

Technology governance in the face of exponential change

To manage risk without impeding innovation or growth, companies need strong technology governance. But this requires directors, whether or not they are in technology businesses, to keep up to date on emerging technologies. Artificial intelligence (AI) is the most recent technology to disrupt economies, companies, and the work environment. Its capabilities and applications are advancing at a rate that even its developers find surprising, and investment in AI is keeping pace. Directors are keen to learn more about the business opportunities that AI creates as well as the risks it introduces.

On July 11, 2023, members of the European Growth Audit Network (EGAN) met with Dr. Harvey Lewis, partner in Client Technology and Innovation at EY, to discuss recent advancements in technology and practical steps for fast-growth companies to build strong technology governance.

This *ViewPoints* synthesizes discussions around four themes that emerged during the meeting and premeeting conversations:¹

- **Generative AI has emerged faster than anticipated.**
- **Current and future opportunities will significantly change the business landscape.**
- **Boards need to familiarize themselves with technology risks.**
- **Companies should build strong technology governance frameworks to manage risks and foster growth.**

For a full list of meeting participants, see Appendix 1 (page 10); for a list of reflection questions for audit committees, see Appendix 2 (page 11); for guest biography, see Appendix 3 (page 12).

Generative AI has emerged faster than anticipated

The explosion of AI in news headlines could give the impression that it is a recent creation, but its origins began in the 1950s and its development was long and slow. Dr. Lewis said, “*The first practical AI program was presented at a 1956 conference at Dartmouth College. The technology developed slower than anyone anticipated: there have been 40–50 years of AI “summers” and “winters.” But everything changed on November 30, 2022, when OpenAI released ChatGPT.*”

Dr. Lewis set the stage by describing OpenAI's GPT-4 (Generative Pre-trained Transformer 4), the latest version of ChatGPT: *"The world's most powerful large-language model, it includes all the text content of the World Wide Web, plus thousands of books. The computer power needed to train these models is astonishing: GPT-4 probably required 20,000 of NVIDIA's latest graphic processing units, running for 400–500 days, for a total cost between \$200 and \$300 million."* He explained that the technology gets its name from mimicking human intelligence—or thinking in a more human way—and it is generative AI's large and intricate neural networks, to some extent modeled on the human brain, that enable it to complete complex tasks. *"Artificial neural networks are used in the field of AI to teach machines to spot patterns in data and apply them to make predictions or generate content. The GPT-4 neural network has 1.6 trillion artificial neurons. Brains have 85 billion biological neurons. There is no direct comparison between artificial and biological neurons, but it gives an idea of the complexity of the models,"* Dr. Lewis said.

Generative AI has garnered public attention faster than even its most enthusiastic proponents anticipated. ChatGPT is estimated to have enrolled 100 million users within two months of its launch, whereas it took TikTok nine months and Instagram two and a half years to reach similar usage levels.² Members noted the difficulty of keeping up with developments. One said, *"Recent conversations at the board have really centered around getting visibility on what's going on in technology."* The pace of recent development has astonished even the developers. Dr. Lewis said, *"There have been two characteristics of technology over the last year. One is speed of development: the pace by which capability is advancing is utterly astonishing. I've worked 30 years in this space, and this is the most rapid, disruptive technology I've seen. The second characteristic is the capability itself to answer the questions—acting in a way reminiscent of the way people act."*

Current and future opportunities will significantly change the business landscape

Members wondered about the opportunities that AI offers fast-growth companies. As one member put it, *"This technology is going to impact us. How are we preparing to take advantage of it?"* Members discussed with Dr. Lewis how AI is positively disrupting the workforce and how companies can leverage it. The following observations emerged:

- **Opportunity for productivity growth is greater than during the industrial revolution.** Using generative AI at work can help employees save time by automating manual and repetitive tasks, interpreting complex data sets, drafting documents, and quickly researching topics. Members discussed with Dr. Lewis how these levels of productivity could significantly alter the workforce. *"These are tools that can enhance productivity to levels not seen since the close of the industrial revolution, and that took 150 years for 25% growth. We're seeing productivity increase up to 50% in some industries almost overnight,"* Dr. Lewis said.

