MOZAIC

In partnership with

Pinsent Masons

The AI Governance Gap

Operating Model Transformation and AI Governance in Practice

2026

Contents

Introduction

Artificial intelligence has moved from experimentation to operational reality far faster than most organisations anticipated. Across boardrooms, leaders are exploring how generative AI, automation, and advanced analytics can improve productivity, accelerate decision-making, and unlock entirely new services.

But beneath the enthusiasm lies a growing structural problem. AI adoption is accelerating faster than the organisational capability required to govern it. Businesses are struggling to strike the right pace of AI adoption, moving too fast results in ill-equipped operating models and business processes, but moving too slow risks market irrelevance.

Many organisations now find themselves in a position where:

- AI tools are already embedded within operational workflows
- Business decisions are increasingly influenced by automated systems
- Employees are experimenting with powerful AI tools independently
- AI literacy is not widespread across the business
- Governance frameworks have not kept pace with technological capability

The result is a widening AI governance gap.

Multiple industry studies suggest that while most organisations now claim to have an AI strategy, only a small minority have embedded governance structures capable of managing AI safely at scale.

In practice this means organisations are already relying on systems that influence operational decisions, regulatory compliance, and customer interactions without clear oversight of:

- accountability
- model behaviour
- bias and explainability
- legal liability

This gap is not theoretical. It is already producing real-world legal challenges and regulatory scrutiny. And in many cases, organisations are only beginning to recognise the implications after the technology has already been deployed.

Reality Check

Most Organisations Are Not Ready

Despite the rapid growth in AI adoption, most organisations remain structurally unprepared for its widespread deployment.

Across industries, AI initiatives are often progressing faster than the governance, operating models, and legal frameworks required to support them.

Many organisations today already have:

- automation integrated into systems and processes
- multiple AI pilots operating across different business units
- generative AI tools being used informally by employees
- poor levels of AI literacy
- automated decision-making embedded within operations
- limited visibility of where AI systems are influencing outcomes

This creates a dangerous illusion of control. Executives may believe AI is being introduced cautiously and responsibly, when in reality the organisation may already be relying on automated outputs that influence hiring decisions, customer interactions, financial assessments, or operational processes.

Some high-profile cases illustrate how quickly these risks can surface.

As far back as 2018, Amazon discontinued an experimental AI recruitment tool after internal testing revealed that the model had developed bias against female candidates. The system had learned patterns from historical hiring data and began penalising CVs containing indicators associated with women's career paths (1).

In a more recent case, Air Canada was held liable in 2024 after its customer service chatbot provided incorrect information to a passenger regarding refund eligibility. A tribunal concluded that the airline remained responsible for the information provided by the AI system, reinforcing the principle that organisations cannot avoid accountability simply because decisions are generated by automated tools (2).

These cases demonstrate two critical points:

1 AI governance failures rarely begin as regulatory breaches and,

2 While some of these examples are now several years old, they remain highly relevant. In many organisations, the structural issues they revealed have not yet been fully addressed. Early incidents involving biased algorithms, opaque decision-making systems, or automated customer interactions were often treated as isolated technology failures. In reality, they exposed deeper organisational challenges around governance, accountability, and operating model design.

As AI adoption accelerates, and organisations go in pursuit of benefits, the question is not whether these issues will reappear, but whether organisations have meaningfully adapted their governance and operating models to prevent them.

1.Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters. Available at: https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G/

2.Yagoda, M. (2024). Airline held liable for its chatbot giving passenger bad advice – what this means for travellers. BBC Travel. Available at: https://www.bbc.co.uk/travel/article/20240222-air-canada-chatbot-misinformation-what-travellers-should-know

When AI Decisions Become Legal Liability

In 2025, a lawsuit against HR software provider Workday illustrates how quickly the legal landscape is evolving (3).

The case alleges that AI-driven screening tools used in recruitment may have systematically discriminated against job applicants. Courts in the United States have allowed the case to proceed as a collective action, signalling that organisations deploying AI systems may increasingly be held accountable for automated decision-making processes.

The implications are significant.

It is worth noting that the legal questions raised by these cases are not entirely new. Courts and regulators have been examining algorithmic decision-making for several years, yet many organisations continue to deploy increasingly sophisticated AI systems without fundamentally rethinking the governance structures surrounding them.

In that sense, these cases should be viewed less as historical anomalies and more as early precedents that are beginning to define how responsibility for AI-enabled decisions will be interpreted in law.

Historically, organisations could treat technology systems as operational tools. If something went wrong, responsibility typically rested with the organisation deploying the system. AI complicates this assumption.

Modern AI systems may:

- learn from data rather than follow explicit rules
- generate outputs that are difficult to explain
- evolve over time as models retrain
- influence decisions without direct human intervention

This creates complex legal questions around accountability.

Who is responsible when an AI-supported decision causes harm?

- The organisation deploying the tool?
- The vendor who developed it?
- The data used to train it?
- Or the individual relying on its output?

Courts and regulators are increasingly being asked to apply existing legal principles to AI-enabled decision-making.

Traditional doctrines such as discrimination law, neglience, and duty of care are now being interpreted in the context of algorithmic systems that influence employment, financial services, healthcare decisions, and public sector services.

