How can responsible AI bridge the gap between investment and impact?

8 October 2025

EY 安永

Shape the future
with confidence
聚信心 塑未来

How can responsible AI bridge the gap
between investment and impact?

8 October 2025

AI pays off when it's embedded responsibly: greater profits, happier employees, and fewer costly mistakes.

In brief

- Companies with oversight and real-time monitoring are turning AI from a risk into a growth engine.

- Significant financial loss from AI risk is real and here to stay for those that don't have the right controls in place.

- Leadership blind spots can leave firms exposed; it is essential to have oversight of citizen developers.

The companies outperforming with artificial intelligence (AI) aren't just building better models – they're building smarter guardrails that let them seize outsized market opportunities. The latest EY Global Responsible AI Pulse survey reveals that organizations embracing responsible AI – through clear principles, robust execution, and strong governance – are pulling ahead in the metrics where AI-related gains have been most elusive: revenue growth, cost savings, and employee satisfaction. These gains aren't marginal, they're the difference between AI as a cost center and AI as a competitive advantage.

Nearly every company in the survey has already suffered financial losses from AI-related incidents, with average damages conservatively topping US$4.4 million. But those with governance measures like real-time monitoring and oversight committees are seeing far fewer risks – and stronger returns.

Responsible AI isn't a compliance exercise. It's a performance lever – and the latest data proves it.

About the research

The EY organization has launched a Responsible AI Pulse survey to provide a regular snapshot of business leaders' real-world views on responsible AI adoption. This article is the second in a series and builds off the findings of the June 2025 pulse survey.

In July and August 2025, the global EY organization conducted research to better understand C-suite views around responsible AI – for the current and next wave of AI technologies. To underpin this research, we conducted an anonymous online survey of 975 C-suite leaders across 10 roles. All respondents had some level of responsibility for AI within their organization. Respondents represented organizations with over US$1 billion in annual revenue across all major sectors and 21 countries in the Americas, Asia-Pacific, Europe, the Middle East, India and Africa.

The responsible AI journey: companies are taking a comprehensive approach

Responsible AI is best understood as a journey – one that moves through three stages.

- First comes communication, where organizations articulate a clear set of responsible AI principles internally and externally.

- Next is execution when those principles are translated from words to action, through controls, key performance indicators and workforce training.

- Finally comes governance, the oversight needed to help ensure actions and principles stay aligned, through measures such as committees and independent audits.

Most companies have started this journey. The second wave of the EY Global Responsible AI Pulse survey asked C-suite leaders about responsible AI adoption steps across these three stages. On average companies have implemented seven out of 10 measures.

