YOUR PRECOCIOUS INTERN

How to use generative AI responsibly in corporate reporting

INSIG
·AI

Falcon Windsor

Why this research?

Truth and accountability are the bedrock of corporate reporting. Early in 2023, when ChatGPT started to gain widespread attention, Claire Bodanis, founder and director of specialist reporting and advisory firm Falcon Windsor, first anticipated that the wholesale use of generative AI would raise issues for the accuracy of corporate reporting in general, and for directors' duties in particular in terms of ensuring that reporting is fair, balanced and understandable.

From initial research and guidance in 2023...

Claire began a campaign for the responsible use of AI in reporting, first by submitting a proposal to the UK Government consultation on AI regulation in June 2023. Then, with the input of 40+ FTSE, small-cap and private companies, advisors and investors, she turned that proposal into guidance, published on 22 November 2023: *A responsible approach to using AI in corporate reporting – Guidance for Boards and management on approach and disclosure.*

...to increasing adoption of gen AI inside companies and the need for more practical recommendations in 2024

That guidance paper was picked up by Diana Rose of tech company Insig AI. With generative AI quickly being adopted inside companies, and many changes coming to reporting, Claire and Diana realised that practical recommendations, based on more detailed, up-to-date research, would be helpful for everyone. So, based on a research plan developed with Imperial College London, the Falcon Windsor and Insig AI teams conducted a research project with FTSE companies and investors to look at how generative AI could be used responsibly in reporting.

We hope this paper will reassure audiences everywhere that, whatever generative AI brings, when it comes to reporting, adopting a responsible approach will mean reporting can continue to provide the truthfulness and accountability so essential to our global system of capital markets.

For more information on our AI campaign visit falconwindsor.com

The research team

From Falcon Windsor
Claire Bodanis *(overall project lead and author)*
Miti Adu-Ampoma
Simon Harper

From Insig AI
Diana Rose *(Insig AI team lead)*
Steve Cracknell
Katie Liddle
Claire Walls

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.

Contents

Meaning of 'generative AI'

Our research focused on generative AI – as popularised by ChatGPT, Microsoft Copilot, Perplexity, et al. Generative AI is often seen as a 'black box' because its workings are not easily traceable or able to be validated. We did not look at other types of AI, such as machine learning (e.g. large scale pattern recognition) or general AI (machines whose creators claim can 'think').

The crucial thing to understand about generative AI tools is that they 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). Unsurprisingly, perhaps, generative AI seems so far to have found most traction in the corporate world in use cases where there is no right or wrong answer (e.g. marketing copy) or where accuracy is easy to check (e.g. programming).

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.

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.

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

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

Recommendations in summary

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.

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.

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.

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.

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 the annual report

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.

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.

> "We've got a very strategic leader interested in experimenting with new technologies, so this is the time."
>
> FTSE 100 corporate data and AI lead

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.

The case for using generative AI in reporting

In this section:

Making a case for using generative AI in reporting may seem unnecessary – after all, if the technology is being adopted by companies, won't it just happen to reporting anyway? Our research suggests that it's only a matter of time before it will, but without proper thought and planning, it won't be done in the most responsible or effective way. Some companies are already investigating automation and the use of AI in the finance function in a measured, risk-focused way; rather than letting generative AI just 'happen' to narrative reporting, it's equally important that companies take a thoughtful, measured approach here too.

Given the importance of reporting as a legal document of record, and the fact that, for FTSE companies at least, producing the annual report tends to be a complex, nine-month (if not year-round) project, process change tends to come to reporting pretty slowly. And, as we know from our experience in this field over the last 25 years, people responsible for reporting tend to be relatively cautious, making sure that new technologies and process changes will genuinely be beneficial and avoid unintended consequences. This mindset is important, given the potential ramifications of introducing generative AI to reporting. It also means that, even in the rush to adopt the latest technology to support business processes, when it comes to reporting we still have time to think through its use properly, and get it right.

The investor perspective

This should be a comfort to investors who, as we know from our research, have a keen interest in how companies are using generative AI in meeting broader business objectives of productivity and efficiency, including in how they've produced their reporting. On this latter point, the general feeling was that its use is becoming essential for dealing with the vast quantities of information now related to reporting, but that it should be used thoughtfully to ensure that reporting remains accurate, and that the views and opinions of management and the Board are truly theirs, and not a product of an AI tool. Importantly, they want to know *how* generative AI is being used by companies, including whether it's being used to inform any kind of decision-making.

With this in mind, it's important that companies are thoughtful in how they use generative AI in reporting, and that they can explain how and why they are using it, what the benefits are, and how they are identifying and managing the risks.

In this section we set out the case for using generative AI in reporting, what the benefits and the risks are, and look at how it's being used in companies currently.

> "I would be very wary about [AI being used in] forward-looking statements, or anything that is based around an opinion or a judgement."
>
> Institutional investor

> "I'm not too worried if AI has been used in the process of collecting data or drafting the report. The only thing I care about is that it is accurate and that it does reflect management's views."
>
> Institutional investor

The appeal of generative AI for reporters

For anyone involved in the increasingly laborious process of preparing corporate reporting, the draw for using generative AI tools is obvious. The burden of reporting – the number of disclosures, their complexity, and the confusing nature of reporting requirements – has increased significantly in the last few years.

Taking number of sentences as a crude measure, since 2016, the average length of an annual report amongst the FTSE 350 has more than doubled. And, in the UK at least, unless the Government's long-promised overhaul of reporting comes to pass, reports are likely to get longer still with the incoming wave of sustainability reporting requirements.

More work, fewer resources

In the EU, the new European Sustainability Reporting Standards (ESRS) have created immense complexity and a considerable burden of both time and cost, and will continue to do so even as they are revised and slimmed down. Larger UK-listed companies are caught up in these too, and many other UK companies eventually will be affected by the new sustainability reporting requirements coming from the International Sustainability Standards Board, IFRS S1 and S2, which are likely to be endorsed for application in the UK, possibly as soon as from the financial year beginning 1 January 2026.

Yet we have no more time available to us to produce all this information – and quite often it's the same teams doing all the work. So, as suggested by all who took part in the research, anything that can help reduce this workload is going to get a serious look-in. It's no wonder that many are looking towards generative AI tools like Copilot and ChatGPT that can read documents and summarise them for you, summarise meetings and take notes – the administrative task list appears endless. Not to mention the promise that such tools can also do your analytical and research work, like turning raw information into useful datasets, analysing and drawing conclusions from data, providing a gap analysis; writing drafts and editing your work.

With all these potential uses, it's clear that generative AI, if used well, could make us more productive and efficient in producing corporate reporting. But in turning to it as a solution, it's essential that we don't introduce problems that end up making the reporting task more rather than less complicated. So where do we start?

> "The number one thing everyone in my company has said is that [Copilot is] super useful for transcribing things."
>
> FTSE 100 communications director

> "I see some real benefits in using it as a summarising tool, for instance where you've got big groups of data."
>
> FTSE 250 company secretary

Ensuring its use supports the purpose of reporting

To determine the best way of using generative AI in corporate reporting, we need to start with the purpose of reporting. Starting here helps us think through where and how generative AI could help us, and where it could potentially detract from our objectives.

If you ask people what the purpose of reporting is, many will say it's 'to meet regulatory requirements'. The annual report must indeed meet regulatory requirements; the report is a legal document of record for which companies and directors can be held to account, and through which they demonstrate how they have operated in accordance with company law.

But the more important question we need to ask to be able to report well is, what is the *purpose* of those regulatory requirements? Why are companies required to report *at all*? This takes us to the true heart of reporting, and its ultimate 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.

There are two concepts here that must be upheld for reporting to be authentic and useful. One: truth, in the context of truthful information and telling honest stories, and two: relationships, in the context of the annual report giving investors and other stakeholders insight into the minds of management and the Board – particularly those stakeholders who do not have personal, face-to-face interactions with those individuals. These concepts are at the heart of the UK Financial Reporting Council's requirement that annual reports be 'fair, balanced and understandable'.

> "We have to be a bit careful about demonising AI when as investors we've been naturally accepting humanly-curated opinion [from proxy agencies] that's sold back to us as fact."
>
> Institutional investor

Accurate data and a truthful story

The annual report serves these concepts of truth and relationships, and achieves its purpose, by providing two types of information, in a way that is clear and easy to find and understand:

- Accurate data and disclosures in accordance with reporting requirements
- A truthful story, namely the opinion of management and the Board as to the meaning of that data and those disclosures for the company and its future prospects.

Of course, not all reporting today provides these well, or serves that purpose. Some reports are not accurate; some reports do not give the true opinion of management and the Board; some reports, while technically doing both, render themselves virtually useless by presenting the information in a way that is impenetrable. But the aim of this paper is not to argue over whether today's reporting is on the whole good or bad, or indeed whether or not Boards and management on the whole tell the truth. Rather, based on the assumption that reporters *do* want to achieve this purpose, then how do we make sure introducing generative AI will help and not hinder us?

Mitigating the risks of using generative AI in reporting

The two principal concerns raised by both companies and investors about using generative AI in reporting were on precisely these issues: ensuring the accuracy of information given the well-publicised limitation of generative AI tools which are not programmed to give truthful answers, and the authenticity of the voice of management and the Board if generative AI were to write their views for them.

Of course, for this latter point, the counter argument is that, in many, perhaps most, cases, someone else already writes their views for them, so what difference does it make if generative AI were to do it? The difference, as many told us, lies in the process: the way that senior management comes together to think and debate around what the story really is and the way that senior executives and Board members are interviewed for their statements with their voices and sentiments being represented authentically. And, given that generative AI systems used by companies tend to be based on a limited number of foundational models, there is a real risk that all output, from all people in all companies, will end up sounding the same.

Ensuring its use supports the purpose of reporting continued

On this theme, both investors and companies were also concerned about the perception of leadership that 'can't be bothered' with reporting, and outsource this important work to generative AI.

A further concern of investors was how much easier generative AI would make it for companies to include tickbox buzzwords and phrases and thus 'game the system', now that more analysis of reporting, at least in the first instance, seems to be being done by machines. Some also noted concerns over the real origins and sources of information being used by generative AI, given that it is a 'black box'. This was a concern even when companies use their own internal chatbots working on their own internal information – after all, the foundational models on which such internal chatbots are based were created and trained on external information.

So, to retain investors' (and other stakeholders') confidence that the information within reporting remains accurate, and its narrative authentic, we must ensure that we mitigate those risks in the way we introduce generative AI into the reporting process, and the tasks we use it for.

Generative AI is already here: how do we use it well?

Although this sounds straightforward, in practice it's more complicated given that generative AI is already coming into the workplace and therefore will affect reporting, whether or not that's by design. So any recommendations for its use in reporting, as well as mitigating those risks, must start from how generative AI is already being adopted and used. In this way, they are most likely to be realistic and have the best chance of being useful.

It's important to note that our focus in developing this research was on the adoption and use of generative AI in the corporate information ecosystem, and its likely effect on the reporting process, rather than its broader adoption and use across business operations.

The overall message we heard about how generative AI is being used was: 'not very much – yet.' This was consistent across both our quantitative research into FTSE annual reports and our qualitative research amongst report preparers. While companies are investing in AI for pilots and targeted applications within their operations, few appear really to have invested in its use beyond this.

When it comes to the effect of generative AI tools on the corporate information ecosystem and narrative reporting in particular, usage of AI seems largely restricted to some trials of Copilot, with some companies doing more in the finance area. But, given the pace of change, it's likely that, by the time this paper is published, that will have moved on considerably, and it's likely that generative AI will become more and more embedded in corporate systems. It's important to note, however, that only a few (generally larger) companies mentioned corporate reporting as one of the focuses for AI usage.

