistical systems that give you a *likely* answer (they are programmed to give answers likely to satisfy the prompter); they are not deterministic systems that look up and tell you what an answer *is* (they have no concept of truth).

Meaning of 'reporting'

This paper uses 'reporting' to mean any narrative statements to the market on which investors and other stakeholders rely to make decisions, such as the annual and related reports, results statements, trading updates and so on.

Summary of research findings

For details, see the [full report](#)

Quantitative – what annual reports told us

Note: we can't conclude that no companies used generative AI in their reporting, only that none said that they did. But, when taken with the qualitative results, it's reasonable to conclude that it's not having a big impact on FTSE 350 companies' reporting just yet.

- = **Considerable growth in FTSE 350 AI references in reporting, 2021-2024**
 - Number of companies mentioning AI more than doubled, from 156 in 2021 to 292 in 2024
 - Average number of mentions increased from two to nine
 - The term 'generative AI' appears for the first time in 2023
 - Minimal emergence of some language around AI governance (less than 3% of total mentions)
- = **In 2024, most companies talked about AI, although 14% of FTSE 100s did not**
 - In annual reports, 68% of FTSE 350 companies mentioned AI including 76% of the FTSE 100
 - Across corporate publications as a whole, this rose to 73%, and 86% of the FTSE 100

- = **Only two mentions of gen AI in relation to reporting – both relating to imagery**
 - FTSE 100 WPP, which has quite a sophisticated annual report, used generative AI to create its cover artwork to promote its own services
 - FTSE 250 Ruffer Investment Company, which has quite a basic annual report, used generative AI to create an image of an ugly duckling, with a jokey reference thanking AI for creating it
 - There were no references to the use of generative AI in relation to drafting or the reporting process
- = **A rapidly increasing number of Codes of Conduct include risk-based references to AI**
 - In 2024, three FTSE 350 companies updated their Codes to warn employees not to upload confidential, proprietary or personal data into generative AI platforms.
 - Since this is such a fast-moving area, we also ran a search on Codes published between January and March 2025, and found that this number had increased to 17

Qualitative – what report preparers told us

Note: we focused on the adoption of generative AI in the corporate information ecosystem and its likely impact on reporting, not on its broader use across business operations.

General usage of generative AI inside companies

- = Astonishing variation, driven by culture
- = Microsoft Copilot likely to have the most impact on reporting
- = Massive variety of approaches to AI governance and policies
- = Very little formal training
- = Vital importance of the human checker

Usage of generative AI within reporting now

- = A handful of formal projects, but experiments happening
- = Experiments mostly for editorial/administrative uses
- = Little or no governance over experiments
- = Where governance exists, it's most likely to prohibit the use of Copilot or chatbots in reporting
- = In the longer term – keen to use it as an assistant

Views on future use in reporting

- = Debate centred over use in drafting/editing, and how to ensure the opinions of management and the Board remained truly theirs

Benefits

- = Efficiency and time-saving
- = Useful for more structured information
- = Particularly helpful for less confident/able writers
- = Research tool to help 'read' increasing volumes of data/information

Risks

- = Accuracy and reliability of the output – although mitigated by the subject matter expert responsible for checking
- = Authenticity of the company/individual voice
- = Corporate reputation, and the perception of leadership who 'can't be bothered' with reporting if outsourced to AI
- = The value of the reporting process beyond the report itself, with senior management coming together to think, discuss and debate
- = A reduction in critical thinking

How to use gen AI well: Guiding principle – your precocious intern

Accountability is a fundamental principle of reporting: directors are accountable for it, and this accountability cascades down throughout the organisation to those involved in the process. Introducing generative AI does not change this principle, it's simply a tool to help us produce reporting more efficiently. And you cannot hold a tool to account for what it produces.

Guiding principle – your precocious intern

You can rely on your precocious (AI) intern to:

- = Be capable and quick, with relevant skills or knowledge
- = Work hard endlessly without fatigue
- = Learn on the job
- = Do a task well with very structured clear advice and guidelines

What to watch out for:

- = Unpredictable and can come up with seemingly random responses
- = Cannot be assumed to be accurate
- = Can mislead you since it sounds far more confident than its expertise merits
- = Needs constant checking and lots of feedback (i.e. prompts)

Never let it write the opinions of management and the Board.

Check absolutely everything it produces.

In short: never let it loose on its own.

Why the precocious intern?

