



AI's Second Act, Where Safety, Risk and Trust Collide

The first AI cycle was defined by experimentation, ambition and the race to move fast. AI's Second Act will be defined by a harder test: whether businesses can turn AI ambition into credible, resilient and commercially beneficial results

Contents

Foreword	04
Ambition Gives Way To Caution	06
The Implementation Disconnect Increases	09
Agentic AI: The Conundrum Persists	11
AI Risk: A Compounding Matrix	13
National Approaches Pull Apart.....	15
About FTI Consulting	19



AI's Second Act, Where Safety, Risk and Trust Collide

The first AI cycle rewarded ambition and speed. AI's Second Act will reward proof: it has never been more important for companies to ensure what they say about AI matches up with reality.

The first act was defined by experimentation, bold public claims and a race to stay ahead of the competition. This next iteration is quieter, more deliberate and considerably more demanding.

Organisations are no longer asking whether they can deploy AI, but whether they can do so at scale credibly, safely and sustainably and whether the story they tell investors, customers and employees survives contact with reality.

Our research shows businesses advancing with deliberate caution. **Two-thirds of organisations are moving forward carefully**, rather than racing ahead, and a clear majority have already slowed, paused or pulled back a planned deployment because of reputational, regulatory or trust concerns.

In spite of the known risks around AI, only around a third feel very well prepared for an AI-related incident, demonstrating that confidence is still outpacing resilience.

So the Second Act is not a retreat from AI, but a recognition that integrating it successfully costs more than businesses first thought and that the cost is not purely financial. It manifests across operational, reputational and personnel faultlines.

The upside remains substantial: returns are real and rising, and governance is maturing. But value will increasingly only be recognised by companies that combine disciplined investment, credible governance and a clear stakeholder narrative; not to those making the boldest claims. The pages that follow set out where ambition and caution collide, and why it has never been more important to ensure AI narratives and the operational realities match up.

Foreword

Caution is the default.

Two-thirds of organisations are advancing carefully rather than racing ahead. **The challenge is turning that caution into a practical AI strategy that enables progress without taking unmanaged risk.**

Confidence outpaces resilience.

Only around a third feel very well prepared for an AI-related incident. **That makes AI resilience increasingly important, from understanding exposure and preparing response plans to testing how the organisation would react when something goes wrong.**

The cost is organisational, not just technological.

54% say employee use of unapproved “shadow AI” tools is a significant concern. Getting the operating model right can help organisations focus investment on the capabilities, controls and skills that will make the biggest difference.

Narrative risk is now commercial risk.

60% have slowed, paused or pulled back a planned AI deployment because of reputational, regulatory or trust concerns. Aligning what an organisation says about AI with what it can demonstrably deliver is becoming a critical part of protecting trust and reputation.

The opportunity remains.

Returns are rising, with senior executives significantly more likely than middle managers to report strong returns from AI. The opportunity is to move from experimentation to scalable value, combining the right use cases with the data foundations, governance and execution needed to deliver returns.

“We are now looking at a step-change in the way companies approach AI adoption and implementation. No longer is it sufficient to stay ahead of the competition with bold claims and big budgets. Companies must show a considered and credible approach to AI to their stakeholders and upweighting safety and caution is increasingly how they do that.”



Jon Priestley, Partner, Digital, Data & AI

Ambition Gives Way To Caution

The allure of being seen to lead is now weighed against the need to be trusted, and most organisations appear to be choosing credibility over speed. Looking five years out, our research shows that more than twice as many businesses see themselves as a trusted operator than as a clear market leader.

Yet the day-to-day picture is more conflicted. Quick-win automation continues to shape AI strategies, even as more than half of organisations say AI is already changing how they allocate capital for the long term. The commercial implication is clear: tactical AI choices are quietly becoming strategic capital allocation decisions, frequently without the governance that would normally accompany a commitment of that size.

The true investment required for transformational AI also goes well beyond the technology itself. Businesses expect additional investment in cyber security, change management, training and governance, meaning the cost of scaling AI is as much organisational as it is technological.

