



Artificial Intelligence: An Emerging Oversight Responsibility for Audit Committees?

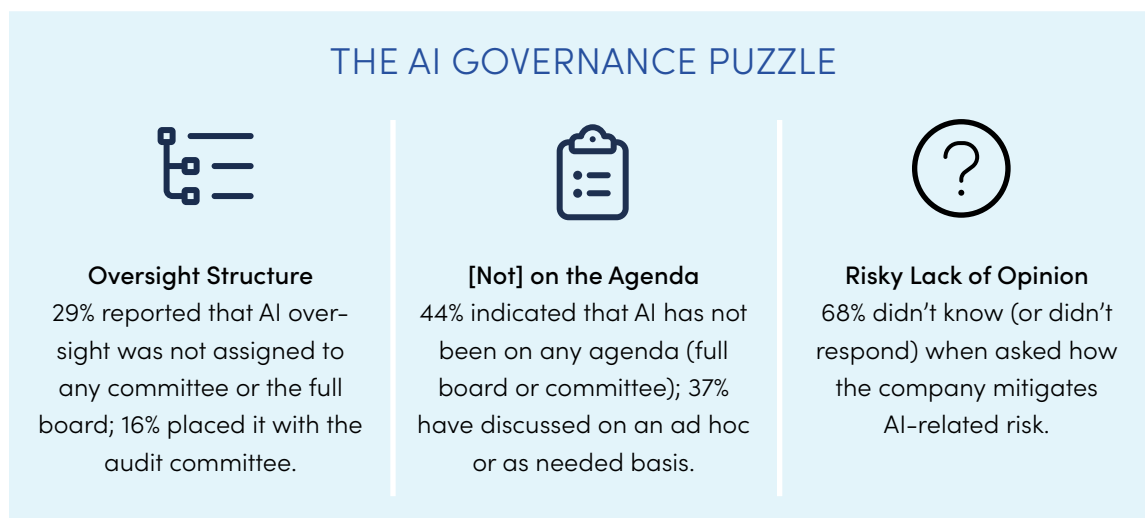
By Brian Cassidy, Ryan Hittner, and Krista Parsons, Deloitte & Touche LLP

The audit committee has many discrete duties, including overseeing financial reporting and related internal controls, the independent and internal auditors, and ethics and compliance, to name just a few. However, these and other duties are part of a broader audit committee responsibility: risk oversight. While the audit committee does not manage all risks, it is responsible for overseeing the procedures and processes by which the company anticipates, evaluates, monitors, and manages risks of all types. Recent developments in artificial intelligence (AI), including the emergence of generative AI, are leading businesses to evaluate AI's potential impact to their business technology strategy. As businesses expand their use of AI, especially into core business processes, the audit committee will need to understand the challenges and opportunities presented by AI to address risks related to governance and stakeholder trust.

WHO'S MINDING THE AI STORE NOW?

According to a [2023 survey](#) conducted by Deloitte and the Society for Corporate Governance, corporate secretaries see AI strategy and oversight as still evolving. The findings show that few respondents (13%) had a formalized AI oversight framework, although many (36%) were considering the development and implementation of AI oversight policies and procedures.

These results are particularly interesting when compared to a 2022 Deloitte [survey](#), in which 94 percent of respondents said AI was critical to their company's short-term success.¹ This may suggest some level of information asymmetry between management and the board, congruent with the notion that AI is in a state of flux. Thus, at least for now, the AI landscape might best be characterized as an abstract governance puzzle.²



RISKS AND OPPORTUNITIES

FAMILIAR AND DIFFERENT SET OF RISKS

With new technology comes the possibility of new risks. Some AI risks present well-trodden challenges that arise in other technology areas and can be overseen and understood in the context of an ongoing enterprise risk management (ERM) process,³ such as the [COSO ERM framework](#). However, other risks may be unfamiliar and/or amplified. A few illustrative examples are highlighted below.

- ▶ **Shadow IT Environments:** Use of IT assets by personnel without the knowledge or oversight of IT security professionals can occur with any type of software or hardware. However, unauthorized use of generative AI by personnel may compound data-related risks. This risk may be increased given the [lack of AI policy](#) in many organizations. Further, employees leveraging generative AI to write code may [inadvertently introduce vulnerabilities](#) through code generated by AI.

¹ Business leaders were defined as company representatives who met one or more of the following qualifiers: (1) responsible for AI technology spending or approval of AI investments, (2) responsible for the development of AI strategy, (3) responsible for implementation of AI technology, (4) acting as AI technology subject-matter specialist, or (5) otherwise stated they were influencing decisions around AI technology. See Nitin Mittal, Irfan Saif, and Beena Ammanath, *Fueling the AI transformation: Four key actions powering widespread value from AI, right now, State of AI in the Enterprise, 5th Edition* report, Deloitte, October 2022.