> **Why annual reports *should* be biased**
>
> A few people commented that using generative AI to analyse the data and write the narrative might improve reporting because, in their view, it would be less biased. This misses the point about what the narrative is for – namely to give insight into the Board's and management's view of why the company has produced these results, and what they mean for the future of the business. The question in the reader's mind is: do I trust this group of people to run the company? In that sense, annual reports *should* be biased: rather then presenting some disinterested analysis of the data, they should present the (truthful) views of those leading the company – within the guardrails of an audited set of accounts, and the regulatory requirement for annual reports to be 'fair, balanced and understandable'.

WHO READS ANNUAL REPORTS – HUMANS OR MACHINES? AND HOW DO WE CATER FOR THEM?

There is a lot of debate over who reads reports today, and whether we should increasingly consider machines – notably generative AI – rather than humans as their primary audience. This research did not cover a detailed analysis of who or what reads reports today, but, given the purpose of reporting, we believe that reports can (and should) meet the needs of both.

After all, except in the specific context of large quantitative and systematic funds that use machine reading without human intervention (not always to a good end!), when reports are being read in the first instance by machine readers, those machines are producing output that will ultimately be used by humans to make decisions based on that output, including, for example, whether or not to read a report for themselves.

The key principle that serves for both type of reader is to use design and writing to distinguish between two types of information:

- Data/disclosure-type information should be structured and designed so that reports can easily be tagged for machine reading, and so that humans can easily find what they need.
- Story-type information, which is largely focused on the human reader, or the machine reader summarising for the human reader, should be very clearly written, engaging and accessible. Clarity is essential for both types of reader, and don't forget that machine readers will summarise what's there, rather than puzzling out what the author might have been trying to say.

Unlike rules for what must be disclosed, there are few rules about how to present this information clearly and well, and conventions of presentation that worked when reports were of an order of magnitude shorter do not always serve these longer reports. A principle increasingly being used in reporting is to separate out disclosure statements from narrative, and make the disclosure statements much more structured.

How it's being used today: summary of what reports told us

Given that AI is such a hot topic, a good starting point for finding out how companies are thinking about it is to read what they're saying in their annual reports. Reflecting the focus of our qualitative research, we focused on the FTSE 350 and analysed all annual reports published as PDFs on corporate websites in the calendar years 2021 to 2024. To test what we found in the annual reports, and to see where and how AI is being talked about more broadly, we also looked at other documents published by FTSE 350 companies in PDF form on their corporate websites, including sustainability and topical reports, policies and results statements, plus a handful of non-UK companies' reports. In total, we reviewed 21,350 documents.

In summary:

- **Considerable growth in AI references, 2021-2024.** Unsurprisingly, the number of FTSE 350 companies mentioning AI at all in their annual report more than doubled between 2021 and 2024, and the average number of mentions increased fivefold over that period. This is driven by general mentions of AI, with the term 'generative AI' appearing for the first time in 2023, along with the emergence of some language around AI governance (less than 3% of total mentions).
- **In 2024, most companies talked about AI.** If we look at 2024, 68% of FTSE 350 companies made some mention of AI in their annual reports (including 76% of the FTSE 100). If we look across corporate publications as a whole, that rose to 73% (and 86% of the FTSE 100). If we look at the number of AI mentions in 2024, 70% can be found in the annual reports, with the remainder scattered across other publications.
- **Two mentions of generative AI in relation to creating reporting imagery (both reports published in 2024).** WPP used AI to create its cover artwork as an example of its creative services and Ruffer Investment Company used it to create an ugly duckling image, with a jokey reference thanking AI for creating it. However, not a single report refers to generative AI in relation to the reporting process.*

We can't, of course, conclude that no company used AI in its reporting, only that none *said* that they did. And it is possible that we missed a reference – after all, no technology tool is perfect, not even ours! But, when taken with the results of the qualitative research, what we *can* say is that it seems unlikely that generative AI was used in any significant way in the creation or production of the documents we reviewed.

Evolution of average number of mentions of AI in FTSE 350 annual reports compared with all corporate documents

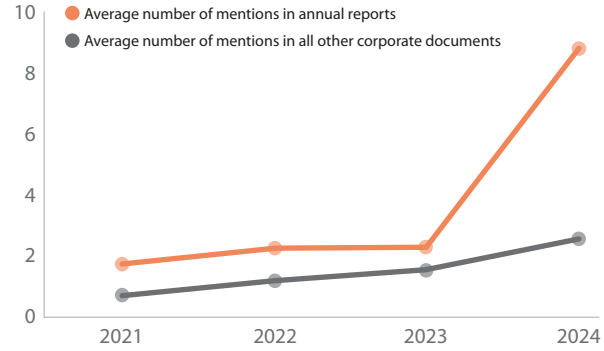

Evolution of number of FTSE 350 companies mentioning AI in their annual reports compared with all corporate documents

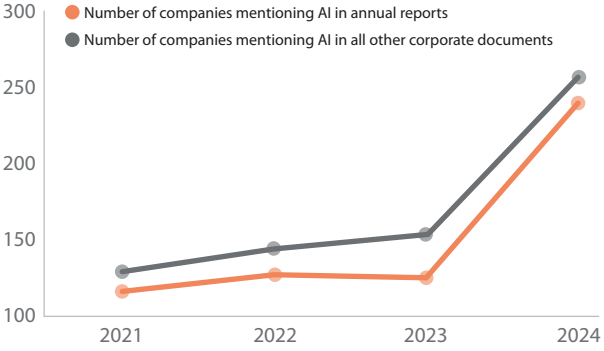

Evolution of AI mentions by type in FTSE 350 annual reports

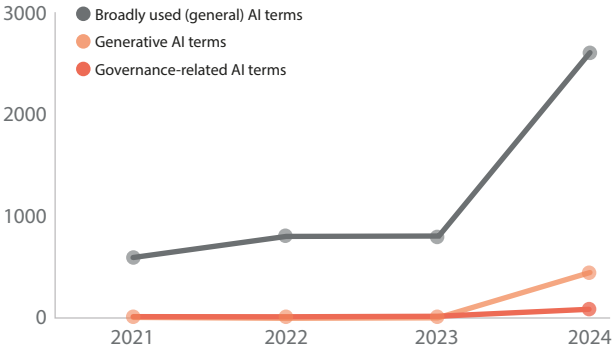

Anecdotally across the global dataset, we found that some companies disclosed the use of generative AI tools for 'content generation', but not specifically for their annual reports.

*Reports published in 2025 are out of the scope of our research period, but, since WPP published their 2024 report in March 2025, we checked to see if generative AI had been used again this year. As last year, it was used to create the cover artwork. Ruffer has a June year end so their next report had not been published at the time of publishing this paper.

How it's being used today: summary of what reports told us continued

Key themes emerging on companies' use of AI

While no company refers to the use of generative AI in the reporting process (aside from a couple of uses for imagery), it's instructive to look at how AI is discussed, since we may be able to infer something about how companies are incorporating AI into their systems and processes, which would ultimately have an impact on reporting. We noted three key themes.

- **Perceptions of AI: risk and strategy.** Mentions of AI across all documents broadly relate to both risk and strategy. The sentiment varies depending on sector. For example, financial services companies mention large language models (LLMs) in their risk assessments, not only in terms of cybersecurity, like all other sectors, but also in terms of concerns of a financial bubble, created by the excitement over the launch of ChatGPT specifically. Banks also mention AI risks specifically in enhancing ransomware and fraud attempts.
- **Efficiency and automation.** In terms of operational use cases, some companies seem to be starting to adopt LLMs in their processes. Within FTSE 350 annual reports published in 2024, 19 mentioned giving their employees access to and training on using LLMs to work more 'efficiently and effectively'. One financial services company stated that they had tried to 'automate the process of answering due diligence questionnaires using an internal large language model.' In all cases, detail was lacking over the extent of usage, how it was monitored and its success or otherwise.
- **References to AI in Codes of Conduct.** A few companies are putting in place policies to mitigate the risks of AI. In 2024, three FTSE 350 companies updated their Code of Conduct 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.

Why we can't draw any conclusions from the two UK examples of AI usage

Media company WPP used generative AI to create the front cover of its annual report published in 2024 (for the financial year 2023) for the purpose of marketing their in-house creative technologies. This was highlighted on their inside front cover, and it's worth quoting in full: 'Our cover artwork – reflecting the reshaping of the landscape by technology – was produced by our in-house creative technologists in collaboration with engineers at NVIDIA. The imagery is an evolution of the WPP brand identity, combining high-fidelity 3D models in NVIDIA Omniverse™ with generative AI using our proprietary AI production studio on WPP Open.' Meanwhile, Ruffer Investment Company, which produces quite a basic annual report, simply used an AI-generated representation of an ugly duckling to illustrate a message, with a jokey reference thanking AI for creating it.

For the purposes of this research, we wouldn't class either as an example of how generative AI may be being used in the reporting process more widely, given that they are highly individual instances and, in WPP's case, from a specific company that is developing AI tools itself as part of its business offering. It's also reasonable to assume that generative AI was not used in any other aspect of creating or producing either annual report, otherwise the companies would no doubt have said so. However, WPP in particular is certainly a good company to watch in terms of potentially being at the forefront of AI usage in reporting.

Quantitative research sample

See appendices for full findings and methodology

- 21,350 corporate documents in total
- All FTSE 350 companies' annual reports
- A range of other FTSE 350 corporate publications
- Anecdotal samples taken from a range of other UK and international listed companies
- All documents published in PDF form on corporate websites
- All documents published in calendar years 2021-2024
- Keyword and phrase search split into three groups:
  - Broadly used (general) AI terms (AI, artificial intelligence)
  - Generative AI terms (e.g. generative AI, gen AI, genAI, large language model, GPT, Copilot)
  - Governance-related AI terms (e.g. AI strategy, AI policy, AI governance, AI principles, responsible AI, AI working group, AI guidelines)

How it's being used today: summary of what report preparers told us

Turning to the qualitative research with FTSE companies, the picture is generally borne out in terms of overall usage, and lack of usage in reporting. Otherwise, the key finding was the astonishing variation across companies in terms of how generative AI is being adopted (or not) and the pace of adoption, although so far, not many companies are doing much with it within the reporting process.

The only determining factor as to the stage companies were at with generative AI seemed to be culture, particularly the personal interests and personalities of leadership and the Board, and a company's attitude to risk. There were no other factors in common such as size of company, industry and so on, although in general the largest companies seemed to have the most going on.

Following is a summary of what respondents told us about how generative AI is being adopted in general, and within the reporting process. For a fuller discussion of companies' views, see appendix 2.

General usage – growing use of Microsoft Copilot and the importance of the human checker

A few companies were not using or investigating generative AI at all, although most were doing something. This ranged from ad hoc usage of chatbots and testing of systems like Microsoft Copilot, to very tightly controlled and planned adoption.