And caution is the mode of the day. **Two-thirds of organisations are moving forward carefully rather than racing ahead**, and a clear majority have already slowed, paused or pulled back a planned deployment on reputational, regulatory or trust grounds. That is a deliberate strategic posture, and one boards should be able to articulate as confidently as any acceleration story.

The risk is that caution lacks a credible narrative. A business that pauses a deployment for sound reasons, but cannot explain those reasons to investors, regulators or its own workforce, looks indistinguishable from one that has simply fallen behind. In this new stage of AI the discipline of the decision counts for little without the credibility of the explanation beside it.


If your next AI decision slows progress in order to protect trust, can you prove it is disciplined strategy rather than lost momentum? How would you create the perfect narrative to communicate that?

“The rush to lead on AI created a powerful incentive to move first and ask questions later. As concerns around safety, jobs and the societal impact of AI become more front of mind, that calculation is changing. The next test of AI leadership won’t just be how much technology a company deploys, but whether it knows where AI creates value and where using it could destroy it.”




Damian Low, Managing Director,
Digital & Insights

Q. In the past 12 months, has your organisation slowed, paused or pulled back a planned AI deployment because of reputational, regulatory or trust concerns? (Please select one response) - per market




In France and Germany there is little slowdown to speak of, because momentum never fully built: three in four businesses are moving carefully and only one in five is moving fast. That caution reflects European sensitivities hard-earned from previous technology and data cycles.

Q. Which of the following best describes your organisation's top-level approach to AI adoption?
(Please select one response)



Q. In the past 12 months, has your organisation slowed, paused or pulled back a planned AI deployment because of reputational, regulatory or trust concerns? (Please select one response) - per sector



The Implementation Disconnect Increases

A persistent divide runs through our data between the C-suite and middle management. Leaders consistently report more speed, stronger returns and tighter alignment, while managers closer to delivery report more caution, more risk and a wider gap between message and reality. Across the data, AI appears more successful the further up the organisation you go.

That perception gap matters because it risks inflating investment cases, weakening governance and quietly eroding internal credibility, because the people being asked to deliver AI recognise the difference between the reported position and the operational one.

Our research finds that nearly two-thirds of leaders feel pressure to emphasise or even overstate AI in corporate and investor narratives. Yet only half the C-suite believes external messaging fully matches what is happening inside the business, and that figure falls sharply among middle managers. More than half of middle managers we

surveyed doubt or are only moderately confident that other senior leaders in their sector genuinely believe their organisation's public AI leadership claims, while only 43% trust claims made by others.

Businesses may therefore be entering an AI credibility arms race, in which public positioning accelerates faster than operational proof. The practical consequence is that the most accurate account of an organisation's AI position usually sits several layers below the people describing it publicly. Closing that gap is a governance task before it is a communications one: boards need an unfiltered line of sight into delivery, and managers need a route to report reality without it being interpreted as negativity or dissent.


What would change if your board heard the unfiltered AI reality from the people closest to delivery before it heard the polished version? And how would you manage that?

"The biggest AI risk for many organisations is no longer a lack of ambition; it is the gap between what the board believes is happening and what delivery teams know is happening. Closing that gap means creating a governance rhythm that brings operational reality into the boardroom before it becomes reputational reality outside it. The organisations that build those feedback loops now will make better investment decisions, avoid overclaiming and earn greater trust from their people and stakeholders."




Lucy Phillips, Partner, People & Transformation

Q: How closely do your organisation's public statements about AI match what is actually happening inside the business? (Please select one response)



Q: Overall, how much do you trust the AI leadership claims made publicly by other companies in your sector? (Please select one response)



Agentic AI: The Conundrum Persists

AI is becoming more broadly embedded in businesses, but humans are reluctant to give up their power when it comes to decision making. Nearly a third of respondents (29%) say humans remain firmly in control with only 13% saying AI is increasingly driving decisions. For all the advances in capability, businesses are not yet comfortable with widespread agentic processes and this represents a commercial and operational conundrum.

The commercial case for autonomous AI rests on the system acting without a person waiting at every step, yet the overwhelming majority of organisations have deliberately designed their deployments so that human judgement is omnipresent. The contradiction is that technology is being bought on a promise of autonomy yet implemented on an assumption of supervision.