Adoption is even higher in sectors like technology, media and entertainment and telecommunications (TMT), where a higher dependency on technology and data for delivering core services makes responsible AI all the more critical. Organizations in these sectors are more likely than others to communicate responsible AI principles to external stakeholders (80% vs. 71%). They are also further ahead on governance: 74% have established an internal or external committee to oversee adherence to these principles (vs. 61% in other industries), and 72% conduct independent assessments of responsible AI governance and control practices (vs. 61%).
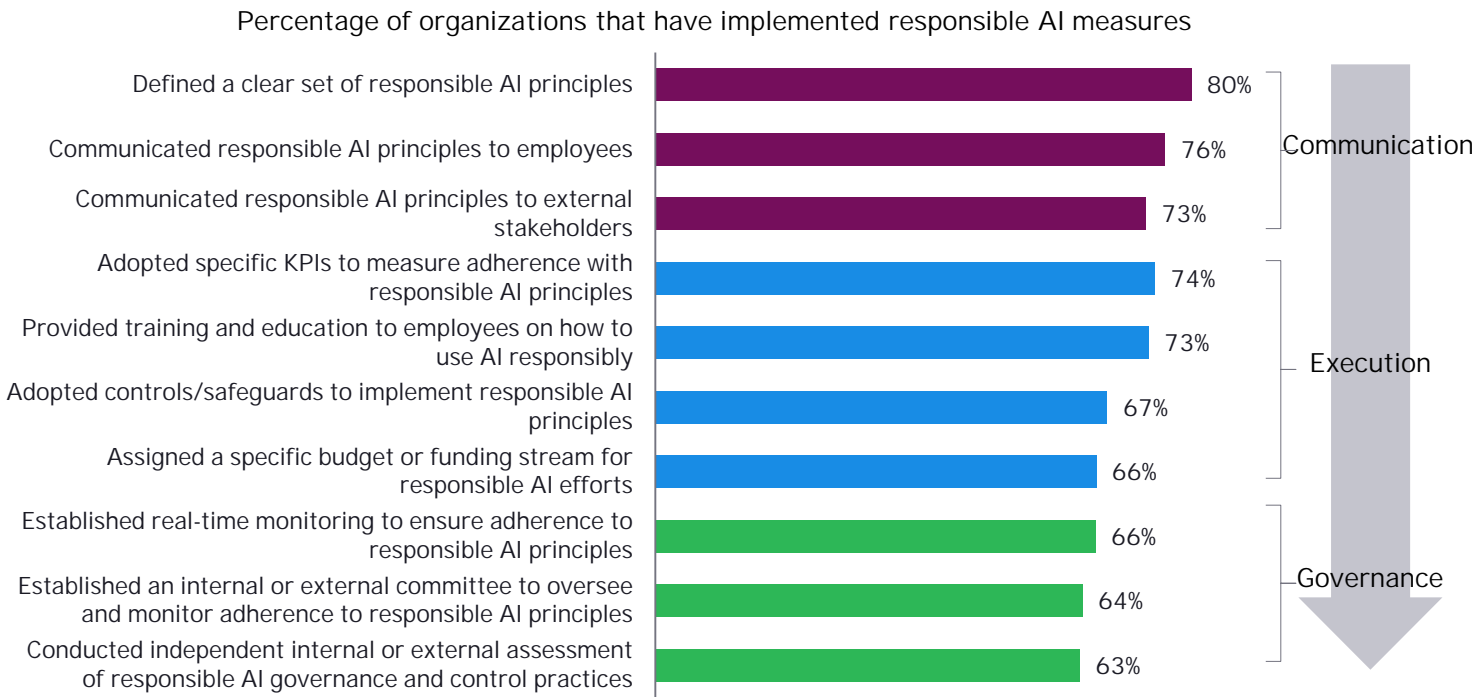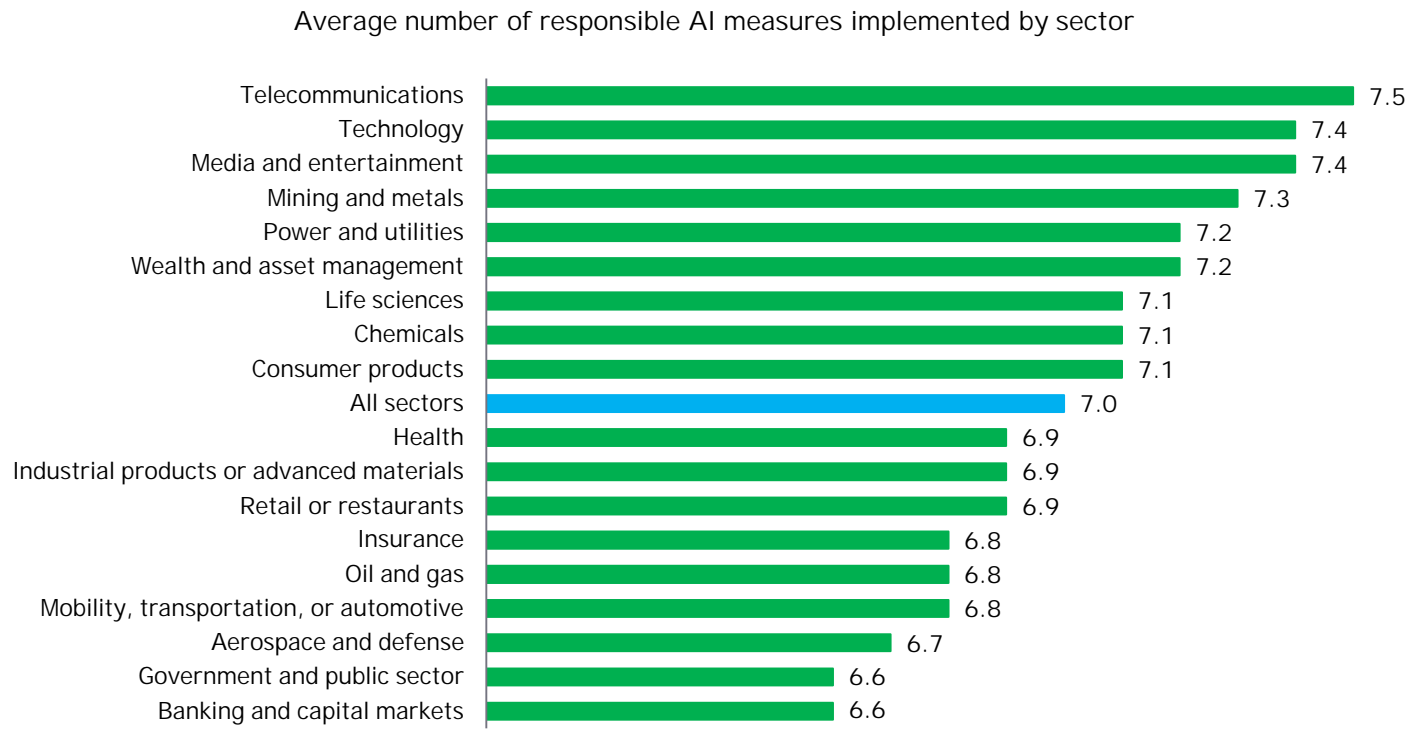
Percentage of organizations that have implemented responsible AI measures

