AI: Looking for ROI, moving toward agency

Boardroom Perspective

KPMG Board Leadership Center

Artificial intelligence (AI) has continued to capture the attention and imagination of companies, policymakers, and society at a remarkable pace. At the 2025 AI Summit in Paris, technology leaders proclaimed that "AI will be the most profound shift of our lifetimes" and that AI will lead to the "largest change to the global labor market in human history." The opportunities and implications, as speculated by Anthropic's CEO, are profound: "Cancer is cured, the economy grows at 10 percent a year, the [federal] budget is balanced—and 20 percent of people don't have jobs."[1]

For many companies and their boards, the opportunities and challenges are stacking up, and the technology's advance is waiting for no one.

As AI reshapes the landscape and becomes ever more critical to a company's strategy and competitive position, a top priority for boards today is to understand where this is all headed—the state of AI, the technology's pace and direction, and where management and the board should be focusing their attention to capitalize on AI's value-creation opportunities.

Our recent surveys[2] and conversations[3] with directors, AI experts, and business leaders offer timely insights and observations on the AI landscape, where many companies are on their AI journey, and where they are likely to go in the near and long term as the technology continues to evolve.

[1] Jim VandeHei and Mike Allen, "Behind the curtain," Axios, May 28, 2025.

[2] See 2025 Survey: Boardroom Lens on Generative AI, KPMG Board Leadership Center, March 2025, and KPMG Quarterly AI Pulse Survey, April 2025.

[3] Observations from our conversations with directors and business leaders were shared, and are included here, under the Chatham House Rule.

Leaning into GenAI as a force amplifier

The boardroom focus on generative AI (GenAI) has evolved. "In the initial years, there was a lot of excitement around GenAI, but it was focused on the 'wow factor,'" according to one expert. "Now the focus is on ROI, getting productivity. We're at the turning point of the J-curve where companies are expecting to get some benefit." Earlier discussions about GenAI had safety as a major theme—and that's still important, but the key message is that companies need to "lean in and move faster to gain the billions in potential benefits of GenAI." The focus is on building business value—by improving the efficiency and productivity of a broader group of employees, by generating deeper insights, and by developing new products and services.

Directors and AI experts emphasized the "GenAI augmentation imperative." Critical to realizing the promised benefits of GenAI is the retooling of workers with broadly available GenAI assistants. Companies need to stimulate the human behavioral change required at scale to adopt existing GenAI tools (large language models [LLMs], copilots, or embedded software) and encourage employee use of these tools in ways that free up significant time, which can be reinvested in something equally or more productive for their company. More than three-quarters of GenAI survey respondents said that the top benefit of the company's adoption and leveraging GenAI is the optimization of the company's operations—efficiency, productivity, and cost savings.

Another takeaway is that AI has not been a "leveler of the playing field," as initially predicted. Experts said that leading firms are pulling away in almost every industry because they are using the technology differently. Companies and their boards should expand the focus on GenAI beyond how the technology can help cut costs to consider how it may be used as a force amplifier—for example, to help expand the company market or deliver its services in a different way.

For many companies, data is a major bottleneck to capturing the benefits of GenAI. GenAI needs access to company data, which may need to be "piped out of the company's basement" and into LLMs. "Where does the company maintain its proprietary data? How much sits in a cloud-based platform? Ask questions about the privacy, governance, and quality of the data."

While only a modest number of directors reported that their companies are deriving revenue from GenAI,[4] in some industries, GenAI has already become key to business strategy for industry leaders, driving innovation and efficiency in areas like product design, supply chain optimization, predictive analytics, and personalized customer experiences. And companies are now deploying tailored, industry-specific applications of GenAI. In our April 2025 Quarterly AI Pulse Survey, 82 percent of respondents expect risk management to be the greatest challenge to their GenAI strategies for the remainder of 2025, followed by data quality (64 percent) and personal trust in GenAI (35 percent).