This is not, at root, a technology or a funding problem. Vendors can demonstrate the capability and boards can find the budget. What agentic AI actually asks for is a mindset shift: business leaders relinquishing direct control over decisions, workflows and customer interactions that have historically defined their authority and their accountability.

Control is not surrendered to AI in a strategy document. It is surrendered through delegated authority limits, a stated tolerance for errors made without a human in the loop, and a governance model that can explain an outcome nobody explicitly approved. Few organisations have built any of the three, which is why agentic pilots tend to stall at the point of removing the human check rather than at the point of technical capability.

Seniority compounds the problem. Managers closer to delivery report more caution and a greater perception of risk than the C-suite, so those most exposed to an autonomous system's mistakes are also the least persuaded it should run unattended. Until that disconnect is reconciled, agentic AI will keep being announced faster than it is adopted.


Which decisions are you genuinely prepared to let AI take without a human safety net and what evidence would make that delegation defensible to your internal and external stakeholders?

"The technology is already far ahead of most people's willingness to trust it. Those who automate with the right guardrails will pull away on efficiency, those who wait will spend the next few years catching up."




Rob Stone, Managing Director, Digital & Insights

Q: To what extent are AI systems influencing or driving business decisions versus serving strictly as tools under human direction? (Please select one response)



Q: How significant are these talent-related challenges in your organisation? (Please select one response per statement) - Gap between AI tools purchased and actual employee usage (response: significant or very significant)



AI Risk: A Compounding Matrix

Not surprisingly, then, the AI risk landscape is broadening faster than the controls built to manage it. Some 54% of organisations that we surveyed cite shadow AI (employees using unapproved AI tools) among their biggest risks for the coming year, on a par with fraud, scams and deepfakes.

Seniority changes the picture again. 61% of middle managers cite unreliable AI outputs versus 53% of the C-suite as a key risk for the coming year, while 57% cite a lack of employee AI skills or governance versus 47% of the C-suite. Those closer to delivery consistently see greater risk in how AI is actually being used and controlled day to day, which suggests the most accurate risk register in most organisations is not the one the board is reading.

AI risk is also spreading well beyond the technology function. While 47% expect increased investment in cyber security, significant numbers also expect to invest more in change management (35%), employee training (30%), AI governance headcount (27%) and crisis management (26%). As exposure extends into reputation, regulation,

people and business continuity, managing AI becomes a whole-business responsibility rather than one that can sit primarily with technology teams.

Preparedness has not kept pace. Only around a third of organisations feel very well prepared for an AI-related incident, even as a clear majority have already slowed, paused or pulled back a deployment on risk grounds. Recognising a risk and being ready for it are demonstrably different things.

The practical implication is that AI creates a risk matrix which needs an owner outside the technology function, a rehearsed response, and a single escalation route that runs from the people using the tools to the people accountable for the consequences.


If your people are already using AI outside approved channels, do you know whether that is innovation you should harness or exposure you must contain? And how would you find out?

“When we classify crises we look at whether they are internally and externally driven. For example, the threat of Shadow AI is internally driven, we would look at it as a governance crisis. So, ownership of its use and responding to its failures needs to be from the very top, the Board & Executive.”



Kate Brader, Partner, Crisis & Litigation

Q: If a serious AI incident occurred tomorrow, how well protected is your organisation? (Please select one response)
 “Very well protected — tested response plans in place”



Q: Which of the following do you see as the greatest AI-related risks facing your organisation over the next 12 months?
 (Please select the top five risks)

