pwc

The AI trust dividend

Creating strong foundations, trusted AI, and real returns

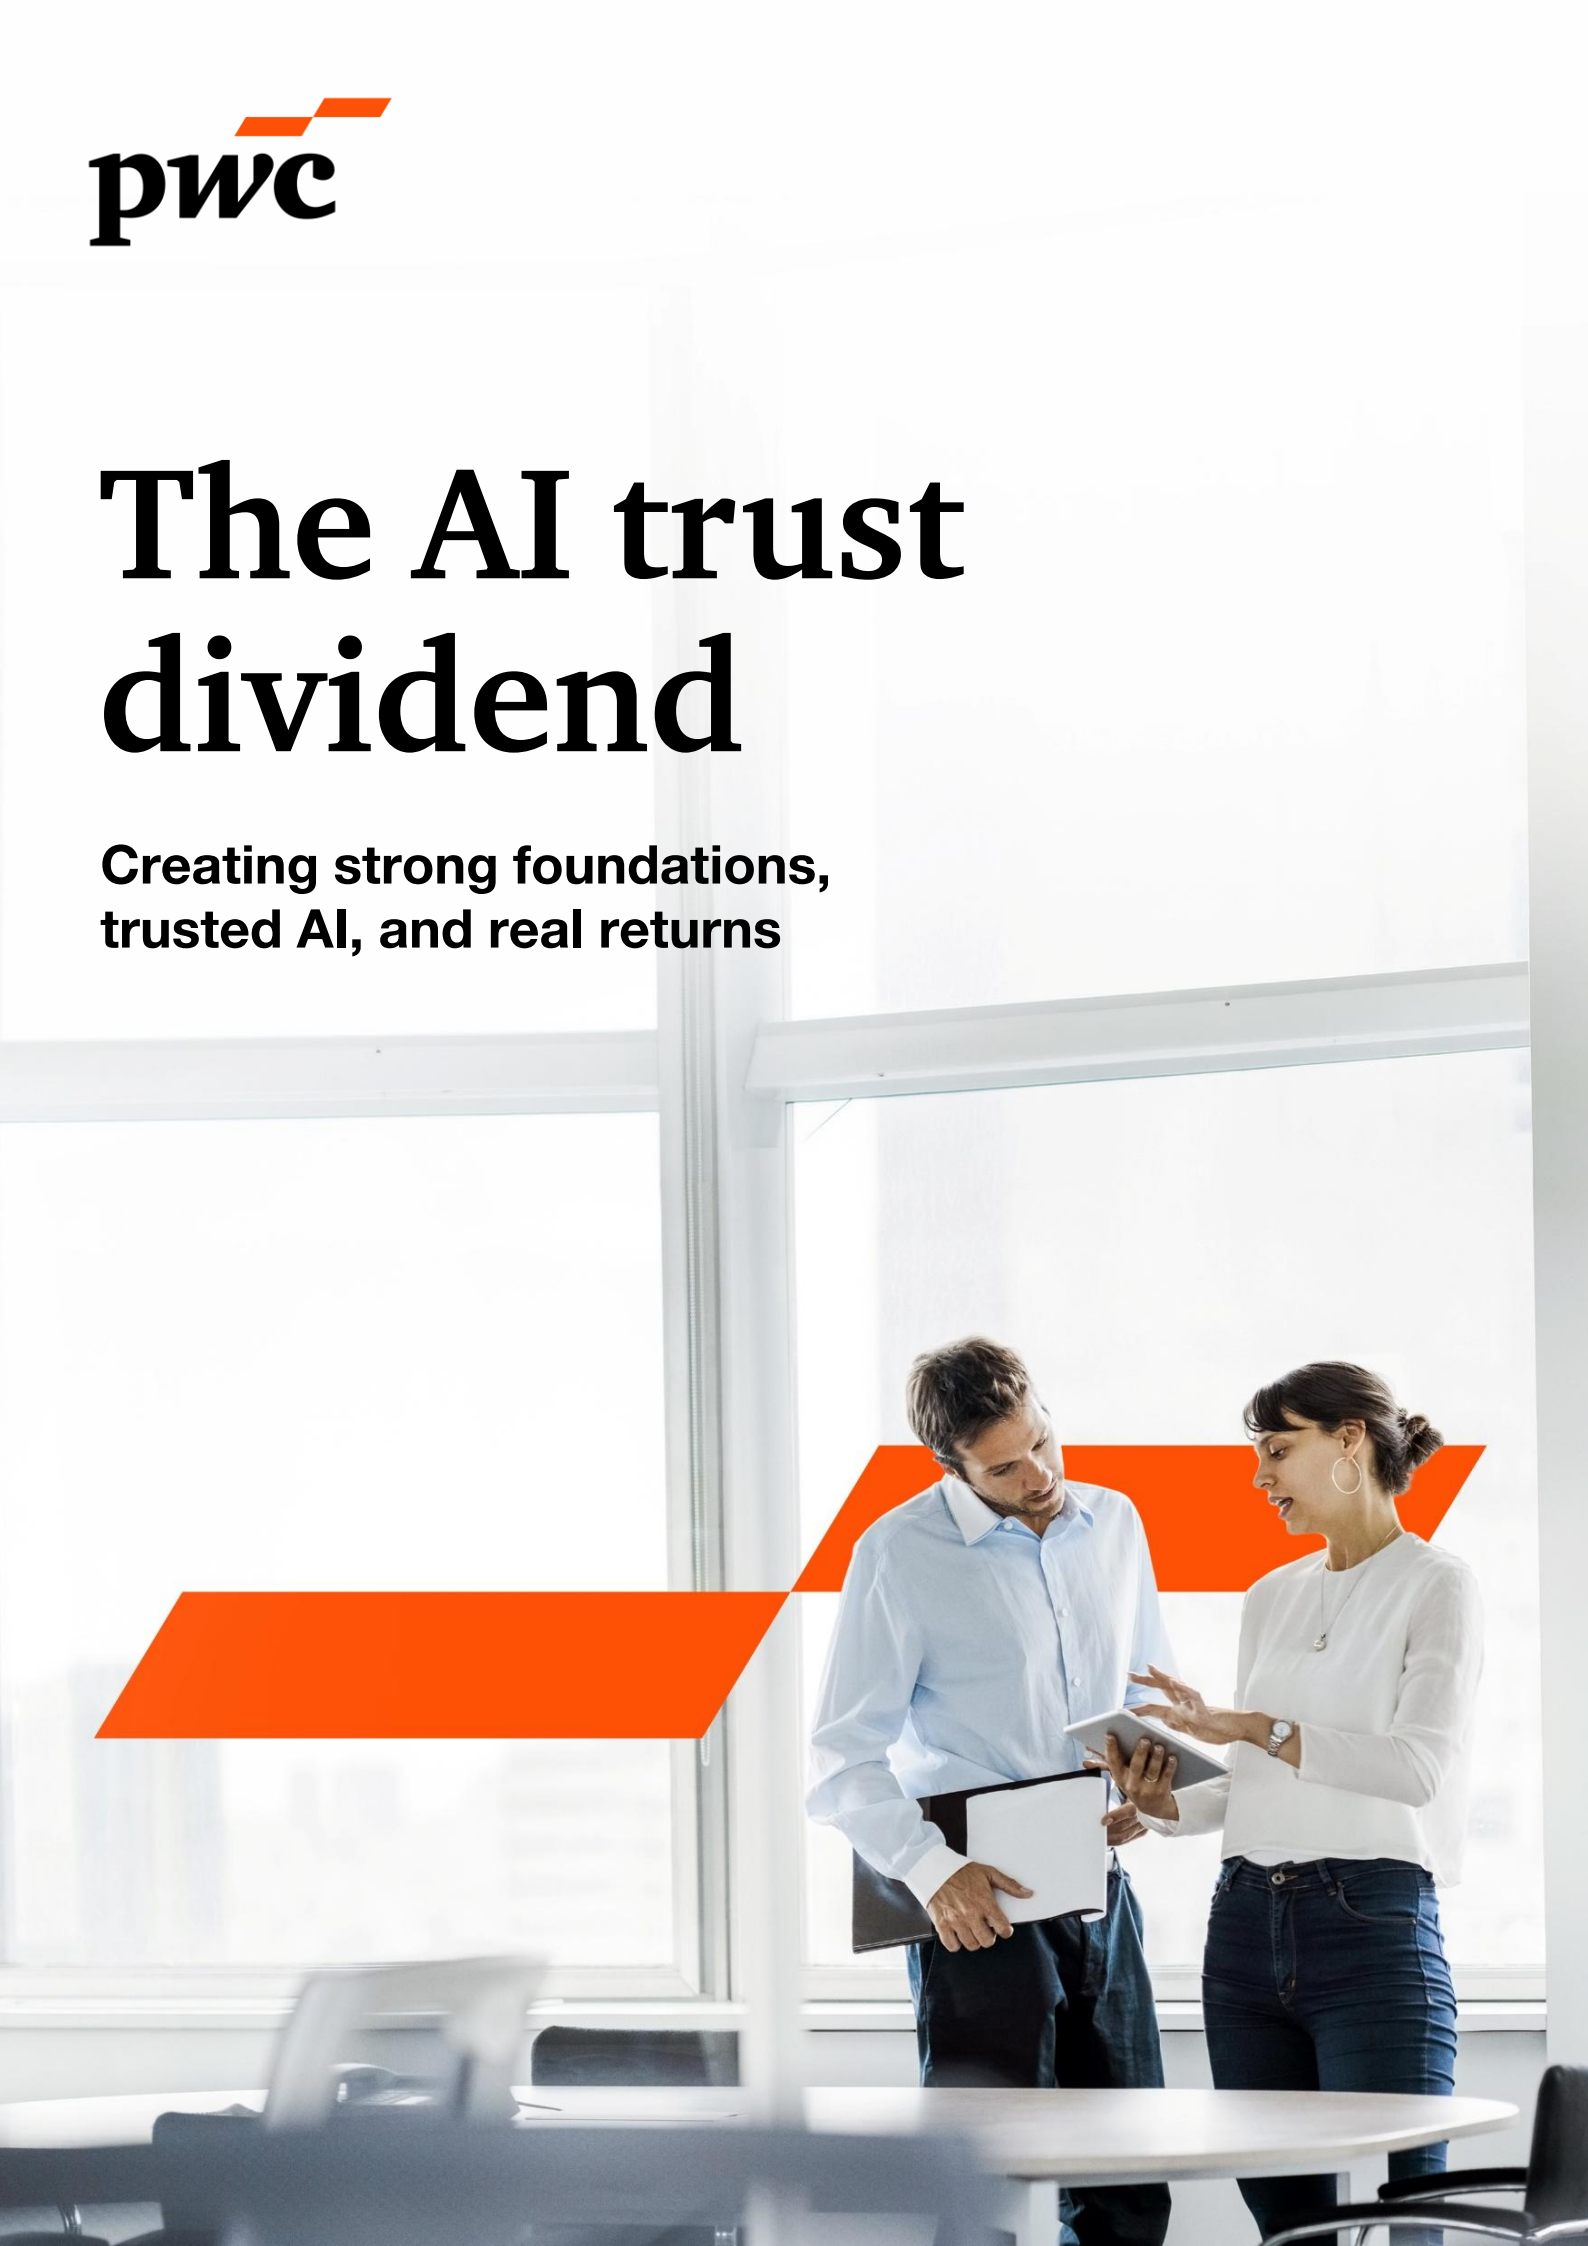

Contents

An AI trust gap is rapidly emerging as the pace of AI capability exceeds what organisations can confidently and safely deploy. While many organisations are caught straddling the gap, some have closed the divide.

By Leigh Bates and Darren Henderson

The most advanced AI adopters are moving fast. They aren't just ahead in AI adoption, they're ahead in their ability to trust in the outcomes AI delivers, which compounds their capability to deploy AI faster, scale further, and capture more value than their peers.

'The AI trust dividend' series will tackle the tough questions organisations are grappling with as they race to achieve tangible return on investment from AI and deploy AI with confidence.