Measure	Percentage	Stage
Defined a clear set of responsible AI principles	80%	Communication
Communicated responsible AI principles to employees	76%	Communication
Communicated responsible AI principles to external stakeholders	73%	Communication
Adopted specific KPIs to measure adherence with responsible AI principles	74%	Execution
Provided training and education to employees on how to use AI responsibly	73%	Execution
Adopted controls/safeguards to implement responsible AI principles	67%	Execution
Assigned a specific budget or funding stream for responsible AI efforts	66%	Execution
Established real-time monitoring to ensure adherence to responsible AI principles	66%	Governance
Established an internal or external committee to oversee and monitor adherence to responsible AI principles	64%	Governance
Conducted independent internal or external assessment of responsible AI governance and control practices	63%	Governance

Fig 1. The responsible AI journey

Average number of responsible AI measures implemented by sector

Sector	Average
Telecommunications	7.5
Technology	7.4
Media and entertainment	7.4
Mining and metals	7.3
Power and utilities	7.2
Wealth and asset management	7.2
Life sciences	7.1
Chemicals	7.1
Consumer products	7.1
All sectors	7.0
Health	6.9
Industrial products or advanced materials	6.9
Retail or restaurants	6.9
Insurance	6.8
Oil and gas	6.8
Mobility, transportation, or automotive	6.8
Aerospace and defense	6.7
Government and public sector	6.6
Banking and capital markets	6.6

Note: Averages are out of 10 possible measures

Fig 2. The responsible AI journey by sector

While there is a drop-off along each stage of the responsible AI journey, the gap is minimal, declining only a couple of percentage points on average from one step to the next. And where measures haven't yet been implemented, companies overwhelmingly say they intend to act. Across all responsible AI measures, fewer than 2% report their organization has no plans to implement.

This progress matters. Responsible AI can't be achieved though principles alone – it requires an "all of the above" approach. Clear articulation of principles, robust controls and strong governance are all essential to ensuring responsible AI moves from words to reality.

Responsible AI is the missing link

AI has already delivered big wins for many organizations. Eight in 10 respondents report improvements in efficiency and productivity – the primary focus of many early use cases. Nearly as many say AI has boosted innovation and technology adoption – helping to accelerate activities in which generative AI excels, such as ideation, discovery, research and development and rapid prototyping. About three in four say it has improved their ability to understand customers and respond quickly to shifting market conditions.

The extent AI adoption has improved company performance

	To a great extent	Somewhat	A little	Not at all
Innovation	37%	44%	16%	3%
Speed of scaling up from pilots or proofs of concept	34%	45%	19%	3%
Speed of tech adoption	34%	47%	16%	3%
Efficiency and productivity gains	32%	48%	19%	1%
Ability to respond to market changes rapidly	31%	46%	18%	4%
Ability to understand our customers	29%	47%	20%	4%
Employee satisfaction	24%	33%	33%	10%
Revenue growth	19%	35%	36%	10%
Cost savings	18%	30%	37%	14%