	All markets	United Kingdom	France	Spain	Germany	Belgium	USA	Ireland
Data privacy or security breaches	60%	72%	58%	60%	64%	38%	66%	64%
Incorrect or unreliable AI outputs leading to poor decisions	57%	48%	64%	62%	63%	51%	54%	60%
Regulatory or legal non-compliance	51%	60%	42%	46%	52%	55%	51%	51%
Reputational damage from AI misuse	44%	48%	49%	44%	38%	41%	50%	37%
Employees using unauthorised AI tools ("shadow AI")	54%	54%	42%	51%	55%	68%	52%	60%
Bias or discrimination in AI-driven decisions	41%	36%	43%	42%	40%	50%	37%	43%
Fraud, scams, or deepfakes targeting the organisation	54%	54%	52%	54%	52%	58%	56%	55%
Operational disruption caused by AI failures	48%	43%	51%	45%	45%	59%	48%	44%
Lack of employee AI skills or governance	51%	45%	57%	50%	52%	52%	52%	49%
Immediate competitive reactions from rivals deploying AI faster	37%	39%	40%	44%	38%	28%	33%	37%
Other (please specify)	1%	0%	2%	1%	1%	1%	0%	0%
Not sure	0%	0%	0%	0%	0%	0%	0%	0%

National Approaches Pull Apart

The race to become a nation recognised for leading the AI revolution is a key focus for many governments, but businesses increasingly define how they view successful AI policies and implementation in very different terms. In the eyes of the business leaders we surveyed, there are diverging views on how governments can demonstrate the strength of their AI commitments, for example, sovereignty is the top priority for those we surveyed in the UK (41%), safety and skills in Spain (48%), and mandatory risk disclosure in the US (40%).

The competitive battle over AI will therefore not be about who regulates more or less, but about which combination of control, capability and confidence each country believes gives it an advantage. These are not variations on one regulatory model but three theories of what makes an AI economy trustworthy, and they are diverging quickly.

Regulatory uncertainty is already a commercial drag. It is increasing risk aversion and slowing adoption, and it sits behind a clear majority of organisations having slowed, paused or pulled back a planned deployment on reputational, regulatory or trust grounds. Businesses

are budgeting for it too: 27% of respondents expect to increase AI governance headcount, alongside further investment in crisis management and change.

The emerging frontier: resilience and sovereignty. Half of organisations feel exposed to dependency on non-domestic AI infrastructure and, while many believe they are prepared for the loss of a critical provider, relatively few have tested alternatives. Preparedness that has never been rehearsed is an assumption, not a capability.

For multinationals this is compliance arbitrage in reverse. A disclosure posture built for the US can read as overclaiming in a market policing safety and skills, while a sovereignty-led answer in the UK may look protectionist elsewhere. One global AI narrative is unlikely to survive three divergent regulatory tests: naming your dependencies and evidencing a tested alternative is the only position that travels.

Can your AI strategy survive contact with three different national tests of trust; sovereignty, safety and disclosure? How are you tracking, understanding and responding to the change in AI regulation?

“Governments want to see the benefits of AI - in improving public processes and structures, and in economic output and productivity. But they are also extremely averse to risk: regulation is an easy tool to reach for - and many voters want to see it used. There is a tension here: how that tension is managed, as nations inevitably diverge in their approach in the sector, will determine who leads in AI and who follows.”




Alex Deane, Partner Public Affairs


Q: What are your top three AI policy priorities for government?

	All markets	United Kingdom	France	Spain	Germany	Belgium	USA	Ireland
Strong AI sovereignty framework	27%	41%	21%	28%	24%	19%	26%	30%
Clear liability rules when AI fails	34%	32%	38%	35%	35%	25%	37%	36%
Faster data-centre and energy infrastructure	33%	36%	30%	31%	33%	33%	35%	35%
A distinct regulatory approach from the EU (UK respondents only)	28%	26%	31%	24%	31%	41%	16%	24%
Stronger mandatory AI risk disclosure	36%	27%	33%	32%	33%	45%	40%	42%
Greater investment in AI safety and skills	37%	36%	33%	48%	32%	34%	41%	32%
Stronger protections against deepfakes, misinformation and AI-enabled fraud	33%	35%	33%	32%	37%	24%	39%	33%
Clearer rules on data, privacy and copyright for AI training	29%	37%	24%	27%	32%	18%	36%	30%
International coordination and standards	8%	6%	9%	8%	9%	4%	9%	10%
Other (please specify)	0%	0%	0%	0%	0%	0%	0%	0%
None	2%	2%	2%	1%	2%	1%	2%	0%

Q: How exposed do you believe your organisation is to dependency on non-domestic AI infrastructure or providers (e.g. US foundation models, cloud, chip supply)? (Please select one response)