Key survey findings: 2025 Boardroom Lens on AI

- **Top benefits of the company's adoption and leveraging GenAI:** Optimization of the company's operations (75 percent); development of new products, services, or line(s) of business (12 percent); and increased revenue from existing products and services (9 percent)
- **State of GenAI adoption**: Integrated into their company's strategy (16 percent), company is piloting the technology selectively (41 percent), company starting to scale GenAI (25 percent)
- **Most significant obstacles facing the company in its deployment of GenAI:** New skills and talent (47 percent); need for workforce transformation, including retraining and/or workforce reductions (42 percent)

[4] 2025 Survey: Boardroom Lens on Generative AI, KPMG BLC, March 2025.

AI agents and agentic AI: "Doing, not just knowing"

Many companies are adopting, or exploring the adoption of, emerging forms of AI—such as the deployment of AI agents and ultimately agentic AI—to automate and re-engineer workflows. In our Board Leadership Center (BLC) director survey, 23 percent of respondents said the adoption of emerging forms of AI, such as AI agents and agentic AI, is a strategic priority for the company for 2025 (14 percent said it was not a strategic priority for 2025, but should be). Another 30 percent said the company was actively exploring emerging forms of AI for adoption over the next two to three years.

Intelligence versus agency

Agents are software entities that fulfill goals by planning and taking actions. Using LLMs to reason, they can break down and understand long-form instructions and use tools to accomplish tasks. With GenAI, it's up to you to do something. Agents act independently in pursuit of an expressed goal, whether finding information, working with data or running operations, while keeping humans in the loop. That's the difference—agents mean doing, not just knowing.

In the software development lifecycle, AI agents can generate documentation or perform code reviews based on established coding guidelines, quality standards, and security practices. In financial services, agents can engage in financial planning and analysis for budget forecasting and variance analysis. In healthcare, agents can examine electronic health records and other information to give clinicians a concise view of a patient, flag potential drug interactions, and synthesize longitudinal records—allowing humans to spend more time on meaningful care.

Excerpted from "Readying the Enterprise for Agentic AI," Q&A with Swami Chandrasekaran, Global Head of AI and Data Labs, KPMG LLP, Wall Street Journal, May 2025.

In the KPMG Quarterly AI Pulse Survey, 90 percent of executives said their organizations are past AI-agent experimentation, of which 33 percent have deployed at least some agents. Just over half of organizations are piloting agents, and some 10 percent are exploring the possibility of using agents. This acceleration, however, reveals continued challenges. The primary obstacles to agent deployment include technical skills gaps (59 percent), workforce resistance to change (47 percent), and system complexity (39 percent). Nearly half of leaders (46 percent) are equally focused on efficiency and revenue growth as it relates to their AI agent strategies.[5]

"AI agentic systems are evolving rapidly, moving from simple 'taskers' designed to automate single functions to 'orchestrators' where multiple AI agents interact to achieve complex tasks at scale. They can operate across different domains—for example, in healthcare they reduce administrative overhead, while in financial services they help optimize complex trading strategies—demonstrating true cross-industry applicability."[6]

For most companies, the deployment of AI agents will be an evolutionary process. The KPMG AI framework[7] classifies AI agents into four key types—taskers, automators, collaborators, and orchestrators—each one leveraging the same foundational tools and capabilities but differing in goal planning, execution, and complexity. This framework provides a structured approach to scaling AI from singular goal fulfillment to fully orchestrated, adaptive multiagent systems that smoothly integrate into business operations—sometimes referred to as agentic AI.

Today, many AI agents still appear to be in the lift-and-shift phase, where the only question being asked is, "How can I take advantage of this technology to improve what I do but without fundamentally changing what I do?" Tomorrow, the question will likely be, "How can I take advantage of this technology to reengineer processes, workflows, and business models?"

While AI agents and agentic AI are tailored to specific workflows and operate largely autonomously, there is a drumbeat about the importance of a human "on" the loop supervising the process to help ensure that AI agents and agentic AI are functioning as intended.

[5] KPMG AI Quarterly Pulse Survey: Q2 2025, June 26, 2025.

[6] "Creating Value with AI Agents," Technically Speaking: A KPMG Blog Series, 2025.