Qualitative research sample

- 60 people responded
- 5 institutional investors
- 1 proxy agency
- Representatives from 40 companies, including 20 FTSE 100s
- 1 representative of a UK investor body
- 1 participant and 1 observer from the UK's Chartered Governance Institute
- 3 other governance/advisory people

Key themes were:

- **Microsoft Copilot featured strongly.** While some, more sophisticated, companies were building their own internal chatbots and AI assistants, most were incorporating generative AI through Microsoft Copilot – and it seems likely that this will have the most significant implications for reporting.
- **A massive variety of approaches to AI governance and policies.** In most cases, the approach was surprisingly lax, with no guidance at all, even for how to handle confidential material (although the early 2025 analysis of Codes of Conduct suggests this could be changing fast). At the other end of the spectrum, some companies had highly developed policies and guidelines in place; most companies were somewhere in between.
- **Generally very little formal training.** With a few notable exceptions, there was a distinct lack of training on how to use generative AI effectively, if at all. Aside from being an issue for report preparers themselves who might want to use it, this lack of training is also an issue for others within the company whom report preparers may rely on to provide information and source material.
- **The vital importance of the human checker.** This was the one consistent theme relating to generative AI usage – perhaps unsurprising for a group of people involved in reporting, whose instinct is to check everything, regardless of the authoritativeness of the source.

Usage within reporting – not much so far; but potential as editorial support within careful guidelines

Overall, generative AI tools, including Microsoft Copilot, were not being used much in reporting, although a few (generally larger) companies had formal projects in place, particularly in the finance and other technical teams. Those who were trying it out – generally in companies with little or no governance over its use – tended to be doing so for editorial purposes, while the legal teams in companies with tighter governance tended to prohibit the use of Copilot or chatbots in reporting.

Most people felt that, in the longer term, they would use generative AI tools like internal chatbots or Copilot to help them with tasks like drafting, editing and proof reading against a company's house style. Most, however, felt that such tools would never be more than an assistant, although a few envisaged a future in which the whole annual report would be created entirely by generative AI, with people responsible only for checking the output.

Some examples of what we heard

"I don't think AI will ever be able to do 'fair'. It will also never replace sitting down with the Chair and saying what's on your mind and what do you want to get across?"
FTSE 100 company secretary

"There's quite a lot in the tone and the style that you can pick up on which I think would get lost with AI, because it would inevitably result in sentences being edited rather than rewritten."
Institutional investor

"I would never consciously use it to write CEO reports or anything where it is a very personal statement which you have to be accountable for."
Institutional investor

"I think the auditing profession is at serious risk of having at least half their job done by AI."
FTSE 250 company secretary

"I could see a role for it in smoothing out sections written by different people, to get to something that is more user-friendly and consistent."
FTSE 250 company secretary

"We're looking at restructuring our whole data set to shift it into SharePoint, to make it more accessible to Copilot. It's a big piece of work and we are in relatively early days."
FTSE 250 company secretary

"Would I be able hand on heart say that none of my contributors had used gen AI to provide the bit they've sent in? I have no idea."
FTSE 100 company secretary

"It's quite depressing because my nine-year old son can use Copilot better than I can – but at least he'll be able to have a job in the new world!"
FTSE 250 company secretary

"There are very clear rules and guidance on how you should use the information. We go straight to the lawyers and say, hey, we'd like to use [AI] in our results, and they say no you can't!"
FTSE 100 investor relations director

"We have some self-directed coaching, but no formal training, and no policies around what you can and can't use [gen AI] for."
FTSE 250 company secretary

HOW REPORTERS FEEL ABOUT USING GENERATIVE AI TOOLS

Personal views about the use of generative AI, both in reporting and more generally in business and personal life, were very mixed.

- **It's inevitable.** Most participants accepted that generative AI will be part of the world of work, but felt it must be used sensibly and responsibly. Generally, those people also had some concerns about its widespread use and were keen for strong governance around it.
- **It seems risky.** A few were very worried about the risks, specifically in reporting but more generally as well.
- **It's got potential.** A few were very excited about its potential to transform the world of work, and confident that any risks could be mitigated.

CHARACTERISTICS OF THE QUALITATIVE RESEARCH SAMPLE AND HOW IT INFORMED OUR RECOMMENDATIONS

Our qualitative research was limited to a subset of reporters within the FTSE 100 and FTSE 250, although that subset covered a range in terms of size, and included companies from most sectors. It's important, then, not to assume that the description of how generative AI tools are being used currently in the corporate information ecosystem is necessarily representative of the FTSE as a whole, or indeed of smaller UK listed and non-listed companies.

Focus on high quality reporting

What united this subset was that it represented people who genuinely care about reporting, who uphold the purpose set out in this paper, and who believe reporting matters. Since the aim of this research is to help companies use generative AI tools to help them do reporting well and uphold its purpose, then we believe the views of this subset are both representative and valuable in informing our recommendations.

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

Reporting is ultimately owned by the directors, who are legally responsible for the accuracy of its content, and must sign that, in their view, the annual report is 'fair, balanced and understandable'. They must also explain how they've ensured that it is. Accountability is thus a fundamental principle of reporting: directors can be held to account for it, and this cascades down throughout the organisation to those involved in the process.

The introduction of generative AI does not change this principle, because (at least for now) it doesn't change reporting itself, it's simply a tool to help us create and produce reporting more efficiently. So those responsible are still responsible, whoever or whatever creates the information and narrative that makes up reporting. The difference of course is that one cannot hold a tool to account for what it produces.

So, to reap the productivity benefits of using generative AI to help us create reporting, we need to find a reliable way of doing it that will enable us to uphold the purpose of reporting and continue to have faith that our data and disclosures remain accurate, and our story truthful. And, we must be able to evolve how we use it and work with it as the technology itself evolves, while ensuring that those responsible retain ownership of the process and of the content, so that they can have confidence in the output for which they are accountable.

Our overarching principle for using generative AI in reporting is therefore that it should be treated as you would 'your precocious intern'.

> "I challenged Copilot to do a summary of our key achievements this year, and its top headline was that we had entered into a Copilot trial agreement, so it has got quite a high opinion of itself already!"
>
> FTSE 100 head of Group external reporting

Everything, 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. Why the precocious intern? Most people in business instinctively know what this means, what kind of strengths and weaknesses this kind of person has and what one should and shouldn't expect of them. In our view, these strengths and weaknesses are mirrored in generative AI, so it's a good model to have in mind when thinking about how to use it.

What is a 'precocious intern'?

Interns are generally bright, capable students who have relevant skills or knowledge in an area that means they can be useful and learn on the job. The best interns are keen and eager, working day and night during their internship because they want to do well and get a job as a result of their internship. They learn on the job and get better at responding to and understanding what you want them to do, the more they do it. With very structured, clear advice and guidelines, they can be set a task and, within the limitations of their knowledge and experience, do it very well indeed. And some simple tasks they can probably even do better than you can because they have the time and the focus to do them. But *everything* they do must be monitored and checked. It's impossible to know quite what they will come up with: depending on how they've understood an instruction, they might come up with something brilliant, or they might come up with something rubbish, having misinterpreted your intention.

Someone who is precocious – a term often used for children or young people – appears far more advanced than might be expected for their age, experience and skills. They also, however, tend to *think* of themselves as far more advanced than they really are, and so often sound far more confident and assertive than their expertise and knowledge merits.

Extrapolate these characteristics to a machine, and you have both the strengths and the weaknesses of generative AI as a tool to be used in corporate reporting.

What does 'your precocious intern' mean for using generative AI in reporting?

A precocious intern, then, can be very useful in reporting, as long as you have your eyes wide open to their limitations, particularly their appearance of certitude which can mislead you into being less careful in your checking than you ought to be. The better you are at instructing and training your precocious intern, the better the output they will produce.

But, unlike a human intern, don't make the mistake of thinking your generative AI tool could ever become more than that. However sophisticated generative AI becomes, however much it evolves in ways that we cannot imagine now, it should forever remain nothing more than a precocious intern – or at least for as long as human beings remain accountable for reporting. So don't forget: you must check everything and satisfy yourself that throughout the report, the data and information are accurate, and the story truthful.

Keeping the purpose of reporting and the principle of the precocious intern in mind, the rest of this section sets out how generative AI, in the forms we're seeing companies adopt it, could be used effectively in creating and producing reporting as it stands today. Our recommendations can't, of course, cover everything, and both generative AI and reporting itself are constantly evolving. But the tasks involved in creating and producing reporting don't themselves change that much, so these recommendations should be a useful starting point.

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.

Good with language; good with generative AI

The better you are with language, the easier you should find it to use generative AI well. First, you'd be able to prompt the tool to give you the right answer – which may itself require new ways of thinking about and using language. And second, you'd be able to understand and judge the output, and how it might be interpreted by both human and machine readers. It may be that those working in reporting end up writing fewer words themselves (and focus their efforts on the opinion pieces and forward-looking statements), but their depth of expertise in their subject, and their ability to communicate it using the right words, becomes more important than ever.

> "[AI] is an endlessly enthusiastic assistant that never runs out of energy, and never gets frustrated that you keep asking questions over and over and over to go into more and more detail."
>
> Governance advisor

How to use generative AI well: Practical recommendations

In this section:

Our focus groups discussed a considerable range of uses in reporting, from administrative to creative. We've set out below the results of that research, but some of these are grey areas, particularly drafting and editing. So do keep in mind the 'precocious intern' principle, and ask yourself, would I ask them to do this? And would I trust their output if I did?

What to use it for

- **Administrative – summarising meetings and calls, creating transcripts, writing minutes, summarising notes.** Using generative AI for these purposes can save a tremendous amount of time and it's generally quite good at doing them. However, everyone who'd used generative AI in this way so far urged caution, especially for writing minutes, because it often leaves out important points, or summarises in a way that is not true to the spirit of a meeting. Two particular notes of caution:
  - 1) People often don't necessarily want to have their words recorded and summarised by an AI tool, so may edit themselves and be less open in what they mean, particularly when sensitive or confidential material is being discussed. It is also worth remembering that any AI summary has a digital afterlife that could come back later on.
  - 2) There's more to what happens in a meeting than the words themselves; a lot of meaning and nuance comes from the tone and the way that people interact with each other, which is something that generative AI will miss.
- **Research.** Generative AI is brilliant at reading vast quantities of material and summarising it for you or responding to questions about the content by sifting through to pick out details. An excellent use is to do the first 'read' of sources for you so that you can then decide whether or not a source is worth looking at. This can save hours – or indeed allow you to look at far more material. But don't forget that it can get things wrong – very wrong in some cases – so it's essential that you check any output for accuracy before relying on it.
- **Drafting and editing material that is not opinion (see 'best avoided' below).** Drafting vs editing was the subject of heated debate. While all agreed that generative AI could produce something reasonable, some felt it should never be used for creating a first draft from a set of points, since that initial thinking process is the most important part of writing a draft. Others felt the opposite – starting with a blank page was so difficult that getting AI to create an initial rough draft was the best place to start, but they would never then use it for finessing. For them, that was where the real thinking and personalisation came in. Regardless, generally most agreed that the tool should be used for routine disclosure and narrative sections, but kept out of matters of opinion.
- **Tidying up and updating disclosure statements.** A lot of narrative content currently required for annual reports is straightforward disclosure that changes little from year to year – for example a description of a risk management process, or how a company has complied with certain requirements. Putting the previous year's statement into a chatbot to be edited and updated with this year's content as a starting point could be a quick win – as long as the output is properly checked.

Use clean, well-structured data

Generative AI is only ever as good as the instruction you give it, along with the information it is working with. This is particularly important for companies whose data is poorly structured, or poor in quality. A generative AI tool cannot fix the underlying problems with that information, nor indeed tell you what those problems are. The model will get confused if you give it, say, a series of different kinds of documents or data formats for the same information, since it may not be able to determine which information you're trying to extract. Be sure to clean up and structure the underlying information properly.