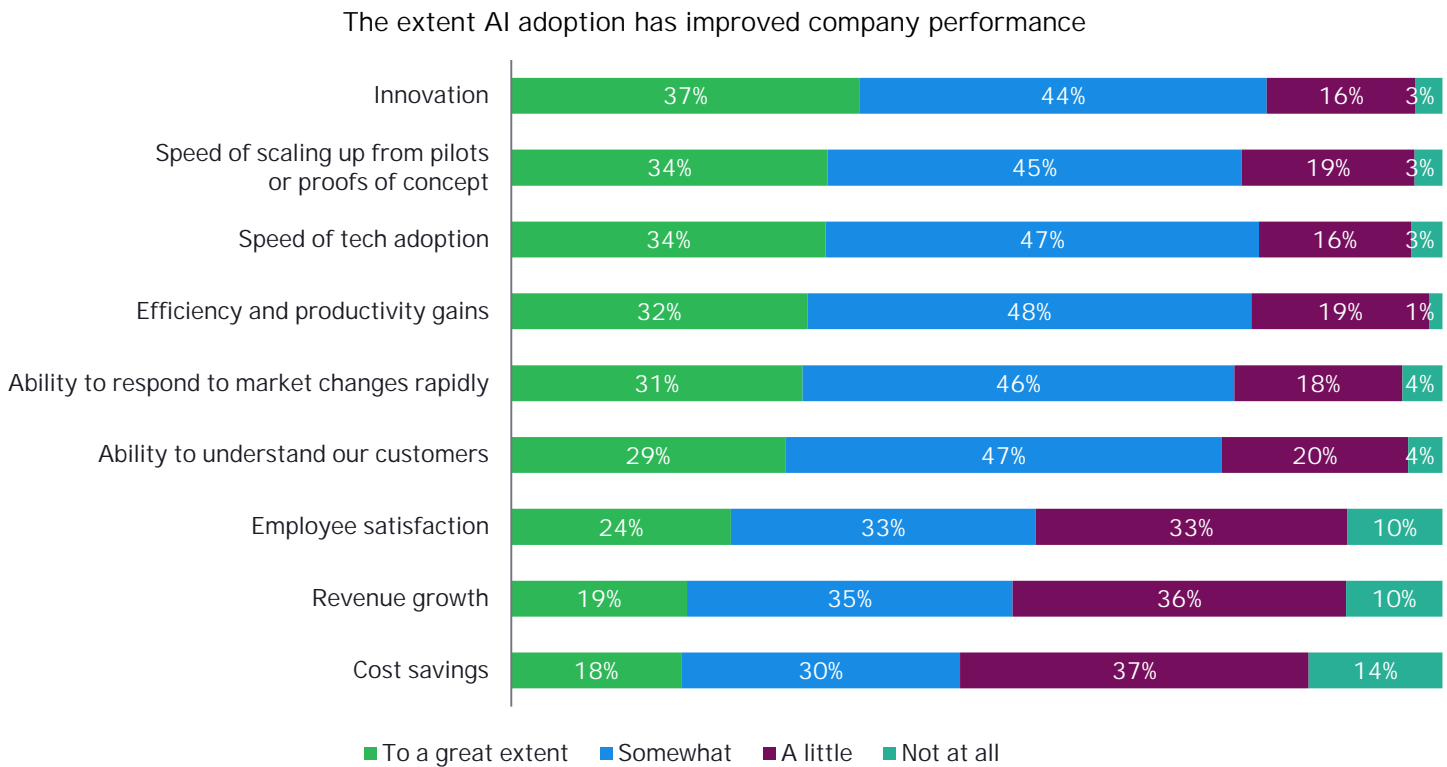■ To a great extent ■ Somewhat ■ A little ■ Not at all

Fig 3. Organizations adopting AI are capturing performance improvements, but key metrics lag

However, in three critical areas – employee satisfaction, revenue growth and cost savings – AI has not delivered similar performance improvement. According to the EY AI Sentiment survey, half of citizens are worried about job loss due to AI and many remain hesitant when it comes to AI's role in workplace decision-making. Translating AI investments into tangible improvements on the profit and loss statement also remains elusive for many firms.

Cathy Cobey, EY Global Responsible AI Leader for Assurance explains, "Organizations struggle to achieve positive ROI on their AI investments due to the complexities of integrating AI into existing processes, which demand re-engineering, upskilling, and continuous data flow investments. Additionally, challenges in legacy technology integration and the need for evolving governance frameworks hinder their ability to realize tangible financial benefits."

When we dug deeper into the data, however, something remarkable emerged: companies that have embraced responsible AI are breaking through where others are stalling. Organizations that are adopting governance measures – specifically real-time monitoring and oversight committees – are far more likely to report improvements in revenue growth, employee satisfaction and cost savings, the exact areas where most are struggling to report a return.