² Natalie Cooper, Bob Lamm, and Randi Val Morrison, "Future of tech: Artificial intelligence (AI)," *Board Practices Quarterly*, Deloitte, August 2023.

³ Alexander J. Wulf and Ognyan Seizov, "'Please understand we cannot provide further information': Evaluating content and transparency of GDPR-mandated AI disclosures," *AI & Society* (2022).

- ▶ **IP Ownership and Infringement:** Generative AI users can input confidential or protected data, which may result in an array of adverse outcomes, including disclosure of such confidential or protected data to third parties. Outputs using this type of data may also constitute infringement of intellectual property.⁴ Furthermore, as generative AI applications are used to craft increasingly sophisticated media across multiple formats, it may not be clear who owns the rights to any resulting intellectual property.
- ▶ **Cybersecurity Bad Actors:** A frequent concern across many types of technology stems from malicious actors who circumvent security protocols. Generative AI use cases may amplify some types of cybersecurity risks. For example, hackers may use generative AI to write code for purposes of infiltrating data environments or create phishing messages that more accurately mimic human language and tone.

Finding the appropriate balance between AI's benefits and risks depends on a constellation of factors. Outputs produced by generative AI change over time as the technology learns from data. But just like with humans, it is possible for this subcategory of AI technology to learn things that are incorrect. For that reason, traditional risk management strategies **may not be well-equipped** for the challenges that arise from generative AI use.

GENERATIVE AI RISK EXAMPLES

			
<p>Low Transparency How generative AI derives its output can be a “black box,” making it difficult to explain and/or audit.</p>	<p>Hallucination Generative AI products and services may generate output that seems accurate but is actually false or cannot be justified.</p>	<p>Bias Potential When trained on nonrepresentative data, generative AI output could exhibit systematic errors.</p>	<p>Value Alignment Even with safeguards, generative AI output may contradict its intended purpose.⁵</p>

⁴ Christian Heinze, “Patent infringement by development and use of artificial intelligence systems, specifically artificial neural networks,” in *A Critical Mind: Hanns Ullrich’s Footprint in Internal Market Law, Antitrust and Intellectual Property*, eds. Christine Godt and Matthias Lamping, MPI Studies on Intellectual Property and Competition Law, vol. 30 (Heidelberg, Germany: Springer, 2023), pp. 489–515.

⁵ Vic Katyal, Cory Liepold, and Satish Iyengar, “Artificial intelligence and ethics: An emerging area of board oversight responsibility,” *On the Board’s Agenda*, Deloitte, 2020.

Regardless of whether the risk is familiar, completely new, and/or amplified, the resultant consequences may be notable. Failure to mitigate any subcategory of AI-related risks may lead to many adverse outcomes such as reputational damage, financial losses, legal action, and regulatory infractions. A starting point for addressing such concerns might include using mitigation strategies that are already known to work in other contexts, such as the [COSO ERM framework](#) referred to earlier. For AI-centric guidance related to implementation and scaling, it may be worth considering the benefit of systems such as the [NIST AI Risk Management Framework](#).

WITH RISKS COME BENEFITS, TOO

If AI presented nothing but risk, it seems unlikely that it would have emerged as “the” technology of the future. Clearly, AI has benefits, some of which may not be known for some time. One particular set of benefits is squarely in the audit committee’s wheelhouse—namely, the potential to streamline and enhance a company’s internal audit, financial reporting, and internal control functions. There are also aspects of generative AI technology that, while still evolving, may one day fundamentally change an organization’s financial systems. While there is much uncertainty, the future transformative potential of generative AI may add much to the current array of use cases. In the shorter term, various subcategories of AI are already capable of improving the quality of financial reporting via reviewing transactions, identifying errors, addressing internal control gaps, and detecting fraud. If AI isn’t being used within these areas, the audit committee might ask if the company is exploring potential use cases—and if the company is not, the committee might ask to hear the reasons behind that decision.

USE OF AI TECHNOLOGY MAY HAVE MANY BENEFITS



Cost Savings

Process automations and improvements may improve task efficiency.



Boosted Revenues

AI-infused products and services may provide new growth opportunities.



Development Time





AI may shorten time to market by increasing the speed of early-stage testing.



New Insights

Appropriate generative AI use may bolster employee creativity.