Q: If a major AI provider your organisation depends on were suddenly restricted or unavailable (e.g. regulatory action, export controls), how prepared would you be to continue operating? (Please select one response)



- Fully prepared, we have tested alternatives in place
- Well prepared, we have identified alternatives but not tested
- Partially prepared, we have some workarounds, significant disruption likely
- Poorly prepared, we would experience major disruption
- Not prepared at all
- Don't know

AI's Second Act, Where Safety, Risk and Trust Collide

Business leaders are at a critical juncture. AI's Second Act will be defined by how organisations navigate how to credibly, safely and sustainably deploy and scale AI. While the business leaders we surveyed remain optimistic about AI's value, our research finds that growing concerns around AI mean that those who align AI ambition with demonstrable outcomes and responsible governance will be best positioned to create long-term competitive advantage. Ultimately, those who stand to succeed in this AI-first world will be asking themselves the right questions and more importantly adopting strategies which help them overcome big challenges:

1. Proving Value

How do we turn AI investment into measurable business outcomes?

2. Managing Risk

How do we innovate with AI while maintaining security, control and resilience?

3. Building Trust

How do we create confidence among employees, customers and investors?

4. Governing Responsibly

How do we stay ahead of regulation, accountability and ethical expectations?

5. Protecting Reputation

How do we communicate AI ambition credibly without creating reputational risk?

Our AI Advisory Offer

We exist to turn AI from a boardroom ambition into an advantage you can prove with credibility and conviction. Our advisory spans strategy, people and process, and is built for global business leaders who need to articulate their vision for AI, while providing evidence of a considered and risk-minimised plan to deliver it and generate the best commercial outcomes.

About FTI Consulting

FTI Consulting is an independent global business advisory firm dedicated to helping organisations manage change, mitigate risk and resolve disputes: financial, legal, operational, political & regulatory, reputational and transactional. FTI Consulting professionals, located in all major business centres throughout the world, work closely with clients to anticipate, illuminate and overcome complex business challenges and opportunities. For more information, visit www.fticonsulting.com and connect with us on social media.

Our Leadership Expertise

C-suites, boards of directors, and business leaders from around the world come to FTI Consulting with their most complex, business-critical issues that require diverse skill sets and integrated disciplines. As part of our Strategic Communications division, we support dozens of senior executives and high-profile individuals with their reputation management, stakeholder engagement and communication strategies. Our offerings help them mitigate risk and enhance their reputation by combining decades of deep subject matter expertise with functional and disciplinary experience. Put simply, we help our clients to communicate effectively – across any medium – to protect and enhance their interests with key stakeholders.

Thank you for taking the time to read this report. For more information on this report, or to better understand how your organisation can move beyond the hype cycle and prove a return on AI, please get in touch.

Methodology and Scope

FTI Consulting conducted an online survey from the 5th to 11th of August 2026 with n=1,600 senior business decision-makers working in large companies across seven markets. The markets included were: United Kingdom (n=250), France (n=250), Spain (n=250), Germany (n=250), Belgium (n=250), USA (n=250), and Ireland (n=100). On average, the annual turnover of these companies was reported at £1.5 billion.

Primary Contact

Jon Priestley

Senior Managing Director, Partner
Strategic Communications
FTI Consulting

+44 (0) 20 3727 1012 T
+44 (0) 7583 033 692
jonathan.priestley@fticonsulting.com

The views expressed herein are those of the author(s) and not necessarily the views of FTI Consulting, Inc., its management, its subsidiaries, its affiliates, or its other professionals. FTI Consulting, Inc., including its subsidiaries and affiliates, is a consulting firm and is not a certified public accounting firm or a law firm.

FTI Consulting is the leading global expert firm for organisations facing crisis and transformation.

FTI Consulting is dedicated to helping organisations manage change, mitigate risk and resolve disputes: financial, legal, operational, political and regulatory, reputational and transactional. FTI Consulting professionals, located in all major business centres throughout the world, work closely with clients to anticipate, illuminate and overcome complex business challenges and opportunities.

© 2026 FTI Consulting, Inc. All rights reserved. fticonsulting.com