Percentage at which organizations are more likely to see the given improvement

Organizations with real-time monitoring are more likely to see the following improvements

Organizations with an oversight committee are more likely to see the following improvements

+35%
Revenue growth

+40%
Cost savings

+40%
Employee satisfaction

+34%
Revenue growth

+38%
Employee satisfaction

+65%
Cost savings
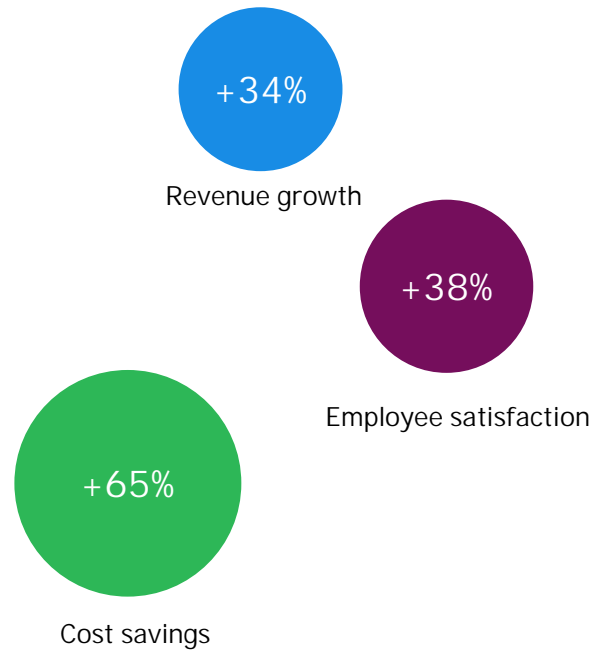
Note: Numbers represent the odds ratios, i.e., improvements in the likelihood of increasing company performance in cost savings, revenue growth, and employee satisfaction. Only RAI measures with statistically significant increases are shown.

Fig 4. Responsible AI adoption drives tangible improvements in AI performance

Identifying the drivers of performance improvements

To identify which responsible AI measures drive performance improvements, we used an ordered logistical regression model to determine which of the ten RAI measures tested increase the likelihood of performance improvements. We used odds ratios to quantify the impact of each measure on a set of performance improvement metrics and identified two responsible AI measures with statistically significant and positive impacts on revenue growth, cost savings and employee satisfaction.

This link suggests a symbiotic relationship. Companies that have moved further along the responsible AI journey are the ones seeing improvements in the areas that need the biggest boost, and it's not hard to see why. Anxious employees may be reassured by a public commitment to responsible AI from their employer. Communicating a responsible approach can build brand reputation and customer loyalty – ultimately driving revenue growth. And robust governance can help prevent costly technical and ethical breaches, as well as reducing recruitment and retention costs – benefits that ultimately flow through to the bottom line and boost cost savings.

For business leaders, the message is clear – increase the return on your AI investments by moving further along the responsible AI journey.

The price tag of ignoring the risks

While responsible AI adoption drives benefits, the converse is also true: neglecting it can come at a steep cost. Almost every company in our survey (99%) reported financial losses from AI-related risks, and 64% experienced losses exceeding US$1 million. On average, the financial loss to companies that have experienced risks is conservatively estimated at US$4.4 million. That's an estimated total loss of US$4.3 billion across the 975 respondents in our sample.

The most common risks organizations reported being negatively impacted by are non-compliance with AI regulations (57%), negative impacts to sustainability goals (55%) and bias in outputs (53%). Issues such as explainability, legal liability and reputational damage have so far been less prominent, but their significance is expected to grow as AI is deployed more visibly and at scale.

Encouragingly, responsible AI is already linked to fewer negative impacts. For example, those who have already defined a clear set of responsible AI principles have experienced 30% fewer risks compared to those who haven't.

C-suite blind spots leave companies exposed

Despite the financial stakes, it's clear that many C-suite leaders don't know how to apply the right controls to mitigate AI risks. When asked to match the appropriate controls against five AI-related risks, only 12% of respondents got them all right.

As may be expected, CIOs and CTOs performed the best – yet even here only about a quarter answered correctly across all five use cases.

Chief AI Officers (CAIOs) and Chief Digital Officers (CDOs) fared only slightly better than average (15%) likely reflecting a background more grounded in data science, academia and model development rather than traditional technology risk management. Consequently, they may have less experience managing technology-related risks than their CIO and CTO counterparts.

Concerningly, CROs – the leaders who are ultimately responsible for AI risks – perform slightly below average, at 11%. And at the bottom end of the spectrum CMOs, COOs and CEOs performed worst (3%, 6% and 6% respectively).