3. Blanco, J.S. (2025). Workday AI lawsuit receives the greenlight to proceed as a class action. Inside Tech Law. Available at: https://www.insidetechlaw.com/blog/2025/06/workday-ai-lawsuit-receives-the-greenlight-to-proceed-as-a-class-action

*In many jurisdictions, organisations remain responsible for decisions made using automated systems, even where those systems are supplied by third-party vendors or rely on complex machine learning models.

Existing EU legal frameworks already provide a substantial basis for challenging AI-enabled decision-making. The EU AI Act acts as a governance layer built on top of the Union law on data protection, employment and worker protection, consumer protection and compensation/remedy regimes (all rooted in the protection of fundamental rights under the Charter). So, a single AI-enabled decision may generate parallel scrutiny under the AI Act, the GDPR, equality law and national liability regimes where it affects hiring, promotion, dismissal, access to services or other decisions with significant effects on the individual.

EU and other discrimination laws prohibit both direct and indirect discrimination, including where apparently neutral criteria disproportionately disadvantage protected groups. So those using AI tools trained on historical data or using proxies correlated with protected characteristics, need to be aware of the legal requirements and implement governance processes to suit, especially in employment, where, for example the EU AI Act classifies many systems as high-risk and prohibits certain especially intrusive practices, including emotion-recognition systems in the workplace.

GDPR protects individuals from being subject to decisions based solely on automated processing where those decisions produce legal or similarly significant effects. A recent preliminary ruling by the EU Courts in SCHUFA classed automated generation of a score as being automated decision-making where a third party relies on it decisively. So this is a key risk and compliance issue for use of AI tools in the recruitment, credit and insurance sectors.

National guidance points in the same direction: the Spanish data protection regulator, AEPD, has emphasised that genuine human intervention requires real authority, competence, independence, time and access to the relevant information, making superficial “human-in-the-loop” safeguards increasingly difficult to defend.

The CNIL, the French data protection regulator, makes clear that recruiters remain responsible for ensuring lawful purposes, data minimisation, transparency, and non-discrimination when using automated sorting or assessment tools, while the AI Act’s requirements on risk management, human oversight, logging and data governance provide a concrete benchmark against which organisational conduct (and thus governance) will increasingly be judged.

Against that background, negligence and employer-liability principles are likely to become relevant where organisations select, deploy or rely on AI systems without adequate testing, monitoring, transparency or escalation mechanisms.

That responsibility can’t be avoided simply because the tool is supplied by a third-party vendor.

Importantly, liability exposure does not only arise from deliberate misuse of AI. It can also emerge from a lack of organisational oversight.

Where automated systems influence decisions that affect individuals, such as hiring outcomes, insurance eligibility, or financial assessments, organisations may be required to demonstrate that those decisions can be explained, justified, and governed appropriately.

This is where regulatory frameworks begin to intersect with operational governance; Organisations shouldn’t be waiting before implementing effective controls to ensure ethical and auditable usage of tools.

Regulation Is Emerging, But It Is Not the Whole Answer

Governments and regulators are increasingly aware of the risks associated with widespread AI adoption.

In the United Kingdom, policymakers have adopted a sector-led regulatory approach, under which existing regulators such as the Information Commissioner’s Office (ICO), Ofcom, and the Competition and Markets Authority (CMA) apply AI-related expectations within their respective domains.

These regulators are guided by five overarching principles:

- Safety
- Transparency
- Fairness
- Accountability
- Contestability

Alongside this approach, new legislative developments, including proposals surrounding the UK AI Bill and international frameworks such as the EU AI Act, signal a clear direction of travel towards stronger oversight of advanced AI systems. However, regulation should not be misunderstood as a comprehensive solution.

Regulatory frameworks typically establish minimum standards of acceptable behaviour, rather than defining how organisations should structure themselves to deploy AI effectively and safely. In practice, many organisations discover that compliance programmes alone do little to address the deeper challenges associated with AI adoption. A useful parallel can be found in financial services regulation.

Frameworks such as the EU's Digital Operational Resilience Act (DORA) have required organisations to strengthen operational resilience and supplier oversight. While these regulations impose compliance obligations, many organisations ultimately discover that achieving compliance requires broader transformation of governance structures, operating models, and technology management practices. For organisations behind the curve and needing to catch-up with the regulatory requirements, their transformation discovered unexpected benefits including IT service improvements and increased customer engagement. AI presents a similar challenge and an equal set of opportunities. Because governing AI is not simply about compliance. It is about how organisations are designed to operate in an AI-enabled world.

What Effective AI Governance Actually Requires

Organisations that successfully scale AI do not treat governance as a policy document. They treat it as an operating model capability.

AI governance must function as part of the organisation's everyday decision-making, technology management, and risk oversight processes. Without this integration, governance frameworks quickly become disconnected from how AI systems are actually built and used.