- **Potential for automation is vast and far reaching.** Generative AI’s capabilities to automate tasks that are historically performed by humans creates wide-ranging possibilities. Dr. Lewis explained that within two weeks of launching GPT-4, developers had been able to automate the work of multiple departments. He said, *“With GPT-4, we have a range of developers who chained a sequence of models together to perform much more complex tasks. We see the emergence of AI systems capable of fulfilling complex workflows. For example, I’ve worked with developers who have created sophisticated organizations entirely populated by bots powered by large language models: bots for HR, finance, marketing, governance, and product development. The ‘owner’ types prompts to start product development and a couple of weeks later, it’s built—the finance is done, the social media is written; it’s all there. This was within two weeks of the launch of GPT-4.”*
- **Future opportunities are already visible.** Members discussed with Dr. Lewis the possibilities of AI in the near future and how this could impact their companies and the workforce. Reflecting on the pace of development, one audit chair said, *“We’ve just started this; I don’t think we have a clue what that topic will be in 10 or even five years from now. We have just started to embrace the wave.”* Dr. Lewis noted several future possibilities:
 - **Generative AI could create strategies to find the best responses instead of the quickest.** Dr. Lewis described how, by the end of the year, he believes workers will have the option to use an AI that strategically communicates with the user to find the optimum response, prioritizing the best answer over the fastest one available. He said, *“Later this year, I suspect one of the large tech firms will release a system that combines large-language model and deep-reinforcement learning. The characteristic of large-language models is it’s tuned to give you the best response there and then; this new technology will treat conversation as strategy. If you have an objective to achieve, it’ll take you through a conversational journey to arrive at that outcome instead of giving you a response straight away. In the context of workflows and processes and systems, you can see how you can use these kinds of AI systems for automating a wide range of tasks.”*
 - **Generative AI will be able to complete large and complex documents from scratch within the next year.** Generative AI is already capable of completing complex tasks and producing various types of written content, but Dr. Lewis suspects that it will soon be able to create completed documents. Dr. Lewis said, *“Within the next year, I suspect we’ll see generative AI systems completing finished documents—magazines, brochures, due-diligence reports, tax filings—just completing them from scratch. And creating high-definition photorealistic videos. Consumers could type a prompt like ‘90-minute feature-length film with my favorite characters with a plot of my choosing.’”*

- **Generative AI could help solve some of the world's biggest problems.** Dr. Lewis predicted that generative AI's ability to automate complex and time-consuming tasks will lead to decreased costs, especially on the margin. He emphasized that the resources regained from this could be used to apply AI to a broader range of problems: *"If we can use AI to find drugs within weeks, create new materials for renewable assets more quickly, and get ordinary work done in 50% less time, we will have more time for solving the big problems, plus we'll have AI helping us. Who's looking at the economic impact there? I don't think many people are."* One member said with some irony, *"We'll solve world hunger through AI."* Dr. Lewis replied, *"Hopefully."*

Boards should also be observing and preparing for quantum computing

Members and Dr. Lewis discussed recent developments in quantum computing and the possible impacts for businesses. Dr. Lewis highlighted the computing power of quantum: *"It can be in many states at once; it's very confusing. When you add one qubit to a quantum computer, it doubles the processing power of that one qubit, and it's doubling every 18 months. Classical computers double every 18 months; quantum is double doubling every 18 months."*

Dr. Lewis explained that quantum computing has seen substantial progress over the past year: *"In the last 12 months, we've seen quantum computers bridge the gap from purely theoretical to having utility."* He added, *"The landscape is starting to shift incredibly in the field of new computation. It will take a while for things to be properly advantageous and usable. But it also brings risks."* The biggest risk, he noted, is that fully operational quantum computers will solve the mathematical problems that underly modern cryptography on time scales far shorter than those of conventional computing—which means that encrypted data, even using large keys, could become vulnerable.

He advised that boards begin thinking now about how to stay ahead of advances in quantum computing: *"There are techniques to apply now, such as quantum key distribution, which make standard encryption far more quantum resistant. A lot of bad actors realize that quantum might not be able to do some of these things now but in five years from now, so they think, 'Why don't I grab your data now and decrypt it in five years when quantum comes along?' Companies must think now about quantum protection."*

Boards need to familiarize themselves with technology risks

Analysts estimate that generative AI can add trillions of dollars to the global economy, but it is crucial that management and boards carefully consider the risks associated with rapid

technology adoption.³ As Dr. Lewis put it, *“The challenge is walking the tightrope between the risk of missing opportunity and the risks that flow from using AI.”*

Members and Dr. Lewis identified several risks:

- **Regulation and ethics around using AI are complex and evolving.** The recent surge in AI use prompted a conversation on evolving AI regulations and how it can be difficult for companies to navigate these changes, particularly those that operate in various jurisdictions. Dr. Lewis said, *“We’re seeing governments around the world increasing their focus on regulation of AI. Right now, there’s a patchwork quilt of regulations. This presents challenges for multinationals or even those operating across different jurisdictions. The European AI Act is implemented in the Republic of Ireland, but the UK government is taking a different approach to regulation. How do organizations navigate that?”*

Reflecting on the speed at which AI is developing, one member questioned how to approach determining ethical use, as well as the longevity of any formal definition: *“Are there parameters? There don’t seem to be any. How do you begin to talk about ethical behavior and what it looks like? And is it for the next 30 days, months, or years?”* Dr. Lewis replied that the AI regulations currently proposed in jurisdictions around the world are *“built on a common set of principles from the OECD [Organization for Economic Co-operation and Development] that are adopted by almost every country in the world. These focus on topics such as fairness, privacy, and transparency. It’s important that boards consider the same principles when putting in place policies and guidelines.”*⁴

- **Ownership of data used to train models and the content it generates is a contentious legal issue.** While regulations lag behind the speed of AI development, establishing intellectual property rights is contentious, both for the data used to train AI models and for the words, pictures, videos, and music that the models produce. Dr. Lewis said, *“There’s currently a big legal debate about two things: The first is the ownership of data used to train models and whether its use infringes copyright law, and the second is ownership of the output. Who owns it? Maybe you’re liable for the output and no one wants to be liable when the model might ‘hallucinate.’ This is a legal liability risk and where that responsibility should fall is unclear.”* He added, *“Currently, large language models like GPT-4 are trained on web data. Therein lies a multitude of problems—not just the liability of using the data on the web but the ownership of data on the web. It’s a legal grey area on whether the responsibility is on organizations training models or those that use the models; it’s a very difficult area.”*
- **Generative AI can provide highly convincing but entirely false responses.** Owing primarily to mismatches between queries and training data, large-language models can produce not only false content but also fake references and supporting material that make the responses appear convincing. Dr. Lewis described a legal brief, produced by ChatGPT and submitted to a New York court without fact checking, which contained six completely bogus

case references. He noted that these “hallucinations” can cause significant issues for a company if the data is not verified: *“It raises real challenges, considering the general nature of the technology and all the ways it can be used in an organization, such as automation of documents, processes, writing proposals, etcetera.”*

- **AI is being used to generate fraudulent material, pictures, and video content.** AI can be a powerful enabler of fraud and other misconduct against organizations and individuals. New models, easily and cheaply accessible to malefactors, can generate believable documents, pictures, and videos. Early hopes of automated tools to detect fraudulent AI-produced content have been dashed. Dr. Lewis said, *“It takes a few tweaks of a prompt to make faked text completely undetectable. We’re living in a world where AI is creating imagery, text, voices, and videos that are indistinguishable from that written by humans.”* Dr. Lewis provided a real example to demonstrate the difficulty of recognizing AI-generated material, even for experts: *“You can already get photorealistic systems. One won the open category in the Sony World Photography Awards; the winner then announced that it was generated by AI.”*⁵

Companies should build strong technology governance frameworks to manage risks and foster growth

Members recognized the importance of technology governance, but they pointed out challenges in developing it. A member said, *“I think boards don’t really focus on technology governance, and it’s the much bigger elephant in the room. Technology is going to transform so much of the business world in the next 10–15 years. How do you get your head around that? How will it impact your business? Is it the board as a whole that has to be educated on how it’s going to impact things such as your business, your supply chains, your customers? How do you bring boards along to realize it’s an important emerging risk for them?”* Another member noted that their board is interested in technology governance, but progress is slow: *“It’s not that they don’t want to [build technology governance], but I think they’re playing a bit of catch-up. They’re not as far along as you would want or expect.”* Resource allocation, because of the size of the company, can be an obstacle. One member said, *“Small boards have small budgets. We’ve focused on discrete issues on technology and cyber risk before, but technology strategy isn’t something small boards have talked about.”*

Members and Dr. Lewis discussed good practices:

- **Boards should consider building on existing frameworks.** A challenging business environment can make it difficult to tackle areas that are not obviously broken, such as building strong technology governance. One member said, *“People have their heads in the sand a little bit. I think it’s because they’re busy running their business; it’s been a tough environment for businesses—the war, COVID, inflation. Thinking about technology risk, how is it going to bite us? They’re overwhelmed and fatigued; they’ve been busy with so many*

things. How do we tackle it? It's so important." Dr. Lewis recommended that boards consider building on existing frameworks to ease the workload: *"One challenge I always see is an organization creating new governance for what is emerging but not accounting for what is already there. New technologies bring new risks but also accentuate and exaggerate existing risks—and we have frameworks already for these."* He added, *"You've got some solid foundations in cyber and privacy; it's important not to throw the baby out with the bathwater. You've got the core aspects—pick up from here and expand on them. In the context of new technologies, consider how these increase risks like data sharing and privacy."* One member agreed: *"Companies have policies on social media; you can easily extend that for a framework on AI."*

- **Smaller boards should add AI to the agenda to understand the full extent of cybersecurity and data-protection risks.** Many EGAN members have tried ChatGPT, but the majority said that their boards were not discussing generative AI or that conversations had only started recently. Given the wide applicability and accessibility of generative AI, boards need to understand its evolving capabilities, stay informed on the latest developments, and understand the potential value of AI. Dr. Lewis said, *"I guarantee some of your employees use ChatGPT."* A recent report from the US Congressional Research Service noted, *"A user may reveal sensitive health information while conversing with a health care chatbot without realizing their information could be stored and used to re-train the models or for other commercial purposes. Many existing chatbots have terms of service that allow the company to reuse user data to 'develop and improve their services.'"*⁶ A member said, *"Cybersecurity and data protection is on the agenda every meeting and given high priority."* Members agreed that adding AI to the agenda will help boards understand the full extent of data-protection risks.

Bring in outside expertise. Members discussed the difficulty of keeping pace with the growing volume and complexities of risks that come from using AI. One member noted that the recent developments in generative AI make it difficult to stay educated on the topic: *"For cybersecurity, there's a system in place; risk remains, but there's a system. On AI, it's very different. I feel a lot of education needs to be done in the boardroom. I feel like we're not as deep as we were in the cloud discussion or cyber discussion. We need to elevate our board members on AI, data, and standards. There's a lot more work to be done."* Another member agreed and noted that inviting an expert to educate the board could help them understand the risks of AI and what questions to ask: *"We as a board really don't understand AI, we need someone to come in and talk to us about the risks. What should we be doing? What kind of questions should we ask, etc.?"*

An expert can help the board ask management better questions, make quicker decisions, and better oversee technology risk. The impact of an expert could be amplified in a fast-growth company that relies on a more hands-on approach from its board members. Some

EGAN members are considering adding IT experts to their boards. One asked, *“To be capable of providing oversight, do you need a technology person on the board—is that what is now required?”* One member highlighted that an expert helped their board to build technology frameworks within the company: *“We had an expert from a university come in and look at our business and what we were doing, but a lot of it was educating us as a board, what we’re responsible for, what questions to ask, what we should be looking for. It was so helpful to have him come in. I’m sure it was expensive, but he was a top expert and was really exactly what we needed.”*

- **Finding the balance between embracing change and moving too quickly involves “thinking fast and slow.”** Although *“12 months is an eternity in this field,”* said Dr. Lewis, companies should be cautious about acting too quickly. One member commented on this challenge: *“Companies need to find the balance between the dark side and the light side—what limitations I should enforce and what processes I can push to become more efficient. It’s a difficult balancing act.”* Another member described the pressure that exists when starting to use AI: *“We haven’t been talking about AI. There’s a lot of buzz broadly in the industry. I saw that if a company talked about AI in the earnings report, they did better—so almost everybody is talking about it. I’m not really with a tech company, so it’s a bit different. I’m sure there’s a role to play in my sector, but we’re not that active.”* While companies may be tempted to invest in the latest exciting technology, it is best to first spend time, rather than money, to understand the technology. Dr. Lewis said, *“One of the worst things to do is to have a knee-jerk reaction and invest a huge amount of time and money in something that quickly becomes obsolete or redundant. If you rush now, which I think a lot of organizations are doing, you’ll probably come up with the wrong answers. You need to do it in a way that’s informed and intelligent; that requires focus before massive investment takes place.”*

About this document

The European Growth Audit Network is a group of directors drawn from Europe's leading high-growth companies committed to a program of learning and problem-solving aimed at enhancing governance of Europe's fastest-growing companies. The network is organized and led by Tapestry Networks with the support of EY as part of its continuing commitment to board effectiveness and good governance.