Percent who identified all of the appropriate controls against five risks

Role	Percent
Chief Technology Officer	26%
Chief Information Officer	24%
Chief AI or Data Officer	15%
Chief Strategy Officer	13%
All respondents	12%
Chief Risk Officer	11%
Chief Financial Officer	11%
Chief Human Resources Officer	9%
Chief Executive Officer	6%
Chief Operating Officer	6%
Chief Marketing Officer	3%

Note: The five risks posed to respondents were hallucinations, algorithmic blind spots or training limitations, bias in training data, legal ownership in autonomous systems, and transparency in AI-generated content.

Fig 5. Many C-suite leaders are not aware of appropriate controls for key AI risks

This lack of awareness has consequences. Firms that have lost more than US$10 million due to AI risks on average report that they have 4.5 out of 10 of the correct controls in place, while firms that have lost US$1 million or less have 6.4. This highlights a clear need for targeted upskilling of the C-suite – particularly as the financial and reputational cost of AI risks continue to rise.

Challenges ahead: Agentic AI and citizen developers

The governance challenge doesn't end with today's models. As agentic AI becomes more prevalent in the workplace and employees experiment with citizen development, the risks – and the need for thoughtful controls – are only set to grow.

"The autonomous nature of agentic AI introduces new risks that can escalate quickly."

Sinclair Schuller, EY Americas Responsible AI Leader

Encouragingly, most organizations are already implementing governance policies to manage these risks. Eight out of the 10 agentic AI governance measures we identified are being implemented by more than 75% of respondents. This includes continuous monitoring (85%), and incident escalation processes for unexpected agentic behaviors (80%). Although organizations have made a great start, there are still challenges in designing effective controls that can adequately oversee systems that operate continuously, adapt rapidly and require minimal human intervention.

Examples of internal governance concerns from survey respondents

- Cross-departmental collaboration: "I notice there is frequently a lack of cross departmental collaboration. Responsible AI is not just an issue for IT, the compliance departments also need to be involved in the process of responsible AI development, and this doesn't occur."

- Ethical guidelines: "Organizations have to set up and adopt strong ethical guidelines before using agentic AI, and not possessing the proper tools for this would be a hindrance."

- Budgeting & resource alignment: "I have observed that budgeting is usually not in alignment with the long term requirements of responsible AI governance, leading to inadequate resources to be in charge of agentic AI throughout its lifecycle."

- Risk management complexity: "The speed of changing complexity in these types of systems exceeds traditional risk management tools and organizations are left anticipating risks and a complete lack of knowing about how to assess or mitigate risks related to agentic AI."

- Lack of incident response planning: "I have noticed many organizations still do not have incident response plan for AI failure or misuse, in cases of inevitable failure they are left unprotected."

- Fragmented oversight: "There is often a disconnect between how one department or functional area works with AI and how the rest of the organization does, resulting in uncoordinated and chaotic evaluation protocols that lead to irresponsible use of agentic AI."

- Ownership: "I have observed a general absence of unequivocal ownership in organizations for continuous responsible AI governance, which leaves gaps in accountability."

Note: EY Insights analysis leveraging an LLM to systematically categorize and tag free-text survey responses according to a predefined curated taxonomy. In total, 36% were tagged to internal governance controls and the examples here represent quotes from survey respondents.

Fig 6. Internal governance issues are the leading concern in preparing for the responsible use of agentic AI

"In the age of agentic AI, where systems operate with increasing autonomy and complexity, organizations must prioritize real-time oversight: Continuous monitoring and rapid response capabilities are essential to navigate the intricacies of these technologies. The autonomous nature of agentic AI introduces new risks that can escalate quickly, making robust controls necessary to prevent costly disruptions and ensure system integrity" notes Sinclair Schuller, EY Americas Responsible AI Leader.

One particular area that is lagging is preparation for a hybrid AI-human workforce. Only a third (32%) say their HR team is developing strategies for managing such environments. Still, given the nascent nature of agentic AI, this figure can be seen as promising since it indicates that companies are starting to think strategically about longer-term implications of the technology.

Citizen developers – opportunity or blind spot?

The rise of "citizen developers" – employees using no-code or low-code tools to create their own AI agents – presents a more complex challenge.