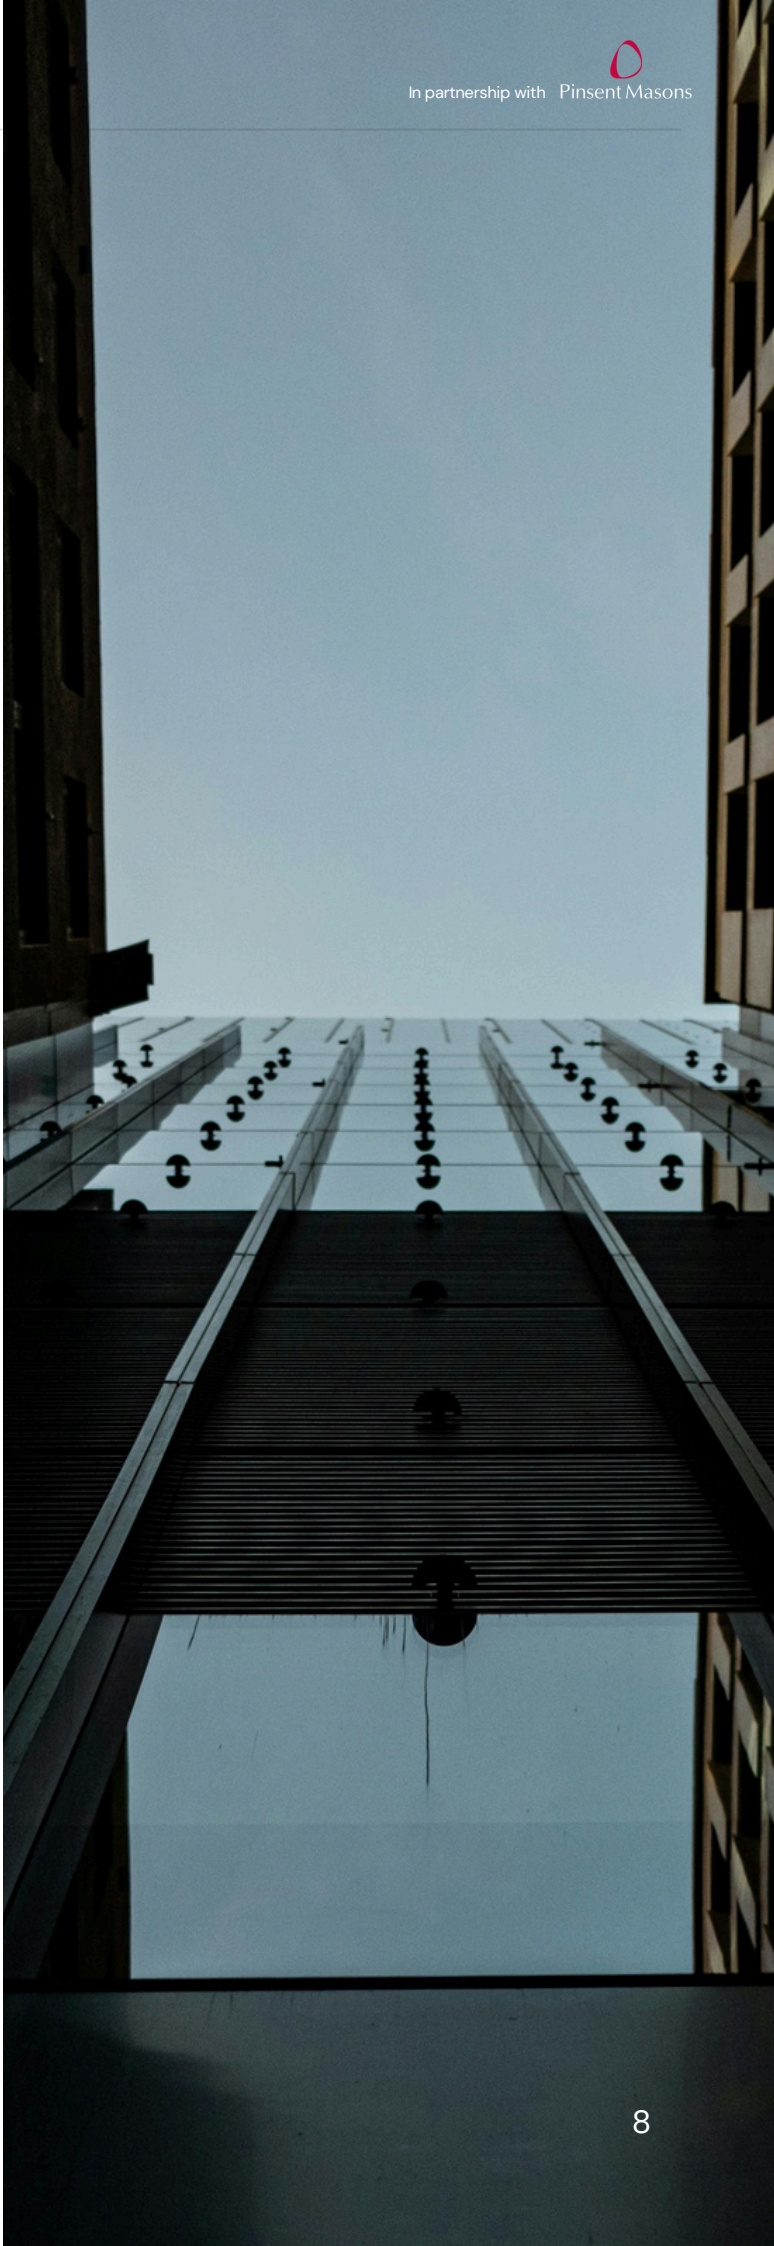

In practice, effective AI governance requires several structural elements working together.