Those responsible for reporting must retain ownership of both process and content, so that they can have confidence in the accuracy and truthfulness of the output. Generative AI tools, while highly proficient at many tasks related to reporting, have inherent limitations. They are probabilistic systems designed to give a likely answer; they are not search engines programmed to find an accurate answer, although the answers usually sound highly convincing. And they cannot engage, as people can, with management and the Board, to discover what's in their minds and ensure that reporting truthfully reflects their opinions.

Which leads us to our guiding principle of how to use generative AI well in reporting: 'your precocious intern'. Most people in business instinctively know what this means; the strengths and weaknesses this kind of person has; what one should and shouldn't expect of them.

Bright, capable, efficient, diligent – yet limited in experience and prone to overconfidence in their own abilities, so must never work unsupervised, or be given tasks that should be the preserve of senior minds.

Bringing your precocious intern into reporting

Everything, from what you ask generative AI to do for you, and how you ask it, to what you allow it to take part in, to how you view its output, should be considered with that role in mind. And the better you are at instructing and training your precocious intern, the better the output they will produce. But never forget that you are responsible for it: you must check everything and satisfy yourself that throughout the report, the data and information are accurate, and the story truthful.

Assumptions behind the recommendations

- = Our focus is on the use of generative AI in the corporate information ecosystem and not on its use in wider business operations.
- = Within that, we focused on the preparation and production of reporting itself, rather than the many processes within companies that may produce information that is ultimately relied on for reporting, although some thought is given to that in the guidelines and governance section.
- = Companies are using 'enterprise' versions of generative AI and not external chatbots (a big no-no). Enterprise versions have been approved by the company's internal technology and/or data team, leaving no risk of sensitive information leaving the closed environment of the business.
- = Our recommendations are not exhaustive and will evolve as companies start to use generative AI in earnest.

How to use gen AI well: Summary of practical recommendations

For details, see the [full report](#)

What to use it for

- = **Administrative** – summarising meetings and calls, creating transcripts, writing minutes, summarising notes: can save a tremendous amount of time and it's generally quite good at doing it, although bear in mind that senior people discussing sensitive matters may not want to have such meetings recorded digitally.
- = **Research**: but check any output for accuracy before relying on it.
- = **Drafting and editing material that is not opinion**: use for routine disclosure and narrative sections, but not for matters of opinion.
- = **Tidying up disclosure statements**: as long as the output is properly checked.
- = **Creating visuals**: thematic or abstract pictures as an alternative to stock photography; experimenting with data visualisation.
- = **Retouching**: photos can be uploaded to generative AI for retouching and airbrushing, but doing so should not fundamentally change the truthfulness of the image.
- = **Interim proof reading and checking**: but don't forget that the output must itself be checked.

Don't use an external chatbot!

Using an external chatbot for confidential information breaches the rules of confidentiality, because you are effectively releasing that information externally.

What not to use it for

- = **Writing the brief**: this should come through discussion and debate from the minds of those who are responsible for it and who own its story.
- = **Opinion pieces** – authored statements, narrative analysis, forward-looking statements: as with the brief.
- = **Photos of real people and things**: like words, photography must be accurate and tell a truthful story too.
- = **Poor quality data or other underlying information**: generative AI tools cannot distinguish quality so are best used on data that's reliable, with verifiable sources.
- = **Final proof reading and checking**: final reading and checking must be done by those responsible for it, who can judge whether or not it's correct.

Get the best out of it: training and practice

Using a generative AI tool is less intuitive than you might think, so the following is essential.

- = **Introduce a proper training programme** with modules for reporting/confidential information, require people involved in reporting to take part, and track participation.
- = **Learn – and practice!** – how to write good prompts, but never forget generative AI is trained to respond to a prompt, not to give a truthful answer.
- = **Become a better reader**, so you can properly judge what the output means and whether it will be understood by your audience.

Ensure it's used well: governance and guidelines

The overarching principle is to develop existing governance and guidelines to accommodate the use of generative AI.

- = **On the agenda**: get the use of generative AI in reporting on the Board's/Audit Committee's agenda.
- = **Reporting policy and training**: develop your policy to include the use of and training on generative AI, not forgetting understanding how it's being used across the organisation and how that might affect reporting.
- = **Communicate the policy** to the whole company.
- = **Document the use of generative AI** during the reporting cycle, and track adherence to the policy.

Explain how you've used it: what to say in your reporting

Audiences, including investors, want to know if generative AI is being used, so companies should explain their position.