[7] Swami Chandrasekaran, Global Head of AI and Data Labs, KPMG LLP, "The 'TACO' Framework For Understanding AI Agents," Artificial Lawyer, February 12, 2025.

Companies should take a risk-based approach when considering the level of human intervention, as there will be lower-skill agents or low-risk areas with acceptable error rates that may be fully autonomous. For higher-skill agents and higher-risk areas, the human on the loop must have sufficient knowledge/expertise to review the output. This is a particular concern today given fears about de-skilling—i.e., concerns that early-career professionals may not have the critical thinking skills and judgment typically acquired in the very early years of their development.

Directors and AI experts emphasized other developments or advances in the use of AI, including:

- *LLMs and small language models (SLMs).* We typically think about GenAI "frontier" models, but LLMs are not needed for all tasks, and in some cases may be less efficient. Many companies have moved away from LLMs to SLMs, which can be fine-tuned for specific tasks and targeted use cases. SLMs are smaller in scale and more efficient, requiring less memory and computational power
- *Company-specific tailored models, or custom AI models.* Many companies are adopting custom AI models, which are designed and trained to address company-specific needs and data. These models can be based on either LLMs or SLMs, but they are further customized through fine-tuning or retraining on proprietary, high-quality data sets unique to the company.
- *Interplay of GenAI and other technologies such as robotics, blockchain, and quantum computing.* One AI expert emphasized the interplay of GenAI and robotics. "We've progressed leaps and bounds in the use of AI for cognitive tasks. Physical intelligence is the next frontier after agentic AI." The same tools that helped with cognitive tasks are now turbo-charging the ability of robots to do many tasks, eliminating the need to build special equipment for every task.

Unique risks posed by AI agents and agentic AI

In *Board oversight of GenAI: Business value, guardrails, and governance*, we discussed the significant risks posed by GenAI—inaccurate data and results, cybersecurity risks, data privacy and security risks, compliance risks, intellectual property risks, and reputational risks—as well as the guardrails and governance policies for the development, deployment, and use of GenAI. We also emphasized that the deployment of GenAI should prompt companies to take a hard look at the quality of the company's data and data governance practices, as achieving the hoped-for productivity and efficiency improvements with GenAI depends on the quality of the company's data, and how it is processed and stored.

Implementing GenAI helps build the necessary infrastructure, AI and data governance, risk management, and organizational experience needed before adopting more complex AI agent systems. As companies consider the deployment of AI agents, it is essential that they reassess and, as appropriate, modify their AI and data-related governance policies and guardrails to address the evolving and unique risks posed by AI agents—primarily risks posed by AI agents' autonomous decision-making and operational integration. AI operates with increasing autonomy, making decisions and executing actions without continuous human oversight.

Reassessing AI and data-related governance policies and guardrails to address the risks posed by AI agents is a time-consuming and resource-intensive undertaking, requiring a comprehensive review of the company's governance, security, transparency, and ethical controls, including security and access controls, continuous monitoring and anomaly detection, adequacy of human-in-the-loop, comprehensive audit trails, and ethical and cultural considerations.

Considerations for boards

As companies increase their deployment of AI and agents, we emphasize four critical areas of board focus:

Full implementation of GenAI

In addition to being vital for competitive advantage—particularly given the efficiency and productivity benefits of the technology—the successful deployment of GenAI with strong governance and risk management is a prerequisite enabler—and the foundation—for the more complex deployments of autonomous AI agents and ultimately agentic AI. Addressing GenAI foundational adoption issues, particularly risk management and trust, is essential in preparing for the deployment of AI agents and agentic AI. As we noted in On the 2025 Board Agenda, the board should understand the company's strategy to develop business value with GenAI and monitor the trajectory of deployment; monitor management's governance structure for the deployment and use of GenAI, including the management and mitigation of GenAI risks; and understand how the company is ensuring the quality and accuracy of GenAI output.