Today's reality

AI experimentation is surging, but many organisations find that AI activity is failing to produce measurable returns. In fact, only

26%

of CEOs in **PwC's 29th Global CEO Survey** said their companies see lower costs because of AI, and more than half (56%) said they have realised neither revenue nor cost benefits.

Despite this reality, some organisations are cracking the code on AI RoI in a significant way. Our **2026 Global AI Performance Study** finds that AI value is currently concentrated in a small cohort: 20% of companies capture

74%

of all AI-driven value. These top-performing organisations, which we call AI leaders, obtain the greatest measurable value from AI across revenue, efficiency, and cost outcomes. In fact, they achieve 7.2x more AI-driven revenue and efficiency gains than their peers.

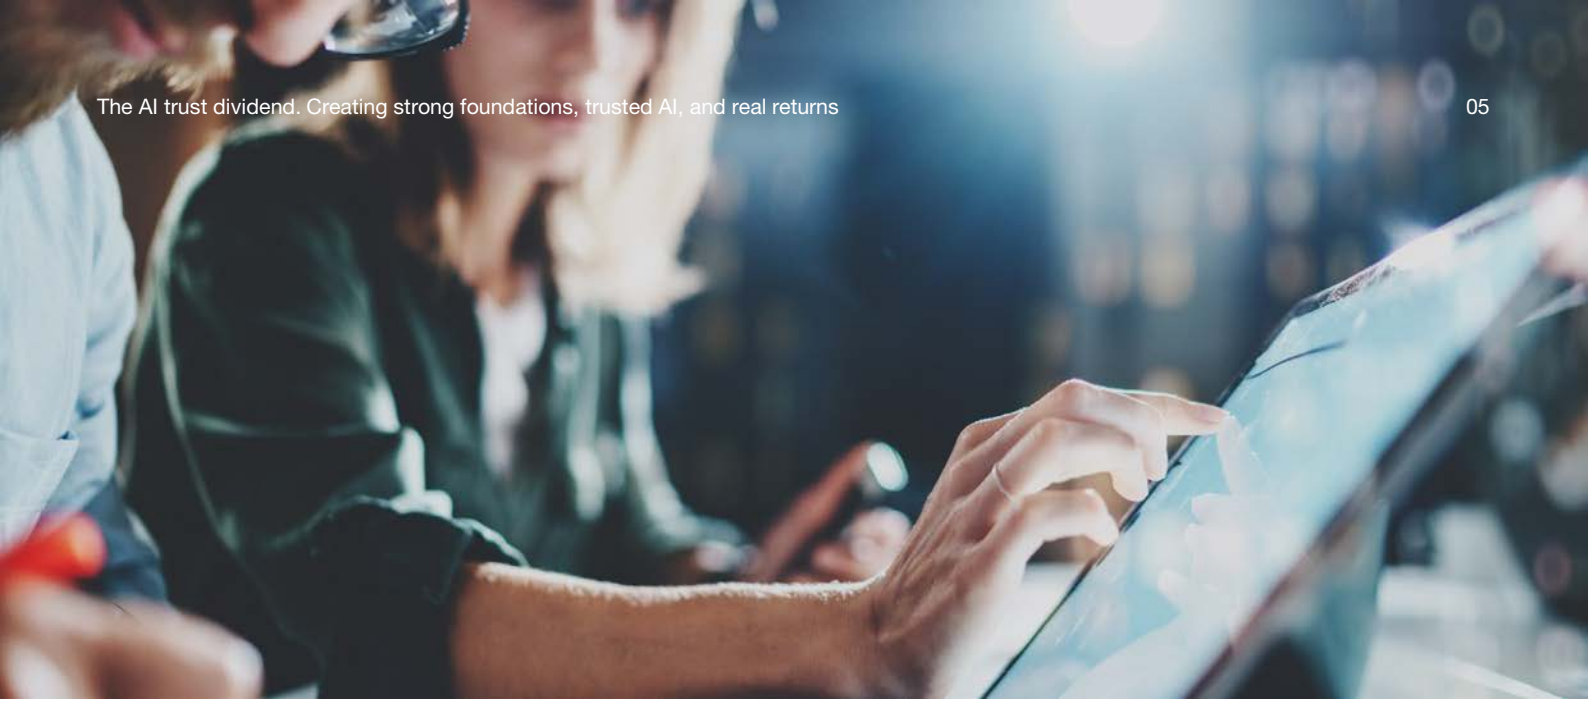

01

Building trust to scale AI

Organisations leading with AI aren't just increasing its use. They're building trusted foundations to scale AI effectively, then strategically applying it for significant financial gains and transformative impact. By doing so, AI leaders are getting nearly double the improvement in AI-driven performance than those with weaker foundations.