ViewPoints is produced by Tapestry Networks to stimulate timely, substantive board discussions about the choices confronting audit committee members, management, and their advisers as they endeavor to fulfill their respective responsibilities to the investing public. The ultimate value of *ViewPoints* lies in its power to help all constituencies develop their own informed points of view on these important issues. Those who receive *ViewPoints* are encouraged to share it with others in their own networks. The more board members, members of management, and advisers who become systematically engaged in this dialogue, the more value will be created for all.

The perspectives presented in this document are the sole responsibility of Tapestry Networks and do not necessarily reflect the views of network members or participants, their affiliated organizations, or EY. Please consult your counselors for specific advice. EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Tapestry Networks and EY are independently owned and controlled organizations. This material is prepared and copyrighted by Tapestry Networks with all rights reserved. It may be reproduced and redistributed, but only in its entirety, including all copyright and trademark legends. Tapestry Networks and the associated logos are trademarks of Tapestry Networks, Inc. and EY and the associated logos are trademarks of EYGM Ltd.

Appendix 1: Meeting participants

The following members participated in all or part of the meeting:



Nadja Borisova
Audit Committee Chair, Blabla Car



Brenda Eprile
Audit Committee Chair, Atlantica Sustainable Infrastructure, plc. and Westport Fuel Systems, Inc.



Christoph Hütten
Audit Committee Chair, Brockhaus Technologies



Sandip Kapadia
Audit Committee Chair, VectivBio



Linda McGoldrick
Audit Committee Chair, Compass Pathways, Audit Committee Chair, Alvotech



Carolyn Schuetz
Audit Committee Chair, OakNorth Bank

EY was represented in all or part of the meeting by the following:



François Langlois
Managing Partner, Markets and Business Development, EMEIA Assurance, EY



Suwin Lee
Private Leader for Europe, the Middle East, India and Africa, EY

Appendix 2: Reflection questions for audit committees

- ? What challenges does your company face in technology governance?
- ? How does your board stay up to date on emerging technologies and their risks, as well as on new regulations?
- ? What good practices have you observed for board oversight of technology opportunities and risk?
- ? What, if any, technology expertise do you have on your board? How does this affect IT governance and risk oversight?
- ? How is your company thinking about the impacts of emerging technologies, especially AI, on its operations and workforce? What policies has your company implemented related to employee use of AI for the generation of work products?

Appendix 3: Guest biography

Harvey Lewis is a senior technology leader with a diverse background spanning rocket science, data science, and research. He has 30 years of experience in artificial intelligence and other emerging technologies. He is a strategist, innovator, and thought leader focused on applying technology to solve problems for clients in the public, private, and charity sectors. He is currently pioneering the use of AI in EY's tax and law practice and leads the firm's emerging technology thought leadership.

Dr. Lewis is an accomplished writer and public speaker, winning the 2017 Management Consultancies Association Award for the best use of thought leadership investigating the impact of AI on the UK's economy.

Dr. Lewis is a former member of the Open Data User Group, the Public Sector Transparency Board, and the Advisory Committee to the All-Party Parliamentary Group on AI. He is a member of techUK's leadership committee for Data Analytics and AI, on the editorial board for the 'Real-World Data Science' platform of the Royal Statistical Society, and is also an Honorary Senior Visiting Fellow at The Bayes Business School in London.

He lives in the South Downs National Park with his wife and sons.

Endnotes

¹ *ViewPoints* reflects the network's use of a modified version of the Chatham House Rule whereby names of members and their company affiliations are a matter of public record, but comments are not attributed to individuals or corporations. Italicized quotations reflect comments made in connection with the meeting by network members and other meeting participants.

² Krystal Hu, "[ChatGPT Sets Record for Fastest-Growing User Base](#)," *Reuters*, February 2, 2023.

³ Michael Chui et al., *The Economic Potential of Generative AI: The Next Productivity Frontier* (McKinsey & Company, June 2023), 3.

⁴ "[OECD AI Principles Overview](#)," Organisation for Economic Co-operation and Development, accessed August 15, 2023.

⁵ Paul Glynn, "[Sony World Photography Award 2023: Winner Refuses Award after Revealing AI Creation](#)," *BBC*, April 18, 2023.

⁶ Kristen E. Busch, *Generative Artificial Intelligence and Data Privacy: A Primer* (Congressional Research Service, May 23, 2023).