


What does AI mean for a responsible business?




The FT Moral Money Forum is supported by

Vontobel

“Rather than centring discussion on the potential of a minimised workforce, we should also focus on how these technological advances will reshape our industry and others”

WHITE & CASE

“To realise AI’s benefits safely and responsibly, boards must be able to navigate the associated legal/compliance, shareholder activism, ethical and reputational risks”



Foreword



Visit the editorial hub at ft.com/moral-money-forum

If anyone had any doubt as to the world-changing significance of artificial intelligence, last year’s one-sentence open letter from many of the top figures in the field should have opened their eyes. The letter, signed by more than 350 AI experts and executives, warned the world that AI posed a “risk of extinction” no less than other threats like pandemics and nuclear war.

For investors and business people who are still getting used to talking about environmental, social and governance risks, the threat of human extinction is hard to top. Yet advocates of AI point out that this area of technology could bring huge benefits for society, whether through improved healthcare and education or a boost to productivity.

The environmental trade-offs are no less stark. Start-ups are already putting AI to work tracking emissions, and optimists think this field of technology will be among our most powerful tools in the fight against climate change. Yet it will also mean a vast increase in demand for energy-hungry computing power, complicating the drive to mitigate carbon emissions from electricity generation.

Things are hardly simple on the governance front either. The board chaos around Sam Altman’s brief expulsion from the leadership of OpenAI in November highlighted the tensions around that company’s unusual governance structure, and raised questions around the best way to ensure appropriate checks and balances.

In all, then, this is a uniquely huge and complex subject to tackle for anyone pursuing a responsible approach to business and finance. Fortunately, Sarah Murray has shed light on the key challenges in this latest of her incisive Moral Money Forum reports, which you’ll find below, along with insights from our Forum partners. Thanks for reading.

Simon Mundy
Moral Money Editor
Financial Times



What does AI mean for a responsible business?

How to navigate the opportunities and challenges posed by a technology few can afford to ignore.

By Sarah Murray

It was what many called an iPhone moment: the launch in late 2022 of OpenAI's ChatGPT, an artificial intelligence tool with a humanlike ability to create content, answer personalised queries and even tell jokes. And it captured the public imagination. Suddenly, a foundation model – a machine learning model trained on massive data sets – thrust AI into the limelight.

But soon this latest chapter in AI's story was generating something else: concerns about its ability to spread misinformation and "hallucinate" by producing false facts. In the hands of business, many critics said, AI technologies would precipitate everything from data breaches to bias in hiring and widespread job losses.

"That breakthrough in the foundation model has got the attention," says Alexandra Reeve Givens, chief executive of the Center for Democracy & Technology, a Washington and Brussels-based digital rights advocacy group. "But we also have to focus on the wide range of use cases that businesses across the economy are grappling with."

The message for the corporate sector is clear: that any company claiming to be responsible must implement AI technologies without creating threats to society – or risks to the business itself, and the people who depend on it.

Companies appear to be getting the message. In our survey of FT Moral Money readers, 52 per cent saw loss of consumer trust as the biggest risk arising from irresponsible use of AI, while 43 per cent cited legal challenges.

"CEOs have to ensure AI is trustworthy," says Ken Chenault, former chief executive of American Express and co-chair of the Data & Trust Alliance, a non-profit consortium of large corporations that is developing standards and guidelines for responsible use of data and AI.

"AI and machine learning models are fundamentally different from previous information technologies," says Chenault. "This is a technology that continuously learns and evolves, but the underlying premises must be constantly tested and monitored."

Some have warned that inappropriate use of AI technologies could prevent companies from

meeting their promises around social and environmental challenges – not least because of AI's hefty carbon footprint, which arises from the energy consumed in training chatbots or producing content.

A 2020 analysis conducted by the journal Nature found that high energy use, along with a lack of transparency and poor safety and ethical standards, could cause AI to erect obstacles to meeting 59 of the 169 targets in the UN's Sustainable Development Goals.

However, the Nature research also brought positive news: that AI could help progress towards 134 of the SDG targets by enabling innovations in areas from sustainable food production to better access to health, clean water and renewable energy.

With its ability to analyse millions of data points at speed and to identify patterns that humans would miss, AI can certainly help to drive positive impact.

For example, by creating "digital twins", it can analyse data from sensors, along with historical and real-time data, to find energy and other efficiencies in building systems. It also offers speed in the development of everything from life-saving drugs to alternative materials for electric vehicle batteries that could reduce reliance on scarce resources such as lithium.

Some see AI as supercharging progress on climate goals through everything from enhancing electric grid efficiency to applying analytics to satellite imagery to map deforestation and carbon emissions in real time.

"It's a very big deal," says Mike Jackson, managing partner at San Francisco-based Earthshot Ventures, which invests in climate tech start-ups. "Things are going to change much faster than people realise – and that's going to be a significant boon for the climate."

With AI holding both promise and peril, the challenge for companies across all sectors will be to temper the instinct to race ahead with appropriate caution.

Businesses will need to commit to thorough testing of AI models, and introduce policies and procedures to address risks of accidental harm, increased inequity and something every organisation fears: loss of control.