Trust is not a soft concept, and its economic impact is clear. Our **2026 CEO survey** affirmed that public companies experiencing the fewest trust concerns delivered total shareholder returns over a 12-month period that were nine percentage points higher than those experiencing the most trust concerns. Trust in AI drives speed to value and the confidence to pursue more transformational use cases.

In most organisations, the challenge surrounding AI isn't building it. It's getting it approved to deploy, adopted by employees, and accepted by customers.

Each of these is a hurdle. Organisations that create a high-trust environment, clear these hurdles in weeks. Others take months or never clear them at all.

When leaders can trust that the organisation has the right control frameworks to protect itself, its stakeholders and its customers, they can create a bold AI vision. Confident that risks they take are within their appetite, leaders can aspire to more impactful and innovative use cases. Not only can they use AI to build greater efficiency; they can begin to reimagine processes and see product and market opportunities. Employees adopt AI tools faster with confidence, and regulators become satisfied there are robust controls in place and accountability to mitigate risk.

Ultimately, trust creates a flywheel that provides a recurring cycle of advantages. Organisations deploying trustworthy AI gain greater stakeholder confidence, leading to greater use, richer data, more value, and reinforced trust.

02

Managing AI risks to build trust

The ways that organisations manage AI risks today are static, linear, and built for a time when one new impactful model would be released per month. As AI permeates every product, process, decision, and interaction, new agents will be created daily and constantly tuned in production. In this environment, a periodic human activity will not be sufficient to manage risk. It necessitates an automated, continuous capability applied at the point in time where it can effectively and efficiently mitigate risk.

The leaders in this new model standardise controls, automate their application, and shift complexity away from end users and development teams. They are also willing to take calculated risks on AI systems for routine operations while minimizing risks related to AI that impacts customers, major financial decisions or regulatory requirements. These approaches are not minor tweaks; it's a fundamental rethink of managing risks in an AI-driven world.

Snapshot: Building an AI risk playbook and testing framework

Recognising the need for consistent, ethical, and risk-aware practices across the AI lifecycle, a global pharmaceutical distributor is strengthening its AI development processes and governance. Together with PwC UK, the team is embedding responsible AI principles in a way that will be usable across different stakeholder groups and technical maturity levels. The company's new AI governance playbook and testing framework will support scaling AI responsibly by ensuring consistent governance, reducing risks, and improving transparency across teams. It will also help align internal capabilities with evolving regulations and industry expectations.

Leaders are wrestling with many areas as they work to build trust, scale AI, and realise its benefits while the technology's capabilities expand. 'The AI trust dividend' series shares the practices and experiences of those who are progressing in these areas. Here, we start with two traits helping advanced adopters manage AI risk to build trust.

01 Clear procedures help move use cases forward based on proportionate risk

Organisations that understand proportionate risk, and how risks associated with each use case will be addressed, are able to move faster than those that don't. This requires interaction spanning areas such as system accuracy, cyber security, data quality and management (including role-based access controls), data ethics, and resilience. Organisations pulling ahead have sufficient and accessible information to help people evaluate AI systems pre-deployment and monitor their behaviours once those systems are live.

AI leaders distinguish themselves by building structured, repeatable controls.

They are

1.7x

more likely to have a Responsible AI framework in place.

Source: PwC's AI performance study

AI leaders are

1.5x

more likely to operate a cross-functional AI governance board embedded into development and deployment decisions.

Source: PwC's AI performance study

Case in point: One large bank accelerated its AI governance process by using automation to extract AI use case information directly from code repositories.