THINKING: WHY A HUMAN WRITER IS DIFFERENT FROM A MACHINE

A lot of debate during our research centred on where the role of generative AI should begin and end in drafting and editing, one of reporting's most critical tasks. When it comes to the question of opinion being from management and the Board, some people argued that using a human writer is no different from using a machine. It may be true that using a poor writer would be no different – perhaps worse than – using a machine. However, we, along with many participants, believe that, if the human writer does their job properly, they will produce something fundamentally different from a machine, and better, because it is indeed the opinion of the relevant individuals, whether or not those individuals were actually the first to write the words.

Good writing: the product of good thinking

This is because good writing is the product of good thinking. 'Doing their job properly' in reporting requires the writer to discuss and challenge those individuals, to prod their thinking and to understand what they mean when they're speaking, which is not always necessarily what they say. This may sound contradictory, but, as any writer will know, the key to good writing is to start from what you want the reader to understand, not what you want to say. And so a good writer will find a way of bringing the true meaning out of those individuals, to ensure that what they really mean is clear to readers.

Generative AI does not think

In our view, if a person hasn't thought a thought, but had it presented to them for comment, then it cannot truly to be said to be their opinion. This is equally true whether that thought has been drafted in the first place by a human or a machine. And let's not forget that generative AI does not think. What it produces is a replacement for what humans produce as a result of thinking, but it gets to that result not by thinking, but by running an algorithm that produces a combination of words most likely to receive a positive response to a prompt.

What to use it for continued

- **Creating visuals.** With the exception of photography, most graphics and visuals in corporate reporting are relatively simple. But, with the exception of highly bespoke tools like those created by media company WPP (found through our quantitative research, see page 13), generative AI tools have not yet produced visuals to a good enough standard for a high quality annual report, perhaps because they have not yet been trained properly on a company's brand. But they do allow you to experiment with data visualisation quickly and easily, shortcutting the briefing process for creation by an expert. Given the sophistication of generative AI, this may be an area that improves quickly. For other types of visuals that are more thematic and abstract in nature, AI-generated visuals may be a good alternative to stock photography, although don't forget that imagery must also pass the accuracy and truthfulness test. This means that pictures of real people and things should be photographs, and should not be created by generative AI.
- **Retouching (see also 'best avoided' below).** Photos can be uploaded to generative AI for retouching and airbrushing. To what extent an image should be airbrushed to make someone or something look better is a grey area, but as long as the image still basically tells the truth, it should be fine. For example, retouching to even out someone's skin tone to make them look their best is fair enough, because they could genuinely look like that.
- **Interim proof reading and checking (although this must itself be checked).** Rather like the checker for spelling and grammar in Microsoft Word, generative AI can usefully do a more sophisticated initial pass on copy for you to check against style guides or other references. However, it's essential that you then check the changes, in the same way that it's essential to check what the Word version suggests, which itself often produces incorrect results.

> **Don't use an external chatbot!**
>
> It may sound obvious, but using an external chatbot for confidential information breaches the rules of confidentiality. Once you have fed in your information and prompts, you have no control over how they are used. Your information becomes part of the information that the chatbot uses in future. Using the precocious intern analogy, would you really hire an intern to help you with confidential corporate information who was working simultaneously for your competitors? And who hadn't signed a confidentiality agreement with any of you?

What not to use it for

- **Writing the brief.** Not all companies write a brief for their annual report or other reporting – they should! A good brief sets out the report's purpose, the intended audiences, the (truthful) story it is trying to convey, what audiences should take away from it – as well as more practical matters such as length, editorial style, timing and so on. It is essential for a high quality report, efficiently produced, since it serves as the guide for everyone involved (sometimes a huge cast). A good brief is the intellectual foundation for the report, and so it 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.** Like the brief, these pieces should come from the minds of those responsible for them. (See page 20 on why the human writer is different from a machine.)
- **Photos of real people and things.** Like words, photography must be accurate and tell a truthful story too. AI-generated imagery of real people and things is not truthful so should not be used. Likewise, fundamentally changing an image, for example airbrushing out a feature or adding something in (for example a safety feature where none exists) should not be done, either with generative AI or any other tool, since the image would no longer tell a truthful story.
- **Poor quality data or other underlying information.** While we'd like to think that all underlying information is of high quality, we know that reporters often have to deal with partial or otherwise less than ideal sources of information. The extent to which you use this kind of information and how you caveat that within the report is a matter of personal judgement, so it would be unwise to use generative AI in this context, since it cannot distinguish quality. Generative AI is best used on data that's reliable, with verifiable sources.
- **Final proof reading and checking.** The word 'final' is used deliberately here. As noted above, there may be a role for generative AI during the reporting process to check drafts against style guides, reporting requirements or source material, but the final reading and checking must be done by those responsible for it who can judge whether or not it's correct.

Getting the best out of it: training and practice

Like anything else, getting the best out of a tool requires first of all that you know how to use it. Let's consider our precocious intern analogy. To get the most value from a new intern, you wouldn't hand them over to the newest, most inexperienced person in the company, who'd have no idea what to ask them to do or how to put them to work effectively. (And if our intern were truly precocious, they'd be unlikely to accept such a job!) Rather, you'd make sure the person managing them was experienced enough to give them the right tasks to do and to train them up properly.

What was surprising in our research was how many companies were bringing in tools like Copilot without formal training or proper explanation, or indeed governance around how it should be used. This matters because the general principle for getting the best out of generative AI is, the better the prompt, the better the output. And that requires training and practice. A number of people who'd been given Copilot to try out gave up because it produced nonsense, and it would have taken far more time and effort to work out how to use it well than just to do the work themselves in the first place.

It's perhaps unsurprising that many companies aren't training people. A lot of people seem to think that a chatbot or Copilot is a glorified version of a search engine like Google, and all you have to do is feed something in, ask your question and you'll get the right answer or the right output. In practice, it's much more complicated than that – and to get good, usable answers you need to give it the right material and ask it in the right way. Rather like how you'd carefully and clearly brief your intern rather than giving them vague instructions with vague information and expecting perfection in return.

> **Work with your IT team**
>
> Enthusiasm for generative AI amongst those who expect to make money from it is predictably high and while this means there is a lot of over-promising and under-delivering, it also means there are a lot of vendors desperate for case studies. If you want to be ambitious with your use of generative AI, there is no better time to talk to your IT/data teams and see if a low-cost proof-of-concept can be arranged. Just don't forget the process and governance components that need to sit alongside the technology.

To get the best out of generative AI, then, the following are essential.

- **Introduce a proper training programme with modules for reporting/confidential information.** Providers of generative AI tools already offer training alongside the tools themselves and larger companies may create their own training programmes. Given the importance of reporting and the risks involved in using generative AI in the process, any basic training should be augmented with specific training for its use in reporting and in dealing with confidential information more generally. This training should include guidelines/governance on how generative AI should/shouldn't be used in reporting (see below). It's also important to make sure people involved in reporting take part, and to track participation so this can be reported if needed.
- **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. The fact that 'prompt engineering' is a highly sophisticated aspect of the AI field is perhaps instructive here. Prompt writing may look straightforward, but it isn't. The more you use your generative AI tool, the better you will become at getting it to do what you want, because you will learn what inputs get the kind of output you need. In the same way that, the more you work with your intern, the better you will understand how they interpret what you ask them to do. But don't forget that, while generative AI is a highly sophisticated language tool, it cannot knowingly give you a truthful answer.
- **Become a better reader.** The problem with a lot of reporting even without the use of generative AI is that it's not particularly well written in the first place. Too many reports are written from the starting point 'what do I want to say', not 'what do I want my reader to understand'. To use generative AI well, you must therefore become a better reader, constantly thinking about what the output means and whether it will be understood by your audience.

Ensuring it's used well: governance and guidelines

Given all we're hearing about the importance of governance and guidelines to using generative AI, it was quite surprising that our research showed a general lack of both, not to mention of training, within companies. Some companies – generally the more safety-conscious, risk-averse – had very strict guidelines in place, although in most cases, it was unclear how usage by employees was being monitored.

Governance to manage risk

This matters because there's plenty of anecdotal evidence that people are using external chatbots without permission. Clearly, companies need to tighten up in this area if they want to ensure the security and confidentiality of their information, let alone be able to account for how generative AI is being used across the organisation. When it comes to reporting, you must be able to verify everything, whether or not you include such evidence within the report itself. To be able to do this, you must therefore know where the information came from, who created it and how, so that you could, if challenged, substantiate it. Which means accounting for how generative AI is being used.

Guidelines to ensure effectiveness

Instructing people about what they should and shouldn't do isn't just about managing the risks and ensuring accountability. Guidelines should be just as much about how people can use generative AI well to help them in the reporting context – while ensuring they don't use it in ways that will cause them problems later on in the reporting process.

> **An opportunity to improve your existing process?**
>
> If you have a well organised, well documented, rigorous process for reporting, introducing generative AI should be fairly straightforward. If your process is somewhat haphazard, then introducing generative AI may introduce risks for validating and verifying information – since it's essential that usage is properly understood and documented. You could therefore use the introduction of generative AI as an opportunity to overhaul your reporting process and create something far more robust and rigorous.

Develop existing governance and guidelines to accommodate generative AI

Generative AI – particularly when it comes to reporting – is not a new, separate endeavour. It is something that is being brought in (we hope) to do what we do better and more efficiently (our precocious intern). Therefore, the general principle here is to build generative AI considerations into the governance and guidelines already in place within companies.

- **Get the use of generative AI in reporting on the Board's/ Audit Committee's agenda.** All Boards are discussing generative AI and its impact on the business; very few, it seems, are thinking about it specifically in relation to reporting. It should be part of their discussions.
- **Develop your reporting policy to include the use of generative AI.** A good policy should cover risks and benefits, how generative AI should and shouldn't be used; who should and shouldn't use it and how; and what training is required. To do this, you will need to:
  - Work with IT to understand and document how generative AI is being used across the organisation and therefore how it might affect the reporting process
  - Decide what uses in reporting you will and won't sanction (see section above on good and bad uses within reporting)
  - Review your existing policy and adapt it to cover the use of generative AI – and make sure it's written very clearly so that people can understand it and carry it out.
- **Communicate the policy to the whole company**, so that everyone involved in providing the information on which reporting relies is aware of the policy, and can ensure it is adhered to.
- **Document the use of generative AI during the reporting cycle.** In the early days of using generative AI in reporting, this will be useful to ensure people keep the principles in mind and use it in the right way, in line with your reporting policy, reinforced by the training. As the use of generative AI becomes 'business as usual', this may no longer be necessary, but don't forget you will always need to be able to verify everything within your report, including all the source material.

Explaining how you've used it

From our research, we know that investors, the principal audience for reporting, have a keen interest in how companies are using generative AI, including in their reporting. Participants from companies also felt that their audiences – not only investors but other stakeholders too – would want to know. This reflected a general view about the importance of knowing what has been created by generative AI, whether in the corporate sphere or beyond, in news items, art, literature and so on.

The discussion about disclosure with both investors and corporates reflected the importance of ensuring that any use of generative AI does not undermine the purpose of reporting, i.e. the accuracy of data and information, and the truthfulness of the story.

Considering the likelihood that generative AI is going to become embedded in corporate systems and therefore touch reporting in many ways, that everyone is interested in ensuring that people remain in the loop, and that most concerns about generative AI usage centre on how it might be used to write opinion, then we propose the following approach for explaining how you are using it. In general, if you're not using it or only experimenting, then say so, because we know people are interested. If – or rather, when – you are using it, explain how you are doing so in a way that upholds the accuracy of the data and information, and the truthfulness of the story .