Handle with care

In 2023, New York lawyer Steven Schwartz was ridiculed in court when it emerged that his brief included fake citations and opinions generated by ChatGPT. For Schwartz, the revelations were deeply embarrassing. But they also raised awareness of the fact that AI programs can make glaring errors, something that is worrying when considering their possible use in industries such as nuclear power or aviation, where mistakes can be fatal.

Even where physical safety is not at risk, AI can introduce bias into decisions such as who to hire, who to arrest or who to lend to. In healthcare, concerns range from data breaches to relying on models trained on data sets that ignore marginalised communities.

For companies, among the biggest risks of getting it wrong is losing public trust. When KPMG polled 1,000 US consumers on generative AI, 78 per cent agreed on the responsibility of organisations to develop and use the technology ethically – but only 48 per cent were confident they would do so.

"You're going in with a level of scepticism," says Carl Carande, US head of advisory at KPMG. "That's where the frameworks and safeguards are critical."

Approaches to AI governance will vary by sector and company size, but Carande sees certain principles as essential, including safety, security, transparency, accountability and data privacy. "That's consistent regardless of whatever sector you're in," he says.

In practical terms, a responsible approach to AI means not only creating the right frameworks and guidelines but also ensuring that data structures are secure, and that employees are given sufficient training in how to use data appropriately.

But responsible AI does not always mean reinventing the wheel. The UN Guiding Principles on Business and Human Rights provide a ready-made means of assessing AI's impact on individuals and communities, says Dunstan Allison-Hope, who leads the advisory group BSR's work on technology and human rights.

"There's been all kinds of efforts to create guidelines, policies

and codes around artificial intelligence, and they're good," he says. "But we suggest companies go back to the international human rights instruments and use them as a template."

Some have not yet implemented any governance structures at all. While 30 per cent of FT Moral Money readers said their organisations had introduced enterprise-wide guidelines on the ethical use of AI, 35 per cent said their organisations had not introduced any such measures.

Reid Blackman, founder and CEO of Virtue, an AI ethics consultancy, sees no excuse for inaction. A rigorous approach to AI does require companies to make change, which takes time and effort, he says. "But it's not expensive relative to everything else on their budget."

While some might turn to the services of consultancies like Virtue or products such as watsonx.governance, IBM's generative AI toolkit, another option is to build internal capabilities.

This was the approach at Walmart, which has a dedicated digital citizenship team of lawyers, compliance professionals, policy experts and technologists. "Given our scale, we often build things ourselves because the bespoke model is the only one that's going to work for our volume of decision making," says Nuala O'Connor, who leads the team.


Whether turning to internal or external resources, there is one element of a responsible approach to AI that so many agree on it has its own acronym: HITL, or human in the loop – the idea that human supervision must be present at every stage in the development and implementation of AI models.

"Let's not give up on human expertise and the ability to judge things," says Ivan Pollard, who as head of marketing and communications at The Conference Board leads the think-tank's development of online guidance on responsible AI.

For Walmart, putting humans front and centre also means treating AI systems used for, say, managing trucks and pallets differently from AI programs that can affect the rights and opportunities of employees. "Those tools have to go through a higher order of review process," says O'Connor.

Executives agree that businesses must start taking AI risks more seriously

Responses to the statement: 'Businesses are not taking the ethical impacts of AI technology on society seriously enough' (%)



Source: Gartner

A vendor in the loop

If companies are still grappling with how to manage AI responsibly, their efforts must extend beyond their own four walls. “The vast majority of companies won’t develop their own AI,” says Chenault. “So they need to ensure they have the right governance and controls in procurement.”

Without these controls, the exposure is both legal and reputational, says Reeve Givens from the Center for Democracy & Technology. “This is a hugely important piece of the AI governance puzzle – and not enough people are thinking about it,” she says. “Because it’s the downstream customers that will have the most at stake if something goes wrong.”

Not all organisations appear to be aware of this. When ranking the risks posed by the adoption of AI and big data, only 11 per cent of the 976 institutional investors polled in a 2022 CFA Institute survey highlighted reliance on third-party vendors.

It was for this reason that one of the first publications of the Data & Trust Alliance was a guide to evaluating the ability of human resources vendors to mitigate bias in algorithmic systems.

The evaluation includes questions on the data that vendors use to train their models and steps taken to detect and mitigate bias, as well as measures vendors have put in place to ensure their systems perform as intended – and what documentation is available to verify this.

The alliance focused on HR vendors for the guidance because many companies’ first foray into AI is for recruitment purposes. “But those guidelines could be adopted for other tech vendors,” says JoAnn Stonier, a member of the Data & Trust Alliance leadership council and chief data officer at Mastercard, which helped develop the guidelines.

“When we’re using third-party vendors, we interrogate them heavily,” she says. “Because we’re ultimately responsible for the outcome of their solutions.”

To make things even more complicated, because AI technologies learn and evolve, vendors cannot know what will happen to their models when trained on the data sets of their clients.

This means that vendor-customer partnerships need to be far more collaborative and long-lasting than in the past. “That will change the supply chain relationship,” says Reeve Givens. “They have a shared responsibility to get this right.”

Watchful eyes

Companies may be trying to demonstrate that they can be responsible stewards of AI technologies. But governments are not leaving it to chance. In fact, for once policymakers seem to be acting relatively swiftly to bring order to an emerging technology.

First out of the regulatory gate was the EU, which in December agreed on its wide-ranging Artificial Intelligence Act, which many see as the world’s toughest rules on AI.