The solution identified key details such as the business purpose, process, data sources, model components, dependencies, and controls, then used this information to generate the risk assessment for dev teams to address risks. This helped reduce manual effort, improve controls consistency and documentation quality, and give governance teams greater real-time visibility of AI systems at scale.

Snapshot: Turning governance into an accelerator

One of the UK's largest utility companies wanted to accelerate AI innovation and adoption while strengthening governance across the business. The organisation recognised that scaling AI successfully meant striking the right balance between unlocking commercial value and managing risks relating to data, reputation, compliance, and trust. Working with PwC UK, the company developed a Responsible AI Framework and Operating Model aligned to its values, code of conduct, and sustainability goals. We provided independent assurance, stress testing, and practical implementation guidance to help embed the framework across the organisation.

The framework applies governance proportionate to the level of risk in each AI use case. By introducing clear risk tiers, ownership and accountability, it gives teams the confidence to innovate faster while ensuring appropriate controls are in place.

The results were clear. The company is better equipped to scale AI responsibly, empowering teams to move from experimentation to adoption with greater clarity, consistency, and trust.

02 Accountability sits in the business

You can get AI to do 95% of the work, but 100% of the accountability sits with us, humans. Moving governance closer to the teams deploying AI provides clear lines of accountability. It also creates the opportunity to embed governance directly into how AI systems are designed and deployed.

With AI, the risks are substantially determined by design decisions made early in development, such as data selection, model architecture, training methodology, and fairness criteria. If the second line takes a traditional review and challenge approach after development, it is too late to influence the most consequential decisions. AI leaders are placing accountability in the first line and embedding governance into the build. This is a major shift since many organisations started their AI deployment as a technology-led initiative in the office of the CIO or CTO. Those moving the fastest with AI have clearly tied risk management to the front line of the business.

At the same time, the third line must develop new assurance approaches. Traditional assurance methodologies are not designed to assure AI systems that learn, adapt, and change continuously. Internal Audit needs new tools, skills, and approaches, to provide continuous assurance over AI models and outcomes, the data and infrastructure supporting them, and decisions made using them. Independent third-party assurance may also be needed for an increasingly wide range of stakeholders who are looking for companies to demonstrate good governance and controls over AI (including regulators, investors, customers, and business partners).

Governance responsibilities are shifting to builders and operators.

56%

of the executives say their first-line teams now lead Responsible AI efforts. Moving forward, we anticipate seeing a broader shift from tech-driven to business-led AI initiatives as experimentation by tech teams rapidly evolves to create tangible returns.

Source: **PwC's Responsible AI Survey**

Case in point: A large multi-national financial services institution implemented an AI governance team within its business-aligned AI strategy function. While this team does not own the AI policy, they are acting as system design thinkers, creating mechanisms to satisfy the requirements of the policy – such as standardising and automating AI risk identification – and embedding those into the core AI delivery lifecycle.

Snapshot: Operationalising AI Governance Across the Enterprise

Faced with fragmented AI governance and the challenge of orchestrating AI systems consistently across the organisation, a global chemicals company is moving its AI transformation beyond design into integrated, organisation-wide adoption. Together with PwC Germany, the team is implementing an AI governance model aligned with the EU AI Act - standardising guidelines, aligning with existing risk frameworks, and establishing a structured management system. Rather than creating a parallel structure, existing governance was leveraged, a dedicated AI governance tool deployed, and regulatory requirements integrated step by step into the live governance setup. The result is a structured transformation that moves AI governance from design to operationalisation - embedding best practices into existing structures and driving sustained adoption across the organisation.

03

The window to act is narrow

The mindset in many boardrooms is to prioritise using AI where it can deliver the most value, while recognising that governance and risk management teams are playing catch up. Without trusted foundations, our view is that this approach can both stop short of delivering returns and expose the organisation to untenable risks. Now is the time to act.

As AI deployment grows, the likelihood of a major AI trust incident becomes certain. Organisations with mature foundations will handle incidents as operational events to be contained, resolved, and learned from. Those without this experience will face existential crises: regulatory scrutiny, leadership accountability, customer loss, and reputational damage.