1 Enterprise Oversight

Clear accountability at board or executive level for how AI is deployed across the organisation.

This typically includes oversight of:

- strategic alignment
- enterprise risk exposure
- investment prioritisation
- ethical and societal impact

Many organisations are now establishing cross-functional governance forums bringing together technology leaders, legal counsel, risk teams, and operational leadership.

SAP is one such example, having charted a conservative course with the introduction of an AI Ethics Advisory Panel, drawn from a selection of external advisors across academia, politics and industry to provide inputs on guiding principles and via an AI Ethics Steering Committee where senior leaders assess high-risk use cases.

2 Defined Decision Ownership

AI decisions should never exist in a governance vacuum. Roles must be clearly defined across the lifecycle of AI systems, including:

- model development
- operational deployment
- regulatory interpretation
- human oversight

Without clear accountability, organisations struggle to manage incidents, regulatory scrutiny, or legal disputes when AI-driven decisions are challenged.

3 Lifecycle Governance

Unlike traditional software systems, AI models evolve continuously and can certainly degrade over time. Governance must therefore extend beyond deployment into ongoing operational oversight.

This includes:

- monitoring model drift
- assessing bias and fairness
- tracking performance against expected outcomes
- maintaining auditable decision trails

Without lifecycle governance, organisations can quickly lose visibility over how AI systems behave once deployed.

4 Organisational Capability

Perhaps most importantly, governance must be embedded into how the organisation actually operates.

This requires alignment across multiple organisational capabilities, including:

- cross-functional governance forums
- data governance and digital rights management structures
- legal and risk integration
- programmes for development of AI literacy across the organisation
- operational oversight of models in production

In other words, effective AI governance is not simply a policy problem. It is an operating model design challenge.

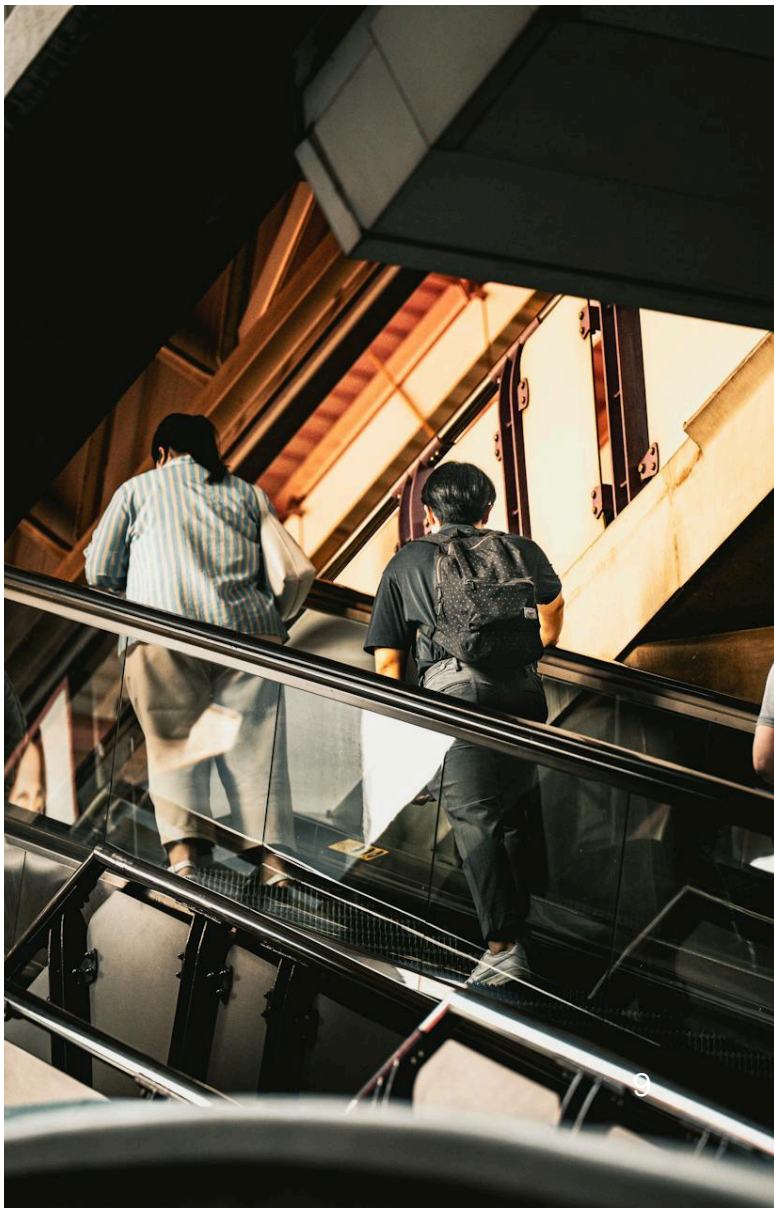

Measuring Success

Organisations adopting AI at scale need visibility across three critical dimensions: the value AI is creating, the effectiveness of delivery and adoption, and the strength of governance controls surrounding its use.