COMMON AI USE CASE EXAMPLES

USE CASE	DESCRIPTION	OPPORTUNITIES	RISKS
 <p>Invoices and Payments</p>	Use of intelligent automation to match invoices to payments, including classification of expenses	The technology may reduce costs by processing a large volume of transactions with a high degree of accuracy.	Poorly designed or maintained systems may generate errors that are time consuming to undo.
 <p>Contract Review or Generation</p>	Leverage of natural language and generative AI processing to create legal documents or review them for errors	By producing the initial drafts or identifying common errors, generative AI may create efficiencies and lower legal liability in a cost-effective manner.	Natural language and generative AI trained on biased data may misapply the law or make up precedent.
 <p>Forecasting and Modeling</p>	Incorporating predictive analytics to improve the accuracy of functions like inventory management and revenue forecasting	Modeling and analytics AI technology may be capable of identifying patterns at a speed that outpaces human-led data analysis efforts.	Lack of robust testing and regular updates can cause modeling and analytics AI to become more inaccurate over time.
 <p>Code Development</p>	Use of generative AI to develop models or applications that create efficiencies for routine personnel activities	Employees may use generative AI to drive efficiencies in day-to-day tasks and help identify possible generative AI use cases.	The technology may expose confidential data with generative AI inputs or may create outputs that involve intellectual property infringement.

AI AND THE AUDIT COMMITTEE

The tendency to assign oversight of emerging risks to the audit committee means it is sometimes described as the “kitchen sink” of the board. However, as noted earlier, this is consistent with the audit committee’s overarching role in risk oversight. It’s also worth considering that it is common for topics taken on by the audit committee at the outset to eventually be overseen by other committees. Some aspects of AI oversight seem more aligned with the audit committee’s work than others. And when it comes to considering such congruence questions, it may be helpful to think about the audit committee’s current levels of technology fluency and comfort. For instance, given the audit committee’s traditional governance areas, it may be prudent for it to oversee AI use in financial reporting.⁶

⁶ The audit committee may want to also think about indirect impacts. Depending on the use case, AI technology may have an array of indirect effects on financial measures (GAAP or otherwise).

In other parts of AI oversight, it may be less clear whether the audit committee is a “good fit.” For example, the impact of generative or natural language AI on the workforce may be more aligned with the oversight of the compensation/talent committee or the full board.

The “temporary assignment” of AI to the audit committee may make sense for other reasons, as well. First, AI remains an emerging technology and is likely to continue to change rapidly. Second, there is extensive governmental interest in AI, which may result in legislation that will require adjustments in its oversight. Thus, determining now that AI, or aspects of AI, should be overseen by another committee or committees may turn out to be premature.

An audit committee might choose to assess its AI risk tolerance across oversight areas such as auditing, financial reporting, and internal control functions. It may be helpful to contextualize that analysis by comparing it to other areas of the company. For example, company divisions that routinely use technology enhancements in client-facing operations may have a higher appetite for risk. But a higher risk tolerance in operational settings does not necessarily correlate with how risks are viewed when it comes to financial reporting impacts.

An important part of the AI governance puzzle for the audit committee is assessing risk. But, at least for now, this task is currently made more difficult by a shifting regulatory landscape. Governments and regulators around the world are considering whether regulation and policy can address AI risks. Their progress toward developing and enacting policies and regulations over AI is uneven across the globe and in different stages of development and enactment. And to make things more complex, stakeholder groups—shareholders, customers/clients, employees, suppliers, and community—all have varying and sometimes conflicting expectations around use and governance of AI. For these reasons, there may be a benefit to continuously assessing AI risks and benefits over waiting for emerging and future legislative proposals or regulatory guidance. But to accurately make such continual assessments, it’s important that the audit committee and the board have sufficient knowledge to ask questions around the organization’s adoption and use of AI.



POTENTIAL AUDIT COMMITTEE OVERSIGHT QUESTIONS TO CONSIDER

- ▶ What are the company’s current and potential future use cases for AI, and do any of them have an impact on financial reporting or other audit committee oversight areas?
- ▶ Has management considered opportunities to use AI that may enhance or improve financial reporting processes?
- ▶ What processes are, or will be, used to evaluate dependencies that may arise in other areas where the audit committee may have primary oversight, like cybersecurity or data management?
- ▶ Are processes for use of AI congruent with the company’s risk appetite in terms of level of proactiveness and mitigation strategy?
- ▶ Given the speed of AI technology development, are existing processes being assessed and updated with appropriate frequency?



Brian Cassidy



Ryan Hittner



Krista Parsons

Brian Cassidy is an Audit & Assurance partner with Deloitte & Touche LLP and the US Audit & Assurance Trustworthy AI leader.

Ryan Hittner is an Audit & Assurance principal with Deloitte & Touche LLP and the US Artificial Intelligence & Algorithmic Assurance coleader.

Krista Parsons is an Audit & Assurance managing director with Deloitte & Touche LLP. She is also the Governance Services coleader and the Audit Committee Program leader for Deloitte's Center for Board Effectiveness.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional adviser. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see <http://www.deloitte.com/about> to learn more.