If you're not using it, or are just experimenting – state your principles

With the 'writing for your reader' principle firmly in mind, if a company has not yet started using generative AI in its reporting, and/or has not yet begun to establish any governance or guidelines around its use, they should include a short, general statement explaining their position. This is because we know that reporting audiences are interested in how generative AI is being used, and therefore it's a good idea to tell them.

To reflect that the use of generative AI in reporting is not properly on a company's agenda yet, it would make sense to include this statement, for example as a box on the inside front or back cover of the report, where people generally expect to find information about how a piece of communication has been produced.

You could consider a discussion in the governance report, but only if there's something useful (and truthful!) to say about it.

If/when you are using it – be specific

Once generative AI is properly embedded in the reporting process then a company should include the following in their reporting:

- **The policy for using generative AI in reporting, and a statement that the report in question has been created in accordance with that policy.** This would be included in the relevant section of the governance report, wherever makes the most sense in terms of the type of company and its relevant disclosure requirements. Here are some suggestions for UK companies, based on where directors are most likely to discuss their responsibilities for accuracy and integrity of information.
  - **For companies reporting under the UK Corporate Governance Code:** within the fair, balanced and understandable assessment.
  - **For companies reporting under the QCA Code:** within Principle 5, which covers risk management, internal controls and assurance.
  - **For companies reporting under the Wates Principles:** within Principle 3, director responsibilities, which covers accountability and integrity of information.
  - **For companies that don't use any of these:** it may make sense to include a section in the directors' report.
- **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.** These are the sections that investors and other audiences most want to know about – and are most concerned about when it comes to the use of generative AI in their creation. Therefore it would benefit companies that don't use generative AI to include a statement along the lines of 'this section was not written by generative AI'; or 'no generative AI was used in the creation of this section/ statement'. This of course would imply that generative AI may have been used in all the sections in which this statement does not appear.

> **Investors: any use of generative AI in decision-making *must* be disclosed**
>
> Investors raised a parallel point about the use of generative AI in decision-making. Since this would affect their view of the company, and therefore the decisions they make about investing in it, they want such information to be clearly disclosed.

THE COMPANY VIEW: WIDESPREAD SUPPORT FOR DISCLOSURE

Almost all corporate participants believed that their audiences would have less trust in reporting created by generative AI than reporting created by people. Some even commented, somewhat tongue-in-cheek, that they'd like to use generative AI to make their reporting sound better, but wouldn't want to disclose it so that people would still assume it was written by them!

It's unsurprising, then, given the commitment of the research participants to high quality reporting, that most felt that the use of generative AI should be disclosed in the annual report, at least in the early days. But to make such disclosure useful, you'd need to explain what type of generative AI has been used, for what purposes (e.g. as a research/analytical tool or to draft or edit sections) and in what sections of the annual report.

THE INVESTOR VIEW: SUPPORT FOR DISCLOSURE IF EXPLAINED WELL

Amongst investors, the feeling was more mixed. Some felt that disclosure would be pointless without a proper understanding of why such disclosure is needed given that, in due course, most people are probably going to use generative AI in some form anyway. Nonetheless, they still felt strongly that the opinions expressed must truly be those of management and the Board, however those are created. And, of course, that all data and information is accurate, with people in the loop to ensure that it is.

Those strongly on the side of disclosure, regardless of context, were particularly concerned about the use of generative AI to write opinion pieces, since that kind of usage tells you something about how seriously those leaders take reporting as a key piece of communication (i.e. not very) implying a wider contempt for their audience. This view did not, however, seem to apply to those leaders working with people who help them write their statements (see page 20 on why a human writer is different from a machine).

Appendices

Appendix 1: Quantitative research findings

What reports told us

In this section:

Insig AI analysed 21,350 annual reports and other corporate documents published by FTSE 350 listed companies in the calendar years 2021 to 2024 inclusive to discover how companies are discussing their use of AI. Our aim was to provide evidence-based insights to complement the qualitative focus group research which aimed to find out how companies are using generative AI, and how that might affect reporting. We looked at two things: general trends over time, across types of documents and between sectors, and detailed contextual searches, particularly for any mentions of the use of generative AI in reporting. Please refer to the methodology in appendix 3 for further details on the process and technology used.

Summary of key findings

- The overall increase in mentions of AI observed between 2021 and 2024 is driven by the use of general AI terms which are found predominantly in the annual report.
- Contextual analysis of these mentions revealed that AI is mainly discussed in terms of risk, strategy and operational efficiency and automation.
- No company referred to the use of AI in the reporting process. WPP used it to create its front cover artwork, promoting its own creative services in AI. Ruffer Investment Company used it to generate an image used in its report.
- Only three Codes of Conduct, all published in 2024, referenced AI for employee use, although an updated search between January and March 2025 suggested this is changing fast, since this number increased to 17.
- It's yet to be seen whether the individual company or sector patterns of high levels of general commentary about AI will influence early adoption of generative AI in reporting.

General trends over time

Our starting point was to get a sense of the trends over time in the evolution of how companies have started to talk about AI. Unsurprisingly, companies have been mentioning AI more and more since 2021. There has also been an increase in the overall number of companies mentioning AI in their corporate documents.

Evolution of average number of mentions of AI in FTSE 350 annual reports compared with all corporate documents

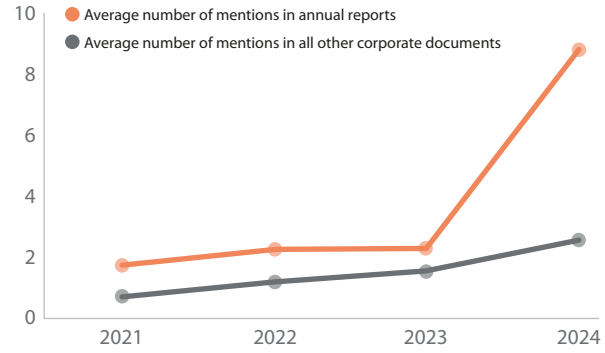

Evolution of number of FTSE 350 companies mentioning AI in their annual reports compared with all corporate documents

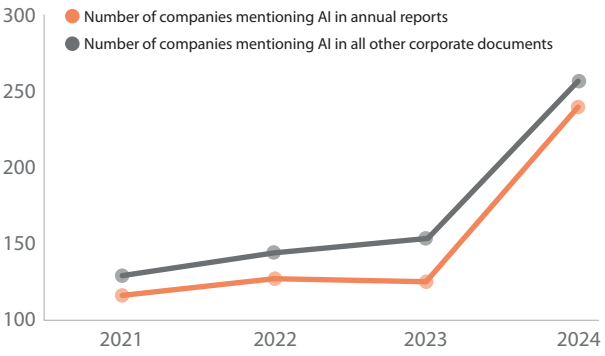

Looking for mentions related to reporting

Our hypothesis was that companies are likely to be talking about AI in quite general terms, and in relation to business operations and overall risk and opportunity, rather than more specifically in relation to reporting. However, we also felt it would be useful to see if there was any discussion about generative AI being used internally within their own information ecosystems, because of course this would then affect reporting, even if not discussed as such.

To test this, and to focus in on the use of generative AI in reporting, we divided our keywords into the following categories (see methodology for detail) and ran the same analysis over time. We chose these categories so that we could then perform more detailed contextual analysis in relation to reporting:

- **Broadly used (general) AI terms** (AI, artificial intelligence)
- **Generative AI terms** (generative AI, gen AI, genAI, large language model, GPT, Copilot, chatbot, Midjourney, Dall-e)
- **Governance-related AI terms** (AI strategy, AI policy, AI governance, AI principles, responsible AI, AI working group, AI guidelines, AI committee, use of AI).

Evolution of AI mentions by type in all FTSE 350 documents

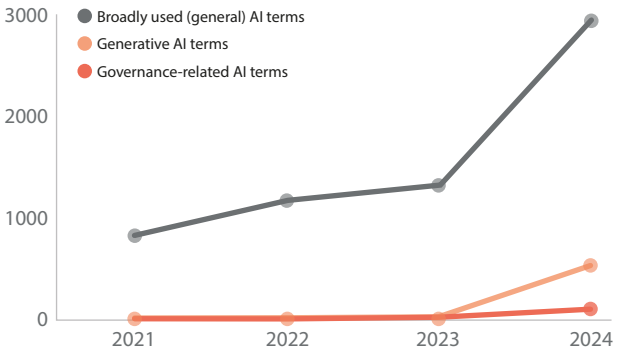

Evolution of AI mentions by type in FTSE 350 annual reports

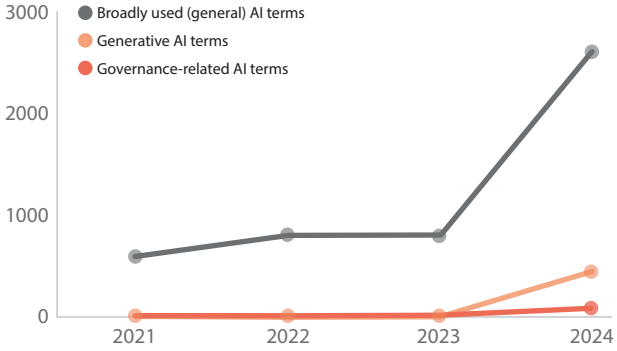

General mentions of AI naturally dominated, and we saw a notable increase in the last reporting cycle. Unsurprisingly, mentions of generative AI don't appear before 2023 since these models only really came into public awareness with the release of ChatGPT in November 2022. The advent of these is likely to also explain the rise in general AI mentions, with applications for generative AI being considered in all industries. AI governance mentions, however, lagged considerably, with some emergence of mentions in the last year, but the numbers remain relatively low.

To find any notable mentions relating to the use of generative AI in corporate reporting itself, or in internal systems that might affect reporting, we manually reviewed and read in context all specific mentions of generative AI terms among the FTSE 350 documents. To cross-check our work, we also used ChatGPT to ask the same question of the same set of information. Both methods revealed no references to generative AI in relation to the reporting process. Rather, this contextual analysis proved our hypothesis that companies were talking about AI in terms of risk, strategy and operational efficiency and automation.

Distribution of mentions across documents

To get some perspective on where FTSE 350 companies are discussing AI, and therefore where we might expect to see further discussion in future, we looked beyond the annual report and counted total references to AI across all documents. Grouping documents into types, the annual report and related financial documents (half year reports, results statements and quarterlies) accounted for 88% between them, with annual reports accounting for the vast majority at 70% of mentions.

Codes of Conduct are included within ESG policies, and, in 2024, just three mentioned AI in terms of guiding its use by employees, although an updated search ran between January and March 2025 showed that this number had increased to 17.

2024: distribution of AI mentions across all FTSE 350 documentation

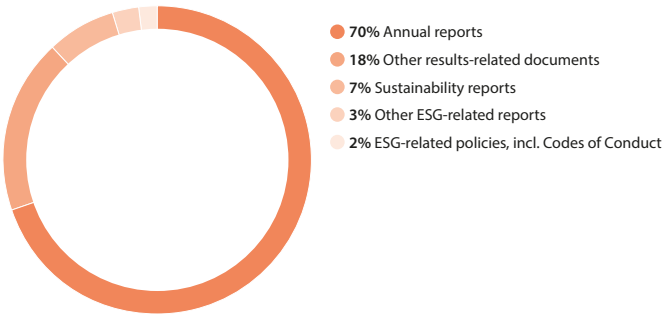

Sector trends

Looking at the FTSE 350 as a whole, we also wanted to see if there were any sector-related patterns about how companies talk about AI, and whether there were outliers that might indicate companies being early adopters, and therefore ones to watch with regards reporting. We analysed the average number of mentions by companies in each sector over time (following MSCI's Global Industry Classification Standard, or GICS).

Unsurprisingly, the information technology sector leads the way in terms of average number of mentions, but we can also see a sharp rise in 2024 in the communication services, consumer discretionary, and healthcare sectors. We felt these warranted further investigation so broke those sectors down into GICS industry groups to see if there were any further trends. The charts here show mentions in annual reports only, but the pattern was mirrored across all documentation as well.

Average number of AI mentions in FTSE 350 annual reports by sector

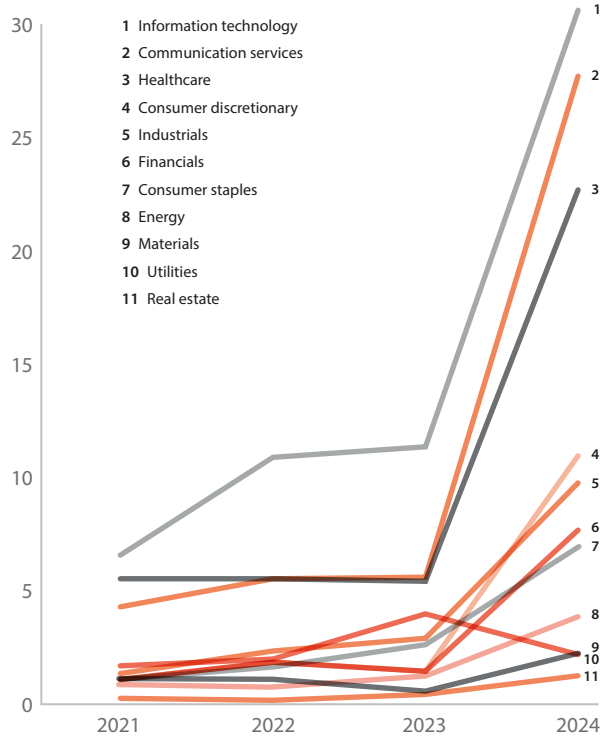

Sector trends continued

The utilities sector is an anomaly with mentions dropping significantly in 2024, so we looked into the detail of that sector to see why. Because there are relatively few companies in this sector, the results are disproportionately affected by a single company, in this case National Grid. Its number of mentions dropped from 11 in 2022 and nine in 2023, to just one in 2024, which caused the steep drop in the overall results seen in the chart. Since National Grid has a March year end, at the time of publishing this paper, the subsequent year's report had not yet been published, so we could not check to see if the 2024 result was itself an anomaly. It will be interesting to see what National Grid says this year, and why the number of mentions dropped so significantly in 2024 in contrast to most other companies.

Evolution of AI mentions in FTSE 350 annual reports: communication services

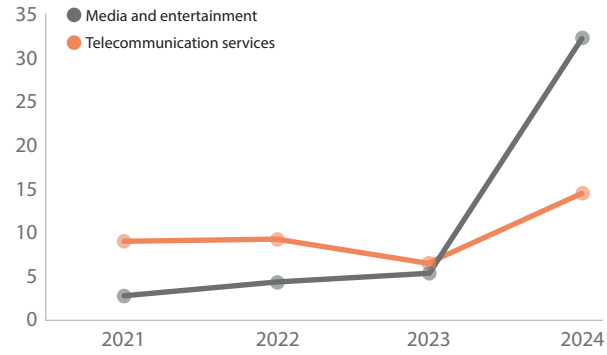

In communication services, the rise is mostly driven by media and entertainment companies. The highest numbers of mentions are made by WPP plc, a communications and advertising company, which states that it uses AI and data to enhance its offering. It did not, however, discuss using it in its reporting, although it did use AI to generate its front cover to promote its own creative services. Since WPP published its 2024 report in March 2025, although the date was outside the scope of this research, we checked to see if generative AI had been used this year too, and got the same result, i.e. they used it to generate the front cover.

Evolution of AI mentions in FTSE 350 annual reports: consumer discretionary

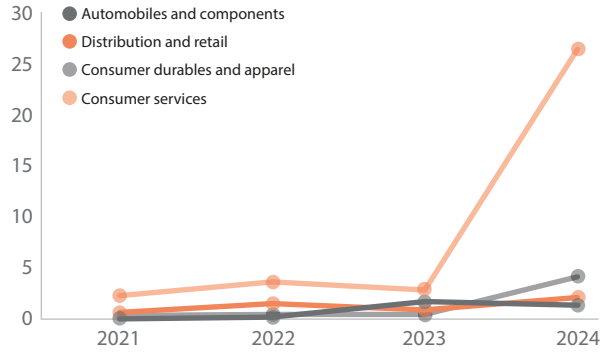

The increase in mentions in the consumer discretionary sector is dominated by consumer services. The highest number of mentions come from Pearson plc, which said it is using generative AI in its offering to support learning.

Evolution of AI mentions in FTSE 350 annual reports: healthcare

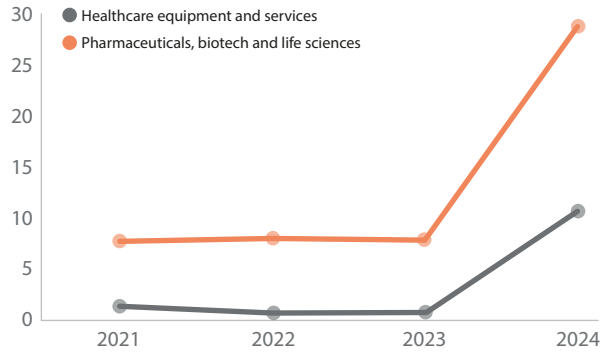

In healthcare, the increase in mentions covers the whole sector.

A watching brief on certain sectors

If we take these high levels of mentions of AI as an indicator of early adoption by certain companies and sectors, it's worth keeping a watching brief to see whether these trends play out in influencing early adoption in the use of generative AI in corporate reporting.

Appendix 2: Qualitative research findings

What research respondents told us

In this section:

The Falcon Windsor team recruited FTSE companies and investors to take part in a series of focus groups to find out how generative AI was or was not being used inside companies and in reporting specifically, and to discuss how it could be used responsibly, considering its potential benefits and risks. In total, 60 people responded, representing 40 companies, five investors and one proxy agency, plus a number of others. Claire Bodanis, the overall research team leader, chaired the company and investor focus groups, which took place between June and November 2024. For full details of how the research was carried out, please see the methodology in appendix 3.

Summary of corporate participants' views

- How generative AI is being adopted by companies (not specific to reporting)
- Current use of generative AI in reporting
- Views on future use – benefits/risks
  - Uses as a drafting tool
  - Benefits
  - Risks
- Views on regulation and disclosure
- Other points emerging

Key theme: astonishing variation, driven by culture and personal views

The key point coming through was the astonishing variation across companies with regards how AI is being adopted (or not), the pace of adoption, and the usage/non-usage of AI within the reporting process. The only determining factor seemed to be culture, particularly the personal interests and personalities of leadership and the Board, and a company's attitude to risk. There were no other factors in common such as size of company, industry and so on.

There was also considerable variation amongst participants with regards to how they feel personally about the use of generative AI (gen AI) tools, both in reporting and more generally in business and personal life. Most accepted that it'll be part of the world of work and it's about making sure it's used sensibly and responsibly; in general those people also had some concerns about the widespread use of gen AI tools and were keen for strong governance around it. A few people were very worried about the risks of gen AI, specifically in reporting but more generally as well, while a few were very excited about the potential of gen AI to transform the world of work, and confident that any risks could be mitigated.

How generative AI is being adopted by companies (not specific to reporting)

- **Range of approaches to adoption.** A few companies are not using or investigating AI usage at all yet, although most are looking at or testing it in some way. This ranges from ad hoc usage of chatbots, to ad hoc testing of systems like Copilot, to more structured adoption of AI technologies, to very tightly controlled and planned adoption, with tightly defined use cases, governance and training.
- **Considerable variety in AI governance, policies and training.** Of the companies that are using or testing AI (most that took part), the range of approaches to governance, policies and training was again extreme. Some companies were very cavalier in their approach, with no guidance at all, and external chatbots permitted for use on internal material. At the other end of the spectrum, some companies – generally those with more risk-averse, safety-driven cultures – had far more caution, with use cases required for any use of AI at all, and highly developed policies and guidelines in place. Most companies were somewhere in between.
- **General view that external chatbots should not be used.** Almost all companies did not sanction the use of external chatbots for corporate information, although only a few had actual policies or controls in place to ensure this would be the case.
- **AI tools in use – mostly Copilot.** While some, more sophisticated companies are building their own internal chatbots and AI assistants, and using a range of technologies, unsurprisingly most companies are adopting gen AI tools in their information ecosystem through Microsoft Copilot. This is likely to have the most impact on reporting given that a) it's perceived as quite safe (because it mirrors company-defined existing data access permissions where they exist) and b) it's on people's desktops and likely to become an everyday tool in how people manage work.

Summary of corporate participants' views continued

- **An experimental tool – training gap on how to use gen AI well.** Unsurprisingly given how new AI tools are, many people felt they didn't really know how to use gen AI tools well, and, so far, training is relatively limited. That was not an issue for everyone – some are happier than others to play around with new tools; others felt experimenting was a waste of their time.
- **The vital importance of the human checker/subject matter expert (SME).** There was unanimous agreement over the importance of the human checker, the SME. Some were concerned about the future of the human checker/SME – how would these people become experts if much work is taken by AI? Despite these concerns, little is being done in practice within HR/talent development so far.

Current use of generative AI in reporting

- **Generally little used; some trials running; some ad hoc usage.** So far, very few people are using gen AI in preparing reporting, although some are experimenting with tools like Copilot for drafting. A few, generally larger, companies have formal projects in place specifically for reporting, particularly in the finance and technical teams. Some people have been trying out Copilot for editorial purposes – either giving ideas for first drafts, or editing drafts already written. These tend to be people in companies with looser governance around AI; companies with tighter governance are, in general, currently prohibiting the use of gen AI/Copilot/chatbots in reporting given the perceived risks to accuracy and compromising confidential information – with such prohibition often coming from the legal team.
- **Direct influence from leadership in a few cases.** A few people whose leaders are very AI-focused commented that they are being challenged to use gen AI tools in reporting; generally those leaders are not heavily involved in reporting so don't see the risk aspects. A small number of people who would like to use gen AI more in reporting are being discouraged or blocked from doing so by leadership or policies that are very risk-averse towards AI.
- **Likely to be used as an editorial tool as companies get more comfortable with gen AI.** Most people felt that, in the longer term, they will use gen AI tools like internal chatbots or Copilot to help them with tasks like drafting, editing, and proof reading against a company's house style. Most felt, however, that such tools would never be more than an assistant, although a few envisage a future in which the whole annual report is created entirely by gen AI, with people responsible for checking the output.

Views on future use – benefits/risks

Almost everyone felt that using gen AI in the reporting process in some way is inevitable, so it's about finding the right way to incorporate these tools to help us do the increasingly challenging job of reporting well. Reporting is only getting more complicated, with increasing disclosure requirements making it even harder to tell a coherent story; and there is no more time (and often no more people) to do the work. So gen AI is very attractive if it can alleviate some of that workload.

In general people focused on gen AI as a drafting/editorial assistant. However, using it as a research tool as well was discussed, and a number of people felt that gen AI tools would be useful in dealing with the increasing 'volume of reading' needed before the actual drafting begins, for example by summarising sources of information, or doing a gap analysis post drafting to check if anything significant has been missed.

Uses as a drafting tool

- **Story vs disclosure-driven information – strategic vs mechanical exercise.** A big area for discussion was what type of information/section would most benefit from or be at most risk from the use of gen AI tools, and there were very mixed views. Some felt that gen AI should never be used to create opinion-type pieces since it would not be an authentic view of management/the Board, but could be used effectively for more disclosure-driven information. Others were much less worried about this. Views often depended on how opinion-type pieces were created currently. If a CEO or Chair currently doesn't have much involvement in his/her own statement, then there was much less concern about outsourcing drafting from a Head of Investor Relations or Company Secretary to an AI tool. Where the development of narrative is a more strategic exercise, interviewing the Chair or CEO, bringing senior people together to discuss the story and message, there was more concern about outsourcing that 'thinking' to an AI tool. An important point in this debate was where the line can be drawn in terms of whether the real author of a piece of text is gen AI or the person prompting it, and indeed whether or not that even matters (see below with regards directors' responsibilities for signing off the report).

Summary of corporate participants' views continued

- **Drafting vs editing – strong views on both sides.** Related to the point about strategic vs mechanical, some felt that the important first draft, the 'creation of the story', should be done by people in discussion as it is now, and gen AI tools could be used usefully later on to help edit. In other words, gen AI would be more of a glorified spell checker/proof reader/editor. Others felt that gen AI is best used to prompt ideas, a starting point that you could then edit yourself rather than staring at a blank page. A number of people felt very strongly about these two approaches and would only countenance one or the other, not both.

Benefits

- **Efficiency/time-saving.** The overwhelming benefit discussed was how gen AI tools can cut the time needed for drafting or editing, although a small number of people were concerned that introducing a gen AI tool would make verification and checking more difficult. There is no solid evidence about the time or efficiency savings, but anecdotally people have found tools like internal chatbots and Copilot useful for editing, with the caveat that everything must be edited/checked thoroughly. It's possible that those who found it most helpful were better at writing the questions/prompts that would give the best output; it's also possible that their own corporate systems were better set up in the first place to make best use of the technology. And, it's possible that a tool like this would be especially useful for people whose native language is not English (assuming that the language of such a tool is English) – although none of these points are proven! Some people, probably more skilled and experienced writers, questioned whether it was really much of a saving to spend the time coming up with the right questions and prompts for a tool to create a draft that would have been quicker and easier to write yourself.
- **More useful for more structured information.** Some felt that gen AI had more to offer parts of the report where information was more structured, and therefore there was less nuance over language and opinion.
- **Could AI improve reporting by reducing bias?** A couple of people suggested that reporting could be more balanced if written by a gen AI tool, since it could be asked to create a balanced response free of human bias. This may be possible although it overlooks the inherent bias within all foundational gen AI models.

Risks

- **Accuracy and reliability of the output – although mitigated by the SME checking.** This was seen as one of the biggest risks, although most felt that with the right human checker/SME the risk of false/inaccurate output could be mitigated.
- **Authenticity of the company/individual voice – it'll all end up sounding the same.** A significant number of people, particularly those responsible for communications/the strategic report, felt that AI-generated narrative would lack authenticity. There were also concerns that reporting across the board would become generic if mostly drafted/written by gen AI. Some thought that training an internal chatbot on a company's tone of voice and information might mitigate that problem, although a question remains over how individual it could really be when the tools themselves are all based on the same underlying algorithm.
- **Corporate reputation, and the perception of leadership who 'can't be bothered' with reporting.** A number of people were concerned about how investors and other stakeholders would view a Board/leadership team who have 'outsourced' the important work of reporting to them to gen AI, although others felt that there's no difference between this and outsourcing it to people within the company. (See note below on perceptions of reporting generated by AI vs written by people with reference to disclosure.)
- **The value of the reporting process beyond the report itself.** The reporting process, when done properly, requires senior management to come together to think, discuss and debate. There is a danger of that valuable process being lost if drafting were outsourced to gen AI.
- **A reduction in critical thinking.** The processes of drafting vs editing are very different, and produce very different results. A few people felt that humans' ability to think critically would be compromised were gen AI to take over the drafting process. Others were much less concerned.

Summary of corporate participants' views continued

Views on regulation and disclosure

- **Regulating the use of gen AI in reporting would be inappropriate/pointless – but guidance would be popular, if even-handed.** Most people felt that trying to regulate how gen AI should be used in reporting would be pointless, considering the nascent nature of the technology and the speed at which it is changing, particularly given how slowly reporting regulation moves. Some felt that, while they would like there to be regulation, they couldn't see how it would work in practice. Most felt it would be inappropriate, while a few commented that reporting regulation is too overloaded already. However, while a few people felt that gen AI usage should be left entirely to companies, most felt that guidance from the FRC or FCA (or both) would be useful. It was noted, however, that any such guidance should be even-handed, and apply to audit firms as well as to companies.
- **Widespread support for disclosing the use of gen AI in reporting, at least for now – but the detail matters.** Most people felt that the use of gen AI should be disclosed in the annual report at least in the early days. But to make such disclosure useful, you'd need to explain what type of gen AI tools you've used, for what purposes (e.g. as a research/analytical tool, to draft or edit sections) and in what sections/parts of the annual report. Some people did not want to disclose the use of AI because they felt people would trust the report less; others felt that disclosure would be irrelevant, particularly those who felt it doesn't matter who or what writes reports as long as directors sign them off. In general, people felt that disclosure may not be needed in the longer term when the use of gen AI tools is established and understood.

Other points emerging

- **Better input = better output.** A general point, not specific to reporting, is that, as we get better at understanding how the tool works, and how it responds to prompts, we will be able to get more value out of it.
- **Directors' accountabilities are no different – but focus needed in the early days.** On the whole, people were not much concerned about the role of directors and their accountability: after all there is nothing different about their accountabilities for reporting; they are responsible for the output whoever or whatever has written it. Most people felt however that, at least in the early days, directors needed more awareness of the potential impact of gen AI on reporting so that they properly understand its role and impact on information for which they are accountable.
- **Narrative reporting may become generic and therefore redundant if it's all written by the same 'mind'.** A few people raised the question of whether, in the long term, the insight/narrative/opinion element of reporting might become generic, and therefore meaningless, if it's all written by the same 'mind', i.e. the algorithm behind gen AI tools. This may be a risk because companies are likely to be using similar gen AI tools based on the same foundational model (e.g. many will be using Copilot).
- **The importance of structured data for gen AI to be used effectively – patchy at best.** Many people commented that the introduction of tools like Copilot was shining a light on the importance of having well organised, accessible and structured data. If a company's information ecosystem is not well organised, gen AI can't work effectively on it, and in general, few companies felt they were really on top of this issue internally.
- **Value for money of the use of gen AI in reporting – not yet proven?** Some people, particularly those in smaller companies, raised the issue of the cost of gen AI tools including Copilot (which is perceived to be expensive), and whether the expense is worth it. Since most companies are only testing such tools at the moment, evidence of the cost/benefit is still to come.
- **Little discussion about the relationship of gen AI and energy usage/carbon emissions/net zero targets.** Only a small number of companies seemed to be considering the potential impact of widespread use of gen AI and its potential impact on their energy usage and net zero targets.
- **Everything must be checked by a human!**

Summary of investor participants' views

- Use of gen AI by investors in their own work/reporting
- Views on use by companies
- Views on regulation
- Views on disclosure of use in creating reporting
- Other points emerging

The general theme coming through all conversations was that generative AI (gen AI) tools are or will be necessary for dealing with the vast quantities of information now related to reporting. The main debate centred around how people use those tools and how to ensure that opinions and decisions remain the preserve of those responsible for them.

Use of gen AI by investors in their own work/reporting

- **Increasing volume of data makes gen AI as an efficiency aid to research essential.** All talked about using generative AI as a research tool to marshal large quantities of information, and how that's become essential. Increasing regulatory/reporting requirements mean that the volume of information about companies on which to make investment decisions is increasing rapidly, and it's impossible for human beings to synthesise all of that without the help of gen AI.
- **Some use of gen AI tools in investors' own reporting.** Some talked about using gen AI to help with their own reporting (e.g. stewardship reports); noted that it's essential to have a human in the loop, and those who are using gen AI in reporting said they wouldn't use it for creating opinion.
- **Investment decisions are still being made – and must still be made – by people.** All agreed that, although gen AI is being used in research, the investment decisions themselves should still be made by people.
- **Gen AI as 'low-level subordinate', and the importance of checking by a human.** Most talked about the use of gen AI as efficiency/productivity tools to help people do their jobs quicker and better, and focus on the important things (in reference to the increasing volume of information). But such usage was likened to a low-level subordinate, someone to do the heavy lifting, and anything produced by it must be checked.

Views on use by companies

- **Focus on general use by companies aside from reporting, particularly if gen AI tools are used in decision-making.** Unsurprisingly, all investors had a keen interest in the general usage of AI within companies – particularly in formulating strategy and making decisions. Investors need to know if companies are using AI in decision-making itself since it will affect their view of the company.
- **Use in reporting – views/opinions of management and the Board should be theirs and not be created by gen AI, but where does 'written by' become 'thought by'?** Very mixed views on this. Everyone agreed that views/opinions expressed in reporting should be those of management/the Board. If those views themselves were created by AI then that was a concern (see point above); but there was considerable debate over where 'written by' becomes 'created by' or 'thought by'. If an opinion is truly that of management/the Board, but the writing of it is assisted by gen AI, then that was largely felt to be OK. But where does that line get drawn? If you use gen AI to create a first draft and give you ideas, does that mean you're no longer thinking for yourself?
- **Concerns over the origin of information.** Given the well documented issues around the 'black box' that is how gen AI works, a few people raised concerns over the origin of information and how we would know where such tools were sourcing their information, and therefore what reporting is actually based on.
- **Accuracy/reliability of output: essential to keep the human in the loop.** Similar to the point about investment decisions being made by people, a common theme was the importance of people controlling gen AI tools, not the other way round.
- **Widespread use of gen AI tools would make it easier for companies to 'game the system'.** Because so much analysis of reporting is done by machines now, particularly around sustainability, people have increasingly been trying to shoehorn buzzwords into their reporting to tick a regulatory box. Using gen AI tools would make doing so far easier and therefore make it easier for companies to tick boxes, regardless of how good the information really is.
- **Corporate reputation, and the perception of leadership who 'can't be bothered' with reporting.** A critical issue for some was the idea that business leaders who outsource their opinion writing to a gen AI tool are demonstrating that they're not that interested in communicating with their investors and wider stakeholders, implying a wider contempt for those audiences.

Summary of investor participants' views continued

Views on regulation

- **Regulating the use of gen AI in reporting would be inappropriate/pointless – but some guidance might be useful.** Everyone felt that trying to regulate how gen AI should be used in reporting would be pointless, considering the nascent nature of the technology and the speed at which it is changing, particularly given how slowly regulation moves. Some also felt that the use of gen AI in reporting is not a problem for regulators to solve. Nonetheless most people felt that at least minimal guidance from regulators might be useful, if only to remind companies of their duties/accountability and the fact that using AI doesn't change those – i.e. use of gen AI in reporting should be within a company's governance framework. Some suggested that for guidance to be useful, it should comment on which parts of the report companies should disclose any use of gen AI.

Views on disclosure of use in reporting

- **50-50 for and against disclosure of use of AI in reporting, at least for now – but the detail matters.** Participants were generally quite evenly split over whether the use of AI in reporting should be disclosed, with those against saying it would be pretty pointless unless we have a fairly clear description of why such use should be disclosed, given that, in due course, most people are probably going to use it in some form anyway. Those on the side of disclosure felt quite strongly that it does matter if companies/directors use gen AI to write their opinion pieces, since that kind of usage tells you something about how seriously companies take reporting as a key piece of communication (i.e. not very seriously) – see point above on corporate reputation.
- **Use of gen AI in decision-making should be disclosed.** A parallel discussion centred around the use of gen AI itself in companies' strategic decision-making; investors generally felt they would want to know about this aspect of AI usage.

Other points emerging

- **Concerns over blind trust in big tech and lack of awareness of inbuilt bias.** One investor commented that in general, we are far too blasé about how big tech companies are taking control of the informational ecosystem and locking us into their architecture, while being unaware of the inbuilt biases in these systems. This is a wider concern for monopolisation and anti-trust.
- **Potential risk to the role of the proxy advisor.** Some commented that the use of gen AI as a research tool could put the role of the proxy advisor at risk – what is their business model if most of the research and analysis is done by gen AI? The proxy advisor commented that they do far more than just aggregate information and present the outputs of gen AI; instead, like investors, people not tech give opinions/advice on investment decisions.

Appendix 3: Research methodology

In this section:

Process for developing the research

2024

January/February – initial proposal developed with Imperial College London. Under our guidance, a team of MSc students from Imperial College London developed a research proposal for us as their corporate partnership project. The team presented this to us in February.

March/April – research plan developed; consultation with the UK's Chartered Governance Institute (CGI). Our research team developed the research plan to encompass qualitative research with FTSE companies and investors, and quantitative research based on all published reports, along with other corporate publications, using Insig AI's bespoke research tool. We discussed our plan with the CGI, who were conducting a wider piece of research on the governance of AI within UK plc.

April/May – research participants recruited. We recruited research participants for our corporate and investor focus groups.

June-August – five company focus groups held and feedback received, some observed by the CGI

- Claire Bodanis chaired a series of focus groups with company executives directly involved in producing reporting and statements to the market.
- Five focus groups were held, on 4, 18, 19, 23 and 29 July.
- In July, the CGI included a call to join the research in their monthly Technical Update.
- Claire also chaired one focus group, on 25 July, with a number of interested parties from the governance and advisory sphere, including a participant from the CGI.
- On 8 August, Claire ran a follow-up session with one FTSE 100 company which was most advanced in creating guidelines and policies around the use of generative AI for corporate information.

September-November – three company and three investor focus groups held and feedback received, some observed by the CGI; one legal session held

- Claire chaired three company focus groups on 9, 23 and 26 September, and two investor groups, on 18 and 25 September.
- On 16 September, she also ran a follow-up session with one FTSE 100 company which was more advanced in thinking about generative AI in reporting.
- The research team compiled a feedback document for each of the corporate and investor participant cohort which was circulated to relevant participants for comment/approval (included as appendix 2).
- On 17 October, Claire ran a session with a legal firm to discuss any legal implications of the use of generative AI in reporting.
- On 11 November, Claire ran a session with a proxy voting agency.

November-December – development of recommendations. The research team met in person to discuss the feedback from the focus groups, to consider initial findings from the quantitative research, and finalise the plan for the paper.

2025

January-April – development and checking of recommendations.

- The Insig AI team completed their quantitative analysis to include all reports published in calendar year 2024.
- The research team drafted the paper and circulated it amongst participants for comments; it was also sent to expert readers for their review.

May – publication: this paper was published on 6 May.

Quantitative research methodology

Objectives

We had two objectives:

1 To analyse patterns in public disclosure on the use of AI terms by FTSE companies in terms of:
   - **a.** Trends over time
   - **b.** Sector trends
   - **c.** Distribution of mentions across documents
   - **d.** Outliers.

2 To identify any notable mentions relating to the disclosure of the use of generative AI in relation to corporate reporting.

Scope of data sample

- All FTSE 350 companies' annual reports
- A range of other FTSE 350 corporate publications (integrated reports, sustainability reports, interim financials, other documents (such as policies, Codes of Conduct)
- Anecdotal samples taken from a range of other UK and international listed companies
- All documents published in PDF form on corporate websites
- All documents published in calendar years 2021-2024

Keyword groups

1 **General AI:** AI, artificial intelligence

2 **Generative AI:** generative AI, gen AI, genAI, large language model, GPT, ChatGPT, Copilot, Midjourney, Dall-e, chatbot (tested and excluded LLM, Co-pilot and Chatbot due to false positives/irrelevance)

3 **AI governance:** AI strategy, AI policy, AI governance, AI principles, responsible AI, AI working group, AI guidelines, AI committee, use of AI (tested and excluded 'acceptable use' due to false positives)

Process

1 Obeying the parameters set out in the scope, we downloaded all sentences containing any of the keywords defined above from Insig AI's ESG research platform, which contains a searchable database of corporate documentation that has been collected and converted from PDF to machine-readable format.

2 We then validated a sample of keyword mentions for relevance by reading sentences in context, and removing false positives.

3. We collected the results in an Excel spreadsheet to model and visualise them.

4 We investigated notable mentions, both manually and then with ChatGPT to test the results.

Qualitative research methodology

How focus groups were run

All sessions were held online (on Teams) and lasted one hour.

- **Pre-session briefing:** ahead of each session, the chair, Claire Bodanis, sent a briefing note and questions to all participants.
- **First 10-15 minutes:** Claire presented the research team's thoughts on the potential for generative AI to be used responsibly in corporate reporting, including ideas of the benefits and risks.
- **Remainder of the session:** open discussion amongst participants under the Chatham House Rule.
- **Contributions:** all participants contributed.
- **Consistency:** all corporate focus groups received the same set of briefing questions and the same presentation; both investor focus groups received the same set of briefing questions and the same presentation.

Briefing note sent to company focus group participants

Purpose of the session

- To gather insight into your understanding of how generative AI is being used inside companies, and any plans for integrating generative AI into corporate processes, including, if relevant, the reporting process
- To gather insight into the generative AI tools being provided by corporate IT and/or data teams, and any accompanying policies and governance
- To gather your views on how generative AI might be used, and how it might help/hinder the reporting process
- Chatham House Rule: just to remind you that this is under CH Rule; while we'll be using the insights, everything will be anonymised, and you will have an opportunity to review/comment on our output. So please be open/honest!

Agenda (1 hour)

- Intro/scene setting – reminder of research aims; introduce generative AI and possible use cases in reporting (facilitator 10 mins)
- Quick round the table comment from each on individual's/company's status re gen AI (10 mins)
- Discussion, using the questions below as a starting point, but please feel free to go off piste (35 mins approx)
- Conclusions/summary – what to expect/what's next (facilitator, 5 mins)

Our thanks and acknowledgements

In total, 60 people responded to the research. These included five institutional investors, one proxy agency, and representatives from 40 companies (including 20 FTSE 100s) that produce reporting. We also spoke to a small number of others from the governance and advisory sphere, including one participant from the UK's Chartered Governance Institute (CGI). An observer from the CGI sat in on some of the focus groups, and a representative of a UK investor body joined one of the investor groups. Six experts from reporting, governance and the investment community read and commented on this paper. Five students from the MSc in Environmental Technology at Imperial College, London, were involved in the research plan: Denise Bartel, Georgia Lavelle, Hebe Morley-Fletcher, Thaïs Ricard and Helena Uthoff. All of us from the Falcon Windsor and Insig AI research team are enormously grateful to everyone who gave their time and thoughts to this research.

Qualitative research methodology continued

Questions

- Have you used/are you routinely using generative AI tools either personally or at work, e.g. ChatGPT, Claude, Perplexity, Copilot?
- What's the status of AI usage inside your company? How confident are you in answering that question? Does your company have an AI policy and governance around AI usage? Do you know what AI tools you are allowed to use and in what contexts? Are there any restrictions? Have you had any training?
- Are you using or planning to use generative AI tools in preparing your reporting (or don't know yet)? And why?
- How do you feel about the integration of AI tools into the corporate environment?
- And specifically, how would you feel about AI tools being used to create the content you rely on/are responsible for?
- How do you think generative AI could help you in doing reporting better? Are there any drawbacks?
- Do you think guidance or regulation on the use of AI in reporting, and its disclosure, would be useful, or is this a matter for companies to decide?
- Any other points you'd like to make.

Briefing note sent to investor focus group participants

Purpose of the session

- To gather insight into your views as investors of the use of gen AI within companies and what role it has to play (or not) within corporate reporting
- To gather insight into how you may be using gen AI within your own companies, and any accompanying policies and governance
- To gather your views on how generative AI might be used, and how it might help/hinder the reporting process
- Chatham House Rule: just to remind you that this is under CH Rule; while we'll be using the insights, everything will be anonymised, and you will have an opportunity to review/comment on our output. So please be open/honest!

Agenda (1 hour)

- Intro/scene setting – reminder of research aims; introduce generative AI and possible use cases in reporting (facilitator 10 mins)
- Quick round the table comment from each on your knowledge of and views on gen AI and its use (10 mins)
- Discussion, using the questions below as a starting point, but please feel free to go off piste (15 mins)
- Insights from the corporate research so far (facilitator 5 mins)
- Discussions/reflections on those insights (15 mins)
- Conclusions/summary – what to expect/what's next (facilitator, 5 mins)

Questions

- Have you used/are you routinely using generative AI tools either personally or at work, e.g. ChatGPT, Claude, Perplexity, Copilot? Do you use gen AI in synthesising information you get from companies?
- How much do you know about how gen AI is used inside the companies you might invest in?
- As investors would you view the information in an annual report, or any statement made by a company to the market, differently if you knew it had been written – or parts of it had been written – using generative AI?
- Do you have any concerns about companies using AI, particularly generative AI in creating reports/statements to the market?
- Do you think reporting could be improved by companies using gen AI?
- Do you think companies should disclose in the annual report if/how they have used gen AI in its creation?
- Does your view on the usage of gen AI in creating reporting vary by type of information included? In other words, would it be OK for some sections and not others?
- Do you think the use of gen AI in reporting should be regulated? Or guidance provided?
- How do you feel more generally about the integration of AI tools into the corporate environment?
- Any other points you'd like to make.

Appendix 4: About Falcon Windsor and Insig AI

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.

Trust me, I'm listed

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 in June 2020. The second edition, with updates on the ESG reporting landscape and the future of digital reporting, was published in October 2021.

Contact

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

Insig AI

Established in 2018 and listed on the AIM, Insig AI plc is a specialised fintech business that focuses on ESG data, machine learning and enabling AI.

A collection of experts in data management, sustainability and financial markets, Insig's team creates ESG solutions that transform complex datasets into AI-powered, actionable intelligence for decision-making.

We work with investors, consultants, regulators and companies to have impact across the sustainable investment information ecosystem.

Our solutions

1. **Machine-readable ESG reports:** instant and cost-effective access to a curated AI-ready database of financial and ESG reports, which enables you to extract meaningful insights from unstructured data and leverage AI for deeper analysis.
2. **ESG research platform:** an award-winning, powerful web-based tool that lets you search, filter, analyse, and benchmark ESG disclosures, giving you instant access to an AI toolkit to transform research.
3. **Automated analytics:** our AI-driven approach systematically processes vast amounts of ESG data, reducing manual workload while maintaining complete traceability. Every insight is linked back to original sources, empowering reliability and compliance.

Contact

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

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 this paper and its corresponding summary paper, with the exception, as discussed in the appendices, of the quantitative analysis of FTSE 350 documentation.

INSIG·AI

Falcon Windsor

May 2025