- = **If you're not using it or only experimenting** – include a short, general statement explaining your position.
- = **If (or when) you are using it** – be specific: explain how you are using it. For example:
 - **Explain the policy for using generative AI in reporting**, and a statement that the report has been created in accordance with that policy
 - **State where it hasn't been used**, notably sections covering forward-looking information and matters of opinion, which should not be written by generative AI.

INVESTORS' VIEWS ON COMPANIES' USE OF GENERATIVE AI

- = Generative AI will be needed for dealing with the vast quantities of information required for reporting.
- = Keen interest in how companies are using generative AI in reporting, although mixed views on the usefulness of disclosure.
- = Debate centred around how to ensure that opinions and decisions remain the preserve of those responsible for them if generative AI were to be used.
- = Concerns over risks to the accuracy of information and the truthfulness of the story if generative AI were to be used.
- = Essential to keep the human in the loop.
- = Any use of generative AI in decision-making must be disclosed.

REGULATION – SOME GUIDANCE WOULD BE USEFUL

Regulators, including those involved in reporting, have all made statements about AI, but so far, its usage specifically in reporting hasn't had much airtime. Companies and investors alike said that they did not expect the use of generative AI in reporting to be regulated, and few would even want it.

Given how slowly regulation moves, and the speed at which technology is evolving, it's likely to be impractical, and reporting regulation is challenging enough already. However, given the far-reaching potential benefits and risks of generative AI, almost everyone said that some guidance from regulators would be useful. Many felt that it would be valuable even just to have a minimal reminder that its use does not change companies' and directors' existing responsibilities and duties, but could have a significant impact on how they discharge them. After all, such guidance gives those tasked with producing corporate reporting the support they need to make the case for change.

About Falcon Windsor

Founded in 2004 by Claire Bodanis, Falcon Windsor is a team of 30+ independent experts committed to helping companies small and large, private and listed, produce truthful, accurate, readable reports that their investors and other stakeholders believe because they tell an honest, engaging story.

We bring together critical thinkers, strategic planners, writers, designers, and production and project managers with impressive credentials in corporate reporting.

Between us, we've delivered hundreds of annual reports and thousands of other communications projects. Many of us have worked client-side too. Reporting's in our bones.

That's why we love sharing our expertise: through our book, webinars, conference appearances; and through working with regulators, and people from every aspect of company life.

In July 2019, the Chartered Governance Institute commissioned Claire to write a book on how to do corporate reporting well. With a foreword by Sir Donald Brydon, and contributions from experts across the reporting world, *Trust me, I'm listed – why the annual report matters and how to do it well*, was published (second edition) in October 2021.



Claire Bodanis
claire@falconwindsor.com
+44 7966 196808
falconwindsor.com

“I’m keen to try anything that will help companies report better, and, used well to support the purpose of reporting, generative AI could certainly do great things. But, if people use it unthinkingly, without proper training or guidelines, it could fatally undermine the accuracy and truthfulness of reporting. I hope our research and recommendations will prompt the necessary debate!”

About Insig AI

Founded in 2018 and listed on AIM, Insig AI is a UK-based technology company specialising in data infrastructure and AI-powered environmental, social and governance (ESG) research tools. We combine deep expertise in data engineering, machine learning, and sustainability to help organisations transform complex, unstructured information into decision-ready insight – with transparency and traceability built in.

Our solutions bring together two core capabilities:

- = Data solutions that ingest, structure and connect large-scale datasets – making them machine-readable, AI-ready, and enabling advanced analytics and efficient access to high-quality data.
- = ESG research tools that support intelligent search, benchmarking, and analysis of corporate reports – helping you extract insight from both the narrative and the numeric sides of ESG data.

We work with investors, corporates, and consultants to enable smarter, more responsible decision-making across finance, ESG, and beyond.



Diana Rose
diana.rose@insig.ai
+41 79 582 96 05
insig.ai

“Amid the hype around gen AI, we need to take a pragmatic approach to adopting it in business, and I hope these focused recommendations help put that into practice. It’s also clear that exploring AI’s role in reporting reflects broader questions around how to handle this technology, so that we harness its power without compromising trust.”

What's next? Please comment!

Generative AI is evolving fast, and reporting is not standing still either. We'd welcome questions, challenge and debate so we can keep this agenda alive, so please share this paper, comment on it, and contact us to develop the conversation.

AI DISCLAIMER

No AI tools of any kind were used in our research project, or in the creation and production of our full paper or this summary, with the exception of the quantitative analysis of FTSE 350 documentation, discussed in the appendices of the full paper.

INSIG
•AI

Falcon
Windsor

© Written, designed and produced by [Falcon Windsor](#)

May 2025