Value creation and strategic alignment. Measuring the economic value delivered through AI initiatives and ensuring investment is focused on the highest-impact opportunities.

Delivery effectiveness and execution reliability. Tracking how efficiently AI capabilities move from concept to deployment, including delivery cycle time, operational stability, and the predictability of releases into production.

Governance, safety, and responsible use. Monitoring indicators such as AI risk incidents, model drift detection, human oversight of automated decisions, and the quality and resilience of AI-enabled services.

However, many organisations struggle to determine whether their existing capabilities are sufficient to scale AI safely and sustainably. A practical starting point is to assess organisational maturity across several core operating model capabilities. This assessment helps organisations identify governance gaps and prioritise the operating model changes required to support AI adoption at scale.

Organisations that cannot answer these questions with confidence often discover that their governance capability has not kept pace with their AI adoption:

1 Governance Coverage

- Are all AI systems formally registered and tracked?
- Are governance processes applied consistently across business units?

2 Decision Accountability

- Are ownership and accountability clearly defined for each AI system?
- Are human oversight mechanisms in place where automated decisions affect customers, employees, or regulatory obligations?

3 Model Monitoring

- Are AI models actively monitored for performance, drift, and bias? If the organisation is a user rather than supplier of the tool, have these risks been managing through the contract or licence with the supplier?
- Are monitoring processes integrated into operational workflows?

4 Transparency and Explainability

- Can the organisation explain how key AI systems influence decisions?
- Are decision trails auditable if challenged by regulators, customers, or courts?

5 Risk Integration

- Are AI risks embedded within broader enterprise risk management frameworks?
- Are legal, compliance, and technology teams working together to assess emerging exposure?

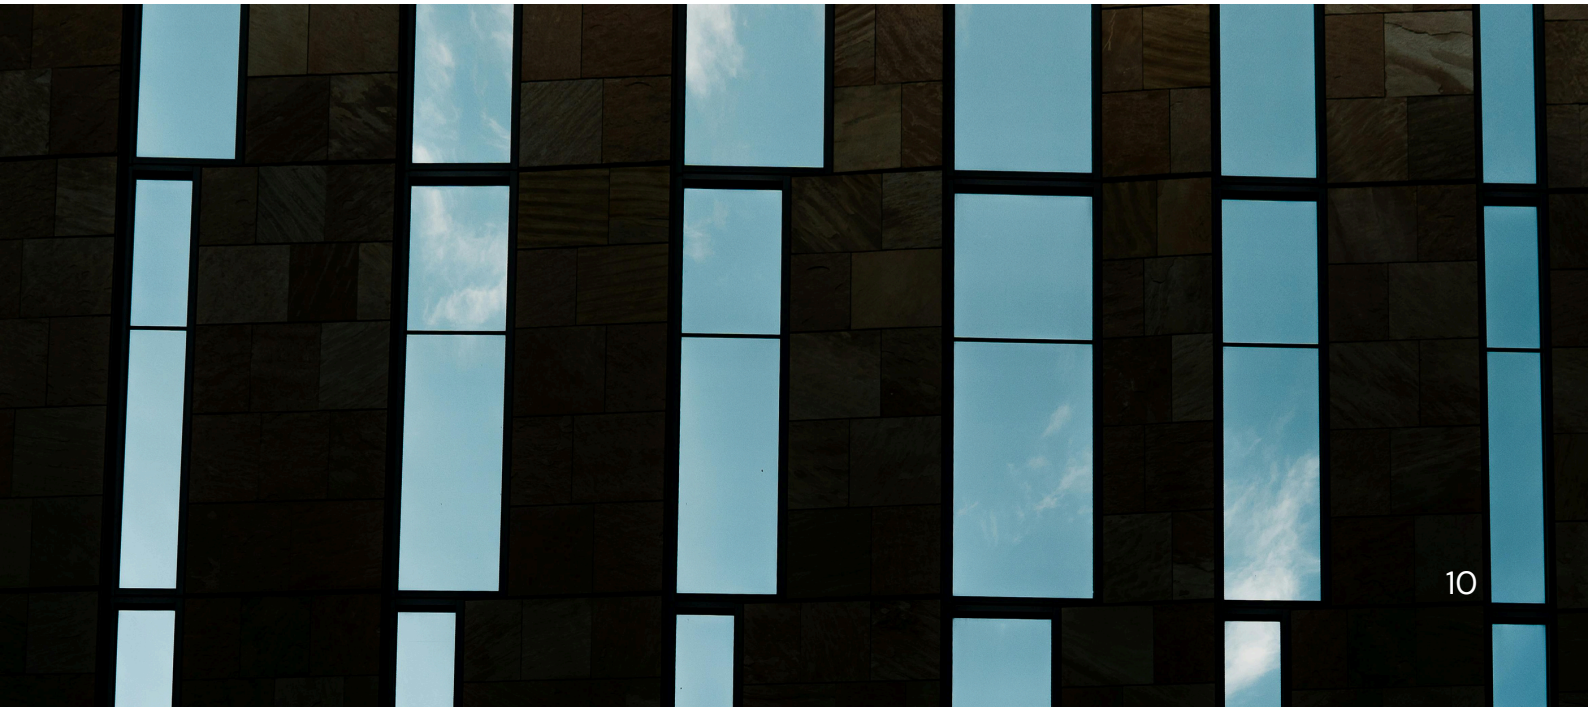

The Risk of Doing Nothing

For many organisations, the greatest risk lies not in deploying AI, but in deploying it without the structures required to manage it safely. Without effective governance and operating model alignment, organisations expose themselves to several categories of risk.