A third (32%) of companies have chosen to prohibit the practice outright. Among the rest, tolerance ranges from tightly limited use cases to active encouragement, with some businesses even promoting best practices across teams.
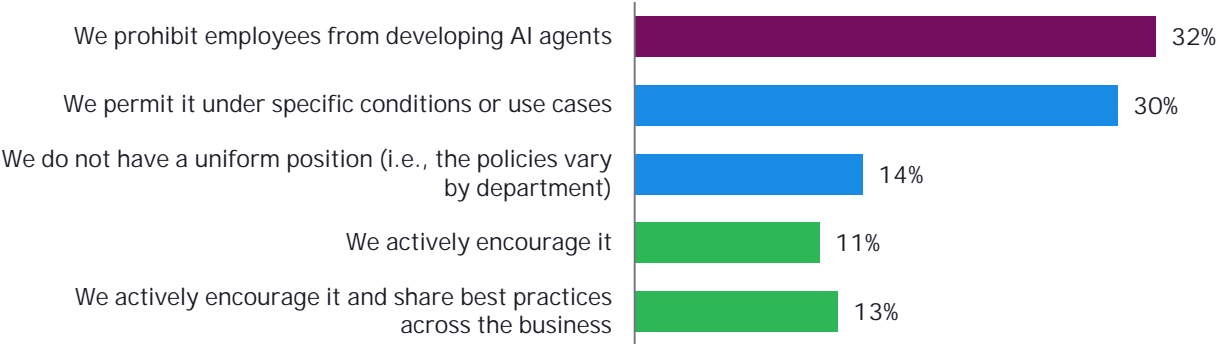

Response	%
We prohibit employees from developing AI agents	32%
We permit it under specific conditions or use cases	30%
We do not have a uniform position (i.e., the policies vary by department)	14%
We actively encourage it	11%
We actively encourage it and share best practices across the business	13%

Fig 7. Two-thirds of organizations allow citizen developers to create AI agents

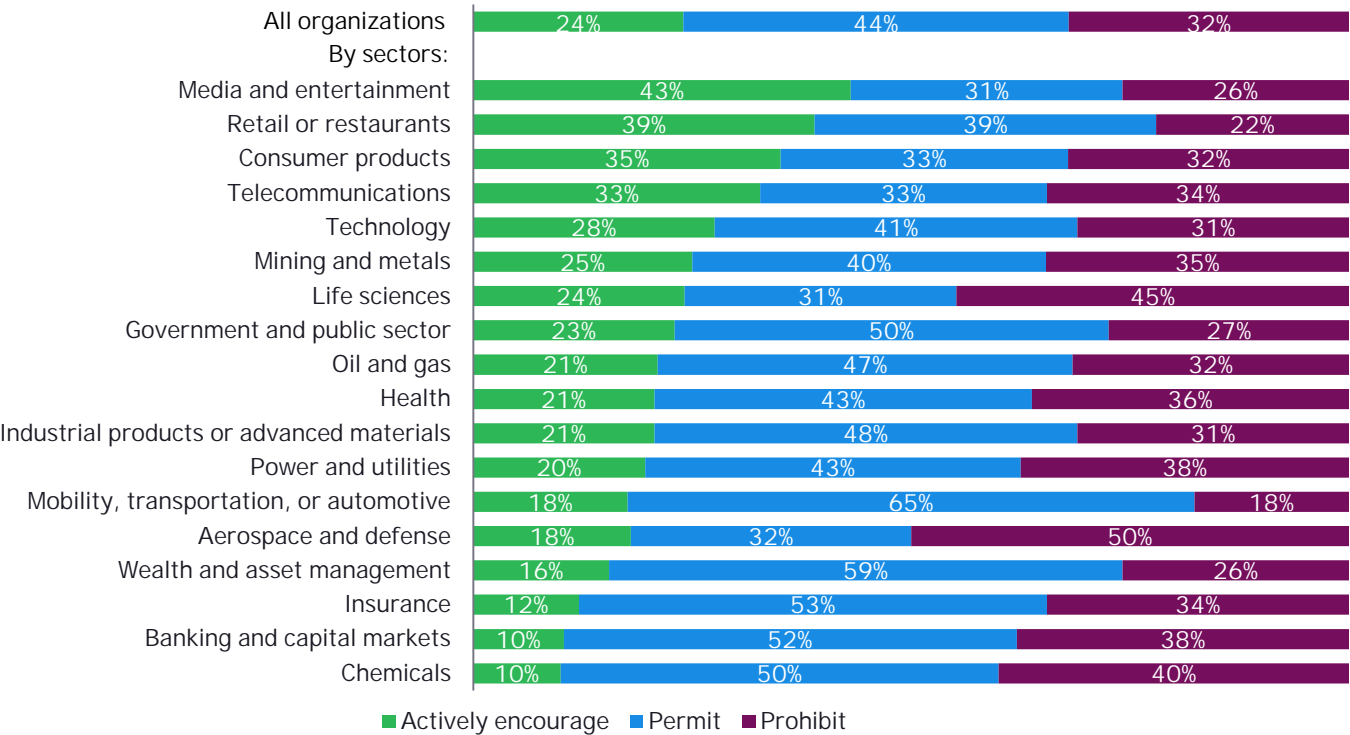| Sector	Actively encourage	Permit	Prohibit
All organizations	24%	44%	32%
By sectors:			
Media and entertainment	43%	31%	26%
Retail or restaurants	39%	39%	22%
Consumer products	35%	33%	32%
Telecommunications	33%	33%	34%
Technology	28%	41%	31%
Mining and metals	25%	40%	35%
Life sciences	24%	31%	45%
Government and public sector	23%	50%	27%
Oil and gas	21%	47%	32%
Health	21%	43%	36%
Industrial products or advanced materials	21%	48%	31%
Power and utilities	20%	43%	38%
Mobility, transportation, or automotive	18%	65%	18%
Aerospace and defense	18%	32%	50%
Wealth and asset management	16%	59%	26%
Insurance	12%	53%	34%
Banking and capital markets	10%	52%	38%
Chemicals	10%	50%	40%