Deployment of AI agents

In addition to the projected productivity benefits, deploying more autonomous AI agents is a key priority because it is an important step—the building block—for deploying more complex agentic AI. The staged deployment of AI agents enables companies to gain operational familiarity, and establish the necessary data and governance frameworks, so the company is better positioned to handle more complex, adaptive, and autonomous AI capabilities, and the transition to agentic AI becomes feasible.

In preparing for the deployment of AI agents, the following three-phase framework from the KPMG blog, Creating Value with AI Agents, may be helpful for management teams and boards to consider:

Measuring AI readiness. Conduct both an outside-in (i.e., market research) and inside-out (i.e., data collection) analysis of AI agent strategies, data and technologies, trust and governance, workforce skills, project management, and value tracking. Identify gaps and develop a plan to address them. For instance, form a cross-functional committee—with members from information technology, legal, compliance, and human resources—to evaluate the company's data posture, AI maturity, and organizational readiness for agentic AI solutions.

Conducting a thorough opportunity assessment. Outline use cases for AI agents, and prioritize those use cases based on risk, value, and complexity. The goal is a prioritized roadmap for use-case implementation, including return on investment (ROI) metrics. Consider starting with "low-hanging fruit," such as chat-based employee support or automated financial reconciliation, then expanding to more disruptive, cross-functional initiatives once initial successes are established.

Implementing an AI agent operating model designed for value, scalability, and sustainability. This foundation is needed to integrate AI agents, foster innovation, and develop a competitive edge.

ROI and the bottom line

Wherever the company is on its AI journey—piloting GenAI, deploying AI agents, or exploring the deployment of agentic AI more fully—ROI should remain front and center. To maximize ROI and increase productivity, many companies are taking a "first principles" approach—recognizing that GenAI and AI agents also provide an opportunity to rethink processes from the ground up, rather than making existing processes slightly more efficient and/or cost effective.

In measuring ROI in the development of deployment of AI agents and agentic AI, Chandrasekaran says companies should focus on broader dimensions beyond obvious productivity metrics like hours saved or cost reduction: "What about revenue growth? Are you seeing more dollars per customer? What about quality improvements? Are there fewer errors and less rework? Look at cycle times in your operations. Are compliance violations down? How about employee and customer satisfaction? The key is to deliberately track these metrics from day one and systematically build metric capture into your systems so you can measure the impact. Otherwise, you won't have the confidence needed to successfully enable the power of agentic AI."[8]

The people aspect

Achieving the transformational benefits of the adoption of GenAI and AI agents at scale hinges on placing people at the core of adoption—prioritizing change management, workforce empowerment, skills development, and cultural transformation. What skills are necessary in an agentic AI and GenAI world today? How do we mitigate the risk of "de-skilling"? It isn't enough for GenAI and agentic AI to be deployed; the technologies must be embraced. Fundamentally changing what people do every day and how they work will require leadership.

Directors also continue to emphasize the importance of director education on AI in all of its forms. As one director noted, "I think trying to keep ahead of [AI developments] is nearly impossible, but we decided we're going to not only have reports from management and use some outside services, but we wanted a couple of our board members that are better in tech to dive a little deeper into this. We asked them to spend some time on education, getting some training so they can bring that back to the board."

[8] Swami Chandrasekaran, Global Head of AI and Data Labs, KPMG LLP, "Readying the Enterprise for Agentic AI," Wall Street Journal Q&A, May 2025.

Contact

kpmg.com/us/blc
T: 800-808-5764
E: us-kpmgmktblc@kpmg.com

John H. Rodi
Co-Leader
KPMG BLC

Patrick Lee
Senior Advisor
KPMG BLC

Anne C. Zavarella
Co-Leader
KPMG BLC

About the KPMG Board Leadership Center

The KPMG BLC champions outstanding corporate governance to drive long term value and enhance stakeholder confidence. Through an array of insights, perspectives, and programs, the BLC promotes continuous education and improvement of public and private company governance. BLC engages with directors and business leaders on the critical issues driving board agendas—from strategy, risk, talent, and sustainability to data governance, artificial intelligence, audit quality, proxy trends, and more. Learn more at **kpmg.com/us/blc**.

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: **kpmg.com**