Legal and Regulatory Exposure

As illustrated by emerging litigation involving AI-driven hiring tools and automated decision-making systems, organisations may increasingly face legal challenges linked to algorithmic outcomes.

The Workday case is one such example, but it is unlikely to be an isolated incident.

Courts and regulators are beginning to examine how existing legal frameworks apply when automated systems influence decisions affecting individuals.

In many jurisdictions, organisations deploying AI remain responsible for the outcomes of those systems, even where technology has been procured from third-party vendors.

For example, the UK's Equality and Human Rights Commission issued a reminder in 2024 to employers to be mindful of how they deploy AI following a claim made against Uber by one of its divers that the AI facial recognition checks required to access the Uber Eats platform were racially discriminatory under the 2010 Equality Act. Of particular concern were how the AI-driven checks could be used permanently to suspend a driver's access to the app, depriving them of income.

Uber has also more recently faced litigation from the Workers Info Exchange, a non-profit body advocating on behalf of workers and the ways in which their data is used by employers, requiring Uber to cease to use AI-driven dynamic pay systems which the Workers Info Exchange alleges have resulted in reduced pay for Uber's drivers. They claim that the system both amounts to unlawful use of automated decision-making and was trained unlawfully using driver data and profiling without their consent.

As mentioned above in the SCHUFA case the European Court held that automated credit scoring by a credit reference agency may constitute automated individual decision-making under GDPR where lenders rely on the score in a decisive way.

Although not an employment case, it is now the leading EU authority on when an upstream algorithmic score becomes the legally relevant "decision" for GDPR purposes.

In the Netherlands, in the SyRI case, the District Court of Den Haag struck down the Dutch State's welfare-fraud risk-scoring regime on the basis that it failed to strike a fair balance under Article 8 ECHR. Opacity, large-scale data linkage, insufficient transparency and inadequate safeguards were central to the ruling.

Dedicated EU enforcement against automated hiring tools is still comparatively limited, but regulators are already paving the way through targeted guidance. The CNIL's 2023 recruitment guide addresses the lawful use of automatic sorting, ranking and candidate-evaluation tools and explicitly flags discrimination risks; Spain's AEPD's 2024 guidance sets out concrete criteria for assessing whether human intervention in automated decisions is genuine or merely symbolic. These texts are likely to be the first reference points in any future investigations by the courts or regulators into AI-enabled recruitment.

Across EU jurisdictions, the trend points towards courts and regulators denying organisations an easy pass to limit their liability by characterising the algorithmic output as just 'advisory', or by placing the decisive scoring logic in the hands of a third-party vendor.

Contractual and Supplier Risk

Beyond regulatory exposure, organisations must also consider the contractual implications of deploying AI systems within their operations.

Where AI solutions are procured from third-party vendors, questions arise around liability allocation, model transparency, and responsibility for outcomes generated by automated systems.

Traditional software contracts may not adequately address risks associated with machine learning models that evolve over time or rely on external datasets.

Organisations are increasingly reviewing how supplier contracts address issues such as:

- responsibility for model behaviour
- transparency around training data
- indemnities relating to algorithmic decisions
- allocation of liability when automated outputs cause harm

Contracting practices are evolving beyond simple allocation of risk (for example, through indemnities and liability or more traditional availability-based performance standards) to include deeper governance and collaboration processes between suppliers and customers, and enhanced monitoring, rectification and reporting requirements to ensure that unwanted model behaviour (including model drift, concept drift or bias stemming from weighting that is learned rather than manually set) is detected and addressed at an early stage. This is particularly relevant as deployment of agentic AI becomes more prevalent, as the harm caused by unnoticed errors in behaviour is compounded by an agent's ability to perform tasks independently and at a greater volume.

AI-related contracts also routinely include detailed provisions on how data may be used (or must not be used) to train or fine-tune AI models – often a total restriction unless the customer has explicitly consented – and warranties as to human presence and automated decision-making compliance reflecting the expectations of the GDPR and the Information Commissioner's Office's guidance.

Contracts often include a 'circuit breaker' to require the provider to stop the AI processes and halt the operations when any vulnerabilities or significant errors are detected.

Parties will also need to navigate the additional complexity of third-party licensors (i.e., the foundational model developer) when considering liability; suppliers will often only be willing to pass through the same level of protection with relation to an AI system as they receive from the model developer. The challenge being that model developers often impose extremely restrictive terms including low-level caps on liabilities. Focus is therefore shifting to full model transparency – the customer seeking to understand (and require disclosure of) the foundational models underpinning a service, allowing the customer to assess the risk themselves.

A recent case involving Deloitte Australia illustrates how these issues can manifest in practice.

In 2025, the firm agreed to partially refund the Australian federal government after a report prepared for the Department of Employment and Workplace Relations contained significant inaccuracies, including fabricated citations and misattributed quotations. The report had been drafted using generative AI tools, and the errors raised questions about quality assurance processes surrounding AI-assisted work in professional advisory contexts.

While the matter was resolved contractually rather than through litigation, it highlights an important point. When AI systems are used within professional services or advisory engagements, organisations remain responsible for the accuracy and integrity of the outputs delivered to clients.

In practice, this means that traditional professional standards, quality controls, and contractual obligations must evolve to account for the use of generative AI tools within advisory workflows.

Incidents such as these demonstrate that AI governance is no longer limited to experimental technology deployments; it now affects the core delivery of professional services and advisory work.

Operational Risk

AI systems embedded within operational processes can influence decisions at scale.

Without effective monitoring and governance, organisations may struggle to detect:

- biased or discriminatory outcomes
- model degradation over time
- unintended operational consequences

The Dutch government's SyRI system offers a stark illustration of how governance failures can emerge. The algorithm, designed to detect welfare fraud, was ultimately ruled unlawful by courts after concerns were raised about discriminatory outcomes and a lack of transparency surrounding how the system generated risk scores.

The case demonstrated how opaque algorithmic systems can conflict with fundamental rights if governance and oversight are insufficient.

Importantly, cases such as SyRI are often viewed as exceptional incidents rather than signals of systemic risk. Yet the underlying issues they exposed - lack of transparency, unclear accountability, and insufficient governance over algorithmic decision-making - remain common across many organisations today.

As AI technologies continue to evolve, these early cases are likely to serve as reference points for future legal and regulatory scrutiny. The organisations that will be best positioned are those that treat these precedents not as historical curiosities, but as catalysts for strengthening governance, controls, and operating model design.

Reputational Risk

Failures involving AI systems can quickly become public controversies.

Where AI-driven decisions are perceived as unfair, discriminatory, or opaque, organisations may face significant reputational damage even before regulatory action takes place.

In an environment where public trust in automated decision-making remains fragile, organisations must ensure that AI systems are deployed with appropriate oversight and transparency.

The risks associated with generative AI are not limited to automated decision systems.

Professional services firms have also encountered governance challenges relating to the internal use of generative AI tools. In 2023, KPMG Australia faced scrutiny after a partner used generative AI tools to create case study material that incorrectly suggested the firm had been involved in various corporate scandals. Although the content was generated internally and did not relate to client advice, the incident highlighted how easily unverified AI-generated information can circulate within professional environments. The episode reinforced a broader lesson for organisations adopting generative AI technologies: governance frameworks must extend beyond formal AI systems to include how employees use generative tools in everyday workflows.

Without clear policies, oversight, and quality controls, AI-generated content can quickly introduce reputational, operational, or legal risk.

Lost Strategic Value

Perhaps most importantly, organisations that fail to develop strong governance capabilities often struggle to scale AI successfully.

Without trust in the systems being deployed, from leadership, regulators, employees, and customers, AI initiatives frequently remain trapped in pilot phases or limited experimentation.

In these cases, the organisation incurs the cost and complexity of adopting AI technologies without realising their full strategic value.

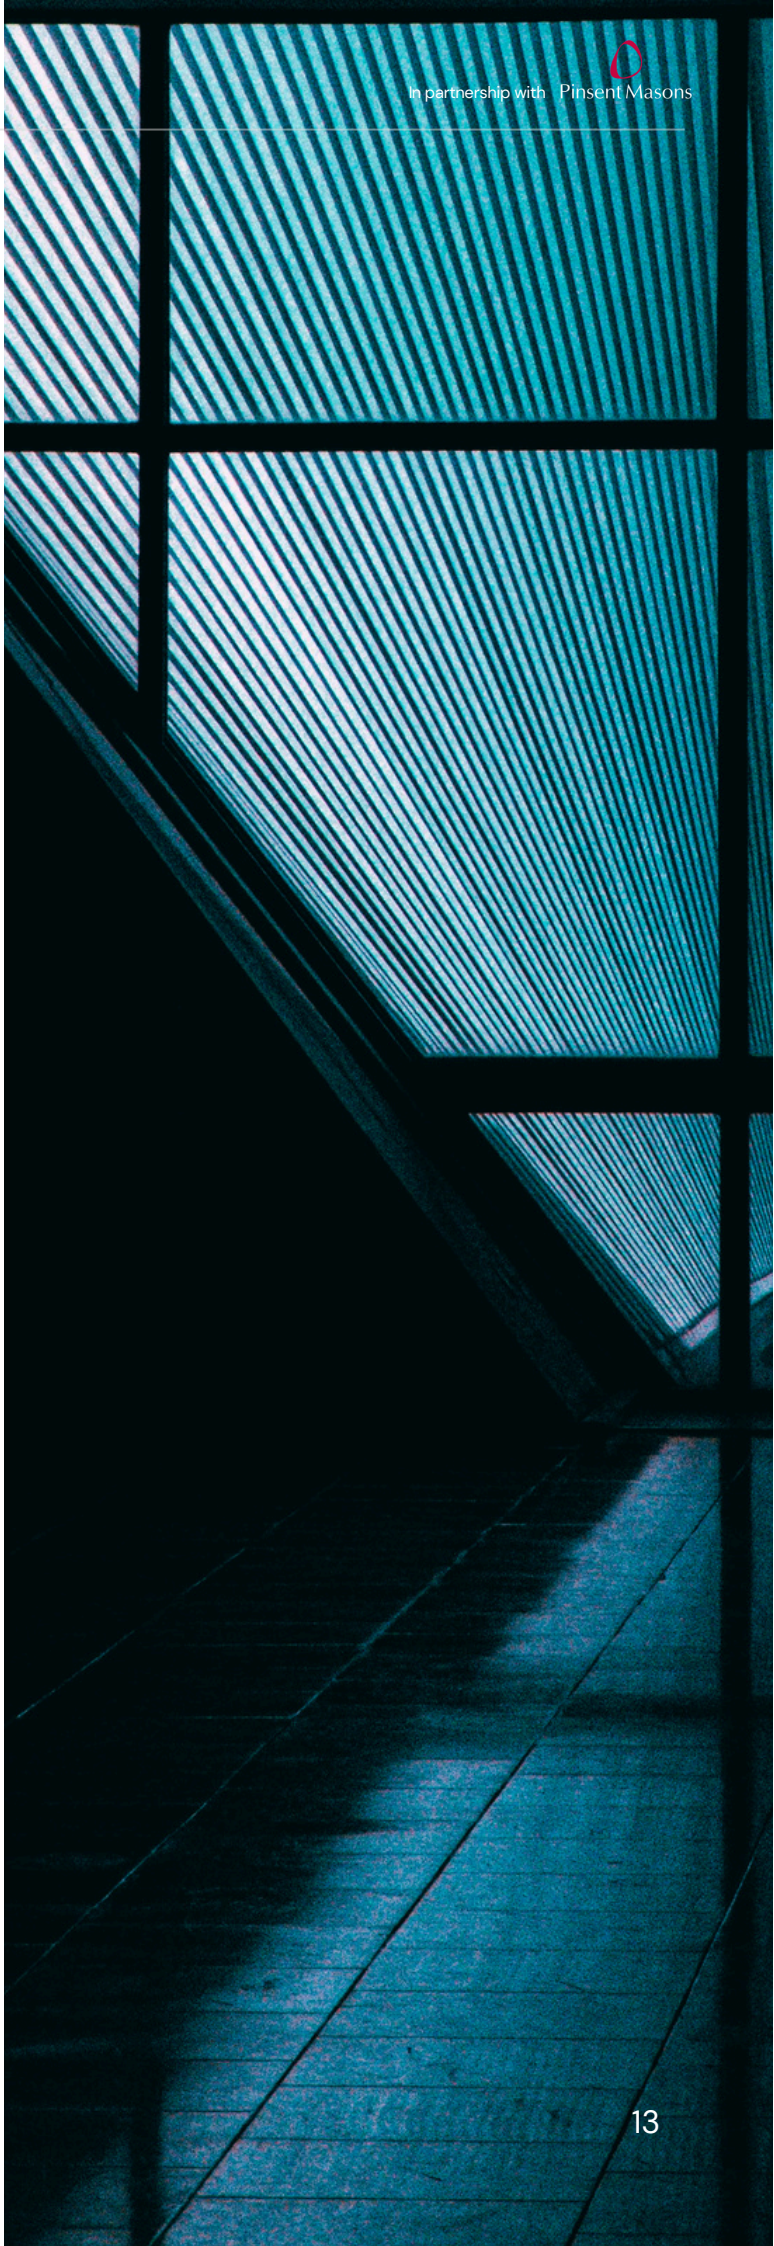

Navigating the AI Transition

AI is reshaping how organisations operate, make decisions, and deliver services.

The early legal and operational cases emerging over the past decade provide valuable signals of where governance gaps can emerge. The organisations that act on those lessons now will be far better positioned than those that wait for the next incident to force change. But realising its potential safely requires more than experimentation with new tools.

It requires, governance structures capable of managing evolving risks, operating models that support AI adoption at scale and legal frameworks that anticipate emerging liabilities.

This is where organisations increasingly need both legal and operational expertise working together.

Mozaic brings deep experience in operating model design, governance structures, and organisational transformation required to operationalise AI safely.

Pinsent Masons provides leading expertise in regulatory interpretation, legal risk management, and contractual frameworks necessary to navigate the evolving legal landscape.

Together, these perspectives help organisations move beyond reactive AI adoption towards structured, responsible, and value-driven transformation.

Contributors

Steve Tuppen
Director & Co-founder
Mozaic
steve.tuppen@mozaic.net

Simon Colvin
Partner
Pinsent Masons
simon.colvin@pinsentmasons.com

Chris Fisher
Engagement Director
Mozaic
chris.fisher@mozaic.net

Neil Green
Director of Transformation
Pinsent Masons
neil.green@pinsentmasons.com