Furthermore, the pace and democratisation of AI adoption is accelerating. Every month an organisation delays investing in the right AI foundations, AI capabilities continue to advance while the organisation’s ability to scale falls further behind.

Snapshot: Testing the Untestable

A UK-based bank aimed to embed a GenAI-assistant productivity tool directly into employees' workflows across risk, HR, lending, compliance, and operations. To do so, it needed to devise a way of testing to understand how a GenAI assistant works, where it boosted productivity, where it faltered, and how to mitigate risks.

Together, PwC UK and the bank developed a testing framework and innovated testing technology, which helped the bank validate and implement the GenAI tool for its staff. The joint team ran live tests over 1,400 assistant outputs across multiple use cases to build a full picture of errors, hallucinations, bias and toxicity, using AI to test AI - or 'LLM As-a-Judge' technology. This complemented the more traditional testing approaches using statistical techniques and specialist human testing experts. Through structured prompt optimisation and post-testing enhancements, the GenAI assistant's accuracy on complex tasks increased from 40% to 85%.

After testing, the team found some risks were easily controlled while others were outside of the bank's risk tolerance. Additional controls and guardrails were developed to address these risks and satisfy regulatory due diligence. With a 'safe use' framework now in place, the business has laid the foundations for an AI-enabled future.

Building trust in AI requires new thinking. Leaders that are accelerating AI RoI are maturing in their ability to do things differently. They recognise that solid foundations are the key to unlocking AI value, and that trust underpins them all.

Here are five questions to ask within your organisation and actions to consider based on where you are on your AI journey.

1 Do you have a clear posture on AI risk appetite?

Build a structured hierarchy that connects board-level posture to specific application decisions and the operational thresholds that govern them. Such a shared framework creates a consistent, board-endorsed space within which innovation can move at the speed the organisation actually intends. As a foundation, document risk appetite. As you advance, automate threshold-based triggers.

2 Is it clear where accountability and responsibility sit?

Accountability and responsibility should be clearly assigned, with ownership sitting primarily with the business area using the AI. They are best placed to understand the intended outcomes, risk tolerance, and how the AI will be used in practice. Clear roles should also be defined for risk, technology, and governance teams

3 Is AI governance and risk management embedded into your AI development with a workforce that understands what is expected?

AI governance and risk management should be integrated into development by design from the start, with controls scaled to the level of risk. Higher-risk AI should receive more detailed evaluation, testing, and monitoring. Teams should have clear responsibilities, practical guidance, and training so they understand safe use expectations and when to escalate issues.

4 Do you know what AI use cases and systems exist and how to keep an up to date view?

Organisations should maintain an up-to-date inventory of AI use cases and systems, starting when ideas move into design. The level of tracking should be proportionate to risk, with priority given to AI that affects customers, key decisions, regulatory obligations, or financial reporting. Higher-risk use cases should have ongoing senior visibility.

5 How do you test AI before deployment and then how do you know it's behaving as intended once deployed?

Before deployment, AI should be tested using real examples, edge cases, and scenarios where it may fail. This helps confirm expected performance, limitations, and key risks. Once live, results and behaviour should be continuously monitored to detect changes, warning signs, or unexpected patterns, so issues can be addressed quickly.

This is the first article in ‘The AI trust dividend’ series. In future articles we'll go deeper on the challenges business leaders need to address at different points in their AI maturity to unlock its full potential.

About the authors

Leigh Bates

Partner, Global Risk AI Leader, PwC UK
LinkedIn: www.linkedin.com/in/lbates/
Tel: +44 (0)7711 562381
Email: leigh.bates@pwc.com

Darren Henderson

Partner, Global Trust and Transparency Leader, PwC Canada
LinkedIn: www.linkedin.com/in/darren-henderson/
Tel: +1 416 941 8379
Email: darren.henderson@pwc.com

Just 20% of companies are capturing 74% of all AI-driven value. We've decoded how, so you can harness AI to drive productivity, reinvention, and growth.

Explore the article:

Decoding ROI from AI

With our trust architecture approach, we build in trust to unlock value across your operations, governance, and tech—driving consistent, credible, enduring impact.

Find out more:

Global trust services

Thank you

Access the online version