■ Actively encourage ■ Permit ■ Prohibit

Fig 8. Two-thirds of organizations allow citizen developers to create AI agents

What should concern leaders most is the inconsistency between stated policy and real-world oversight. Among organizations that allow citizen developers, only 60% have formal, organization-wide frameworks to help ensure alignment with responsible AI principles and just half have high visibility into actual activity. And even among companies that prohibit the practice, 12% admit they lack visibility into actual activity – creating a governance gap where shadow AI development can flourish undetected – meaning they are essentially flying blind.

The emergence of agentic AI and citizen developers underlines a central theme of our findings: responsible AI must evolve in step with technology and workplace behaviors. Clear frameworks, proactive oversight and leadership awareness are critical if organizations want to seize the benefits of these trends without compounding their risks.

Implications for business leaders

Here are three actions C-suite leaders can take to strengthen their AI governance and controls, and boost business outcomes:

1. Adopt a comprehensive approach to responsible AI

The symbiotic relationship between responsible AI adoption and AI-driven performance improvements has a clear message for business leaders. To get more value from your AI investments – particularly in crucial areas like financial performance and employee satisfaction – it's critical for companies to move further along the responsible AI journey. A comprehensive approach includes articulating and communicating your responsible AI principles, executing on them with controls, KPIs and training, and establishing effective governance.

2. Fill knowledge gaps in the C-suite

AI affects every facet of your organization. It's critical for leaders across the C-suite to understand both its potential as well as its risks – and the controls needed to mitigate them. Yet, our survey reveals significant shortcomings in awareness of which controls are most appropriate.

How does your C-suite compare? Identify where the biggest gaps are and fill them with targeted training. At a minimum, ensure the roles closest to AI risks are well versed on the appropriate safeguards.

3. Get ahead of emerging agentic AI and citizen developer risks

Agentic AI promises powerful new capabilities, but it also brings significant risks. It's critical for organizations to identify these risks, adopt appropriate policies, and ensure appropriate governance and monitoring is in place.

There are gaps between companies' stated policies and their insight into whether employees are developing their own AI agents. Understand the costs and benefits for your organization before establishing your position. And – regardless of whether you ban, permit, or encourage the practice – make sure your policy is backed up with real insight into what employees are actually doing.

Summary

As AI becomes more deeply embedded in business operations, leaders face a clear choice: treat responsible AI as a box-ticking exercise or as a strategic enabler. Those that take the latter path are already proving that robust governance, clear principles, and informed leadership can turn potential risks into competitive advantage. The next wave of developments – from agentic models to citizen development – will only raise the stakes. Success will belong to organizations that act now to align responsibility with performance.

Contact us for more information:

Kelvin Gao

Managing Partner
Greater China Cybersecurity Services
Ernst & Young (China) Advisory Limited
kelvin.gao@cn.ey.com

Lena Wang

Partner
Greater China Cybersecurity Services
Ernst & Young (China) Advisory Limited
lena.wang@cn.ey.com

Deft Chen

Manager
Greater China Cybersecurity Services
Ernst & Young (China) Advisory Limited
deft.chen@cn.ey.com

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multi-disciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

ey.com/china

Follow us on WeChat
Scan the QR code and stay up-to-date with the latest EY news.