While the UK’s version is still a work in progress, the AI Safety Summit, convened by Prime Minister Rishi Sunak in November, sent a signal that regulating the technology would be taken seriously.

A month earlier, US President Joe Biden sent a similar message in an executive order directing government agencies to ensure AI is safe, secure and trustworthy. “To realise the promise of AI and avoid the risk, we need to govern this technology, there’s no way around it,” Biden said at the time.

The desire to create safeguards around AI technologies has even prompted a rare moment of collaboration between the US and China. In January, Arati Prabhakar, director of the White House Office of Science and Technology Policy, told the FT the two countries had agreed to work together on mitigating the risks. “All around the world we’re seeing policymakers feel the need to respond,” says Reeve Givens. “I haven’t seen a moment that is as concentrated as this AI policy moment.”

However, she also points to gaps, particularly on standards. “We can’t expect the average manager of a factory or supermarket chain to run a deep analysis on how an AI system is working,” she says. “So what is the approach to certification? That’s a massive global conversation that needs to happen.”

Debates continue over whether new regulations are either appropriately tough or risk stifling innovation. Meanwhile, it appears that they have not yet had much impact on corporate behaviour, at least among FT Moral Money readers, 92 per cent of whom said they had not had to change their use of AI to meet emerging regulations or standards.

Yet there are signs that, having failed to act to prevent the worst effects of social media, policymakers are determined not to let the same thing happen with AI.

“If we let this horse get out of the barn, it will be even more difficult to contain than social media,” Richard Blumenthal, the Democratic senator from Connecticut, said in his opening remarks at a December hearing on AI legislation.

Investing with an AI lens

Regulators are not alone in keeping a watchful eye on how companies use AI. Investors are also starting to ask tough questions. For asset managers and asset owners, responsible AI is partly about building internal governance systems. But it also means finding out whether the companies in their portfolios are using AI responsibly – particularly when investors are applying environmental, social and governance criteria to those portfolios.

“In pretty much every ESG conversation I have, AI is a topic,” says Caroline Conway, an ESG analyst at Wellington Management. “And mostly what I’m trying to get at is governance – how well the company is doing at managing the risk, pursuing the potential benefits and thinking about the trade-off between benefit and risk at a high level.”

Yet if FT Moral Money readers are anything to go by, it is early days for investors: only 19 per cent who identified as corporate executives said investors were asking their company about the use of AI. And 63 per cent of investors in the same survey said AI use did not affect decisions on whether or not to invest in companies.

The responses are perhaps unsurprising given the difficulties investors face in assessing the risks AI poses to

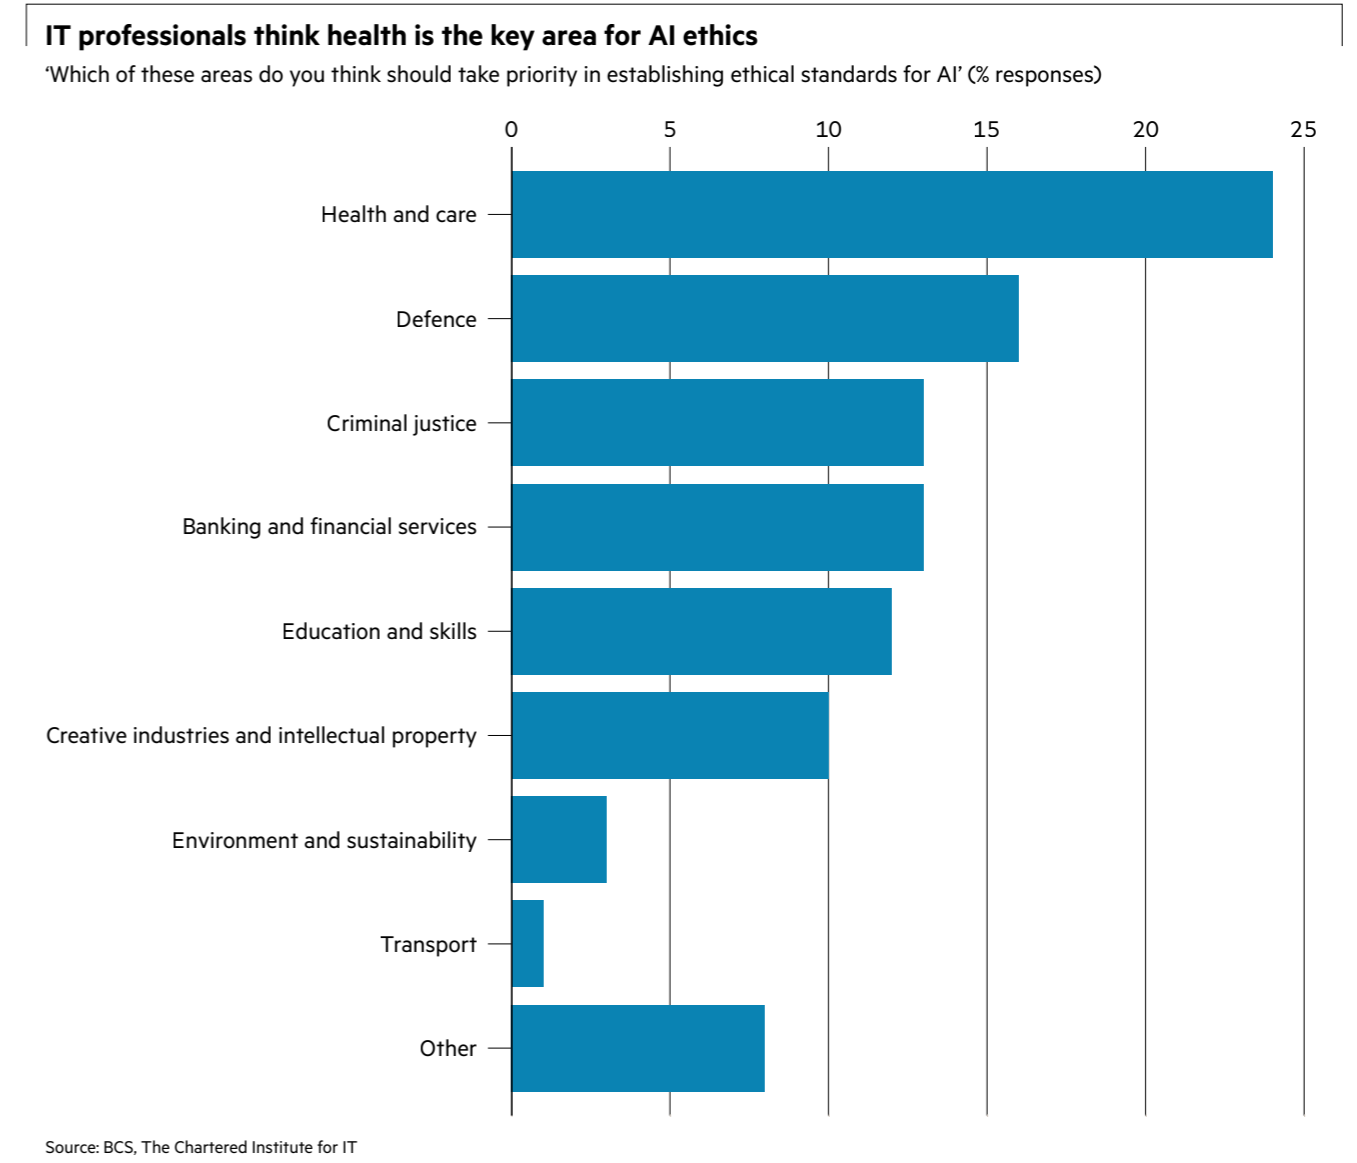
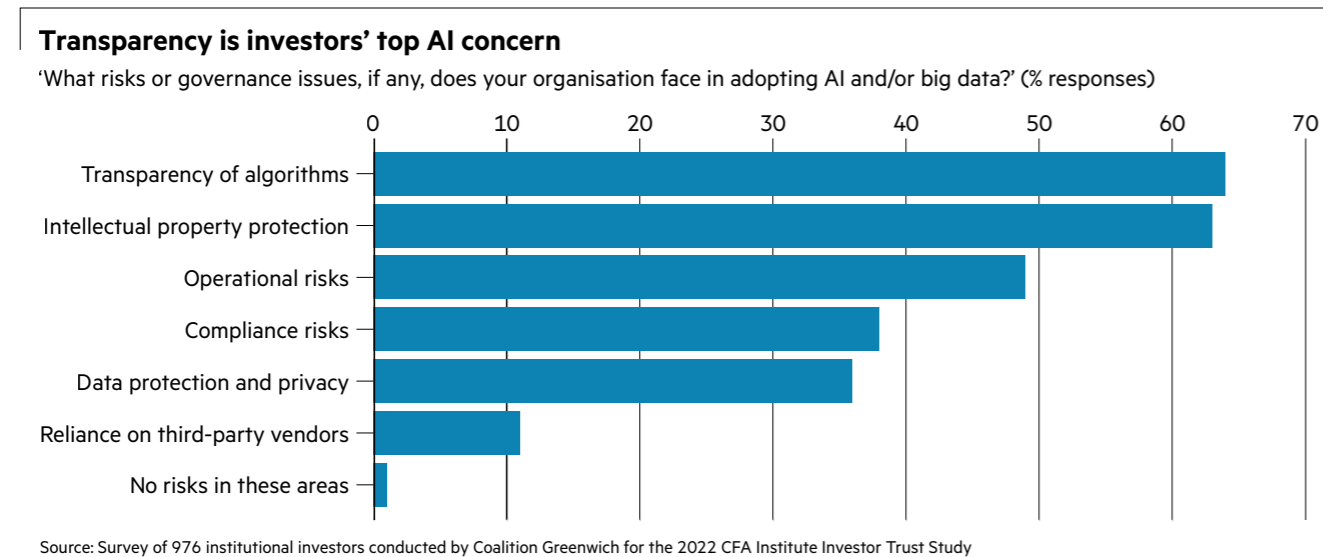
portfolios. “They are seeking basic understanding of how it can be used, which few of them and us truly have, to be honest,” one reader told us.

To help investors navigate this new risk landscape, a group of asset managers has formed Investors for a Sustainable Digital Economy, an initiative to pool resources and generate research on digital best practices in asset management. Members include Sands Capital, Baillie Gifford, and Zouk Capital and asset owners such as the Church Commissioners for England.

Karin Riechenberg, director of stewardship at Sands Capital, suggests investors start by identifying high-risk sectors, which range from technology, healthcare, financial services to hiring and defence.

Then, she says, they should identify high-risk use cases – those where AI will have a significant impact on aspects of people’s lives, such as credit scores, safety features in self-driving cars, chatbots and surveillance and hiring technologies.

“It’s important to look at each company individually and ask what AI tools they are using, what they are intended for and who might be affected by them and how,” she says.



Serving people and planet

If AI technologies are helping to measure social and environmental impact, they are also enabling innovators to create businesses that drive positive change in everything from healthcare to clean technology.

“We see it as a really amazing tool for engineers,” says Jackson of Earthshot Ventures. “It allows us to tease out correlations, to run through millions of simulations much faster and to model things in software before building them in hardware or biology.”

Given these capabilities, it is no surprise that AI technologies are permeating the portfolios of impact-focused venture capitalists and accelerators.

Jackson says AI is being used by almost every company in its portfolio and is at the core of the strategy for at least one-third. The same is true of the portfolio companies at Hawaii-based Elemental Excelsior, says Dawn Lippert, its founder and CEO.

Jackson points to Mitra Chem, which is using machine learning to speed up the development of the iron-based cathode materials needed in energy storage and transport electrification. The company says its technology and processes cut lab-to-market time by about 90 per cent.

Also in the portfolio is California-based KoBold Metals, backed by Bill Gates and Jeff Bezos. The company uses AI to scrape the world’s geological data (even including old hand-painted maps on linen) and deploys algorithms to find deposits of minerals such as lithium, nickel, copper and cobalt.

“To facilitate the transition to electric vehicles, we’re going to need to find a lot more of these resources,” explains Jackson. “Through that ingestion of a tremendous amount of data, AI is helping predict where these resources might be.”

Decarbonising the economy also involves making better use of existing resources — something AI technologies are particularly good at.

The technologies can be used to optimise energy use in buildings or adjust traffic lights to keep cars on the move

rather than idling. “AI technologies find those marginal gains — and they find so many of them that the cumulative value is massive,” says Solitaire Townsend, co-founder of sustainability consultancy Futerra.

AI can also help keep valuable resources in circulation for longer. For example, San Francisco-based Glacier, one of Elemental’s portfolio companies, is using AI technologies to bring greater efficiency and precision to waste sorting, a job for which it is hard to find human workers.

Equipped with computer vision and AI, its robots can identify and remove more than 30 recyclable materials from general waste at 45 picks per minute, a speed far greater than even legions of human workers could achieve. “Recycled aluminium, for instance, generates about 95 cent per cent fewer emissions than new aluminium,” says Lippert, who is also founding partner at Earthshot Ventures. “So it has a huge climate impact.”

By enabling new efficiencies, AI is also spawning a generation of young businesses that aim to expand access to essential services. At 25madison, a New York-based venture capital firm, the portfolio includes companies in the healthcare sector that are using AI to drive operational efficiency.

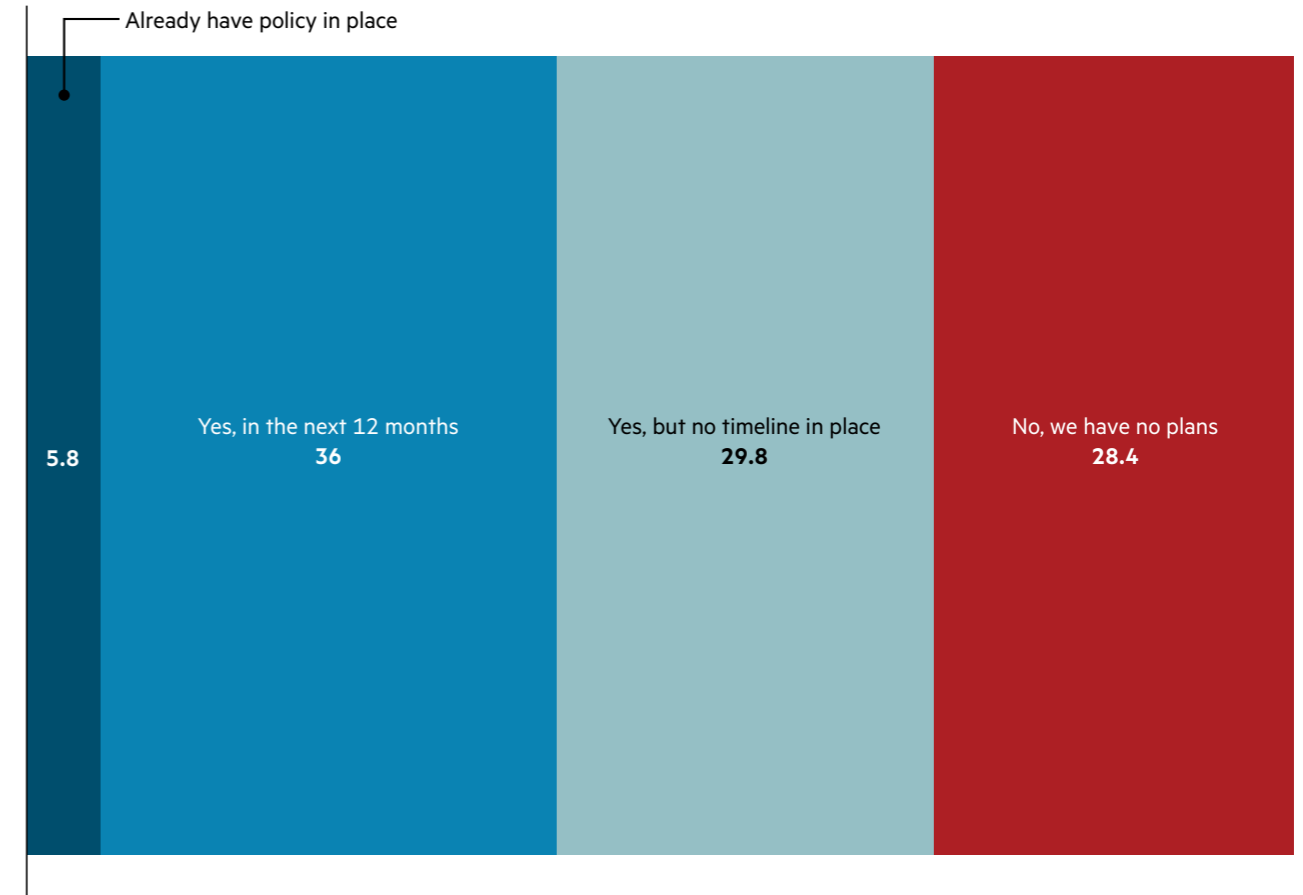
They include Midi, a virtual clinic specialising in perimenopause and menopause that uses AI to manage patient records and billing. The start-up aims to fill the large gap in access that women have to this kind of care, explains Jaja Liao, a principal at 25m Ventures, a fund at 25madison that invests in early-stage companies.

She says AI relieves specialists of time-consuming administrative tasks allowing them to spend more time with patients. “That’s how they make care more equitable.”

As these and other companies are demonstrating, AI technologies can be used for good. But as is the case with KoBold Metals, now valued at \$1.15bn, using AI to benefit people and the planet can also create highly successful businesses.

Many companies have been slow to develop internal AI rules

‘Is your company planning to enact rules or guidelines regarding the ethical and responsible use of AI?’ (% of responses)



Source: Conversica

Where AI meets ESG

ESG ratings have frequently come under fire for being inconsistent, unreliable and part of a confusing “alphabet soup” of acronyms. Now, however, two more letters of the alphabet — A and I — offer assessment tools that some believe could transform the way investors evaluate the ESG credentials of the companies in their portfolios.

Axa Investment Managers, for example, has developed a natural language processing tool that the firm runs over large volumes of corporate documents, including sustainability reports, to enable analysts to assess whether companies’ business activities are helping advance the UN’s SDGs.

“AI can bring super-useful solutions in digesting huge quantities of data,” says Théo Kotula, an ESG analyst at the firm. “That’s not to say it will replace ESG analysts. But it could make our jobs easier and quicker.”

FT Moral Money readers agree. When asked to select the biggest benefits of AI to their organisation’s sustainability goals, the largest group picked the ability to measure and track their positive or negative social and environmental impact.

AI could also improve ESG decision-making for asset managers by incorporating a far broader set of data points. These range from news reports, blogs and social media to data from satellites and sensors that can monitor pollution, deforestation and water scarcity in real time.

At Amundi Investment Institute, the research arm of the Amundi group, Marie Brière says AI harnesses these new forms of data to assess companies’ environmental impact, physical risks, social controversies and potential costs while also uncovering greenwashing.

“You could do this before,” says Brière, who is head of investor intelligence and academic partnerships at the institute. “But it’s now much quicker and uses quantitative tools.”

Moving fast and slow

AI may be ushering in an exciting new era in technological innovation and potential solutions to social and environmental challenges. But as the University of Oxford’s Colin Mayer points out, it is also a gold rush with similarities to previous booms.

“At the moment it’s clear the motive is to become as profitable as possible,” says Mayer, who has spent many years exploring the purpose of business. “The only way to solve this is to align the interests of companies with what we as humans and societies want.”

But with corporate leaders anxious to seize opportunities ahead of the competition, is this alignment possible? “There’s pressure to get it done first,” says The Conference Board’s Pollard. “But with that comes risk — the risk of doing the wrong thing with the wrong tool in the wrong way.”

And while many organisations have appointed chief ethics officers to maintain ethical behaviour and regulatory compliance, they may need to go further. One solution, says Virtue’s Blackman, is to put someone in charge of responsible approaches to AI. “If you’re the chief innovation officer, you want to move fast, but if you’re the chief ethics officer, you don’t want to break things — so there’s tension,” he says. “Someone with a dedicated role doesn’t have that conflict of interest.”

And while large, well-established companies may need to do some organisational retrofitting to put appropriate guardrails around their use of AI, young companies have an opportunity to get it right from the start.

This is something Responsible Innovation Labs, a coalition of founders and investors, is promoting among the next generation of high-growth tech companies. “Responsible AI should be an essential mindset and operating norm in the earliest stage of company building,” says Gaurab Bansal, executive director of the San Francisco-based non-profit.

For Bansal, the right approach is to assess the potential impact of products and technologies on customers and society more broadly. “We think responsible innovation is about designing and accounting for that,” he says. “It’s not about putting your head in the sand or worrying about it some other time.”

Unfortunately, as sluggish progress on meeting climate goals has shown, putting its head in the sand is something business does all too well. The question is whether it will take the same approach with AI. Or can capitalism harness AI for good and use awareness of its risks to prioritise long-term thinking over short-term gain?

So far, the jury is out. Yet there is a sense that, at this early stage of what is expected to be the next great tech revolution, this is a moment when it is still possible to get the governance right.

“We’ll constantly have to tweak it,” says Riechenberg at Sands Capital. “But if we start doing that now, we have the potential to make the most of this technology — to control it and not be controlled by it.”

Advisory Partners

The FT Moral Money Forum is supported by its advisory partners, *Vontobel* and *White & Case*. They help to fund the reports.

The partners share their business perspectives on the forum advisory board. They discuss topics that the forum should cover but the final decision rests with the editorial director. The reports are written and edited by Financial Times journalists and are editorially independent.

Our partners feature in the following pages. Partners' views stand alone. They are separate from each other, the FT and the *FT Moral Money Forum*.

AI's transformative role in asset management and ESG

Christel Rendu de Lint

Artificial Intelligence, one of the most hotly discussed and debated topics of late, continues to advance its technology and capabilities, as well as its impact on multiple facets of business and society. It looks set to affect nearly every industry and has the potential to transform society, yet an open question remains: will it be a net positive or negative overall?

While well-documented fears surround AI, its many potential benefits often attract less focus. As with much of technology, whether it's ultimately advantageous or destructive hinges upon how we manage its emergence and integrate it into our lives.

Encouraging transparency and reducing bias

Introduced carefully — ie with an understanding of its potential drawbacks — AI can lead to much greater transparency and the reduction of bias in decision-making. As AI systems can inadvertently reflect current human biases, and lead to harmful results known as algorithmic bias, developers, owners and users need to urgently prioritise both the awareness and mitigation of this risk. Possible measures include establishing responsible processes, the use of technical tools, and operational practices such as internal "red teams" or third-party audits. Luckily, as awareness of inherent bias and its risks increase, the threat of AI unintentionally reversing progress in crucial areas such as workplace diversity lessens.

If AI is to improve transparency and reduce bias, it must be trusted. Demand from regulators and the public for explanations and clarity around data use is growing. Transparent AI relies upon understandable and clean inputs, making it easier for humans to understand and trust AI decisions, and forming the basis for testing and explaining those decisions to stakeholders. At the policy level, both US President Joe Biden's executive order of October 2023 and the European parliament's Artificial Intelligence Act of December 2023 focus on transparent and trustworthy usage of new technology.

AI's contribution to ESG

A marked opportunity exists for AI to play a transformative role in the world of ESG. Again, success depends on how we guide the technology; namely, we must route it towards making a positive difference. An example is the AI-driven development of technology that aims to remove atmospheric greenhouse gases and help mitigate climate risk. AI's integration of resources into climate change predictions can increase accuracy, help measure and monitor emissions, and provide essential data for climate action. In addition, smart manufacturing can enable significant reductions in energy consumption and waste and carbon emissions, while AI-based forecasts can improve electrical system efficiency via accurate supply and demand predictions.

Boston Consulting Group and Google estimate that with the thoughtful guidance of experts, AI could help mitigate 5-10 per cent of global greenhouse gas emissions by 2030, about 10-20 per cent of the Intergovernmental Panel on Climate Change's interim target for achieving net zero by 2050. Implemented correctly, we believe AI's impact on our journey towards increased sustainability includes faster technological experimentation and, therefore, hopefully also faster than expected breakthroughs.

AI's transformative effect in the investment space

Despite concerns of AI-related workforce reductions in the banking and investment sector, many argue that AI will complement people at work, not replace them. Daron Acemoglu, a professor of economics at MIT, asserts that when AI and humans work together, they can do better than either would alone.

As computers get smarter, their potential to automate complex, repetitive tasks — such as trade execution, performance reporting and compliance monitoring — increases. AI can therefore help identify and implement opportunities for cost savings and operational improvements. This can result in investment professionals having more time for higher-level analysis and strategic decision-making, which is at the core of solid portfolio management. Its potential to process and analyse vast amounts of data at speeds unattainable by humans means AI can enable more sophisticated market analysis, leading to better-informed investment decisions. When it comes to client management, AI can analyse individual client data to provide highly personalised investment advice and recommendations, potentially increasing client satisfaction and retention.

Rather than centring discussion on the potential of a minimised workforce, we should also focus on how these technological advances will reshape our industry and others. The benefits of AI will probably inspire a skill shift: as AI becomes more integrated into the investment management industry, demand will grow for professionals skilled in data science, machine learning and AI implementation. Nonetheless, similar to the medical field, when dealing with matters of sensitive financial planning and investment, it's hard to imagine a world in which chatbot interaction would completely replace human interaction.

So, while AI will enhance efficiency and systematics, it's unlikely to eliminate the human touch in our sector and others. It appears, therefore, that the future we are looking at is one where humans and technology will complement one another. We believe that the careful implementation of AI, with an understanding of its risk landscape, should produce a net positive result for the world of investments and beyond.

Advisory Partner

WHITE & CASE

Legal perspectives on developing and deploying responsible AI

Janina Moutia-Bloom and Tim Hickman

Despite previous calls for an “ethical pause” on AI development, corporate investment in it is expected to reach \$200bn by 2025. As such, board oversight in this area continues to evolve at pace, particularly in light of concerns raised regarding AI including: the potential for misuse of AI technologies; bias amplification; potential lack of transparency and accountability; and the treatment of personal data and intellectual property used in AI systems.

The growing demand for, and availability, impact and stakeholder scrutiny of generative AI is pushing it higher up the board agenda across all sectors. Businesses stand to benefit from embracing AI — for example, it could improve efficiency, assist decision-making and contribute to risk management. Yet, to realise AI's benefits safely and responsibly, boards must be able to navigate the associated legal/compliance, shareholder activism, ethical and reputational risks.

AI is technically complex and fast-moving, and therefore challenging for governments to develop effective regulation, standards and/or guidance. As a result, the international landscape of legal frameworks governing AI is fragmented — with even the definition of AI differing across jurisdictions. However, a new, global phase of AI regulation is starting to emerge, as indicated by the publication of the G7 AI principles and Hiroshima AI Process, the Bletchley Declaration on AI safety, the Blueprint for an AI Bill of Rights, the “first-of-a-kind” Framework Convention on AI, Human Rights, Democracy and the Rule of Law, the UN's B-Tech Project on Generative AI with the UN Guiding Principles on Business and Human Rights, and the UN General Assembly's recently adopted resolution on “safe, secure and trustworthy” AI that will also benefit sustainable development for all (backed by more than 120 States).

The EU's proposed AI Act is designed to provide a horizontal legal framework for the regulation of AI systems across the EU. Once in force, the AI Act's risk-based approach — which has fundamental rights protection at its core — will have global reach and affect actors across the entire AI value chain. However, several of the concepts set out in the AI Act will require clarification by courts and regulators to provide businesses with greater certainty regarding their compliance obligations. Alongside the AI Act, companies operating in the EU must still consider obligations under other applicable sector-specific instruments, such as the General Data Protection Regulation, the Digital Services Act, and the forthcoming AI Liability Directive. Companies should also be aware of regulatory initiatives at national level in the Member States in which they operate. For example, in February 2024, France's competition authority announced that it would investigate big tech's competitive functioning in the generative AI sector, and would issue an opinion in the coming months.

The UK has taken a different approach to the EU, declining to issue new legislation at this stage, and instead adopting a flexible framework of AI Regulatory Principles that will be enforced by

existing regulators. This framework is intended to be both pro-innovation and pro-safety. In February 2024, a Committee of the House of Lords published a report cautioning the Government against a regulatory approach too narrowly focused on AI safety. Days later, the Government published: (i) its consultation response to its AI Regulation White Paper, articulating a principles-based (rather than the EU's risk-based) approach towards regulating AI; and (ii) guidance for regulators on implementing the AI Regulatory Principles.

On the other side of the pond, in late 2023, the Biden administration signed Executive Order 14110 on Safe, Secure, and Trustworthy Development and Use of AI. However, in contrast to the EU's risk-based regulatory approach, the Order places near-equal emphasis on the pressing need to (responsibly) develop and harness the potential benefits of AI as it does on the need to understand and mitigate novel risks. Initiatives also unfold at State-level, including California Senator Scott Wiener's recent proposal for sweeping safety measures for AI in SB 1047, while New York City has already introduced Local Law 144 to regulate the use of AI in hiring decisions.

Shareholders have also started to become active, with AI-focused resolutions having made their debut in the US (for example, calling for tech and motion picture companies to publish an “AI transparency report”). Such resolutions are expected to feature more prominently at future AGMs.

Boards should also be alive to the types of AI-related disputes and class actions being filed in national courts, as judgments in these early legal actions will be instructive in evaluating a company's potential exposure to litigation risk. Disputes have already arisen in relation to issues such as whether training data fed into AI systems infringes copyright, other IP rights and/or personal data protections; alleged bias in the output of AI tools; misrepresentation of AI systems' capabilities; and whether an AI system can itself be an “inventor” under patent law or an “author” under copyright law.

To mitigate the risks explored above, companies should implement effective AI governance. This may involve:

1. developing clear and robust policies which govern — and embed ethical practices into — the use of AI;
2. developing strategies for negotiating AI-specific contractual clauses, including in relation to policies, procedures and testing with respect to the various concerns, and liability attribution for AI failures;
3. establishing a cross-functional team of specialists from legal, compliance, ethics, data science, marketing (among others) to oversee and report to management on AI governance; and
4. undertaking regular risk assessments and audits of AI models and data sets to remediate legal and ethical concerns (e.g., bias).

About the FT Moral Money Forum

The FT Moral Money Forum takes key issues from the ESG debate and explores them for FT Moral Money subscribers.

The forum highlights macro and philosophical questions and explores the experiences and solutions being proposed. We apply an editorial filter to these and present the most interesting ideas and experiences. We also engage our data visual team to find the best form of presentation.

The forum produces regular reports to highlight the ideas, policies and practices that are making a difference.

Find out how to take part in the FT Moral Money Forum by emailing moralmoneyforum@ft.com

[ft.com/
moral-money-forum](https://ft.com/moral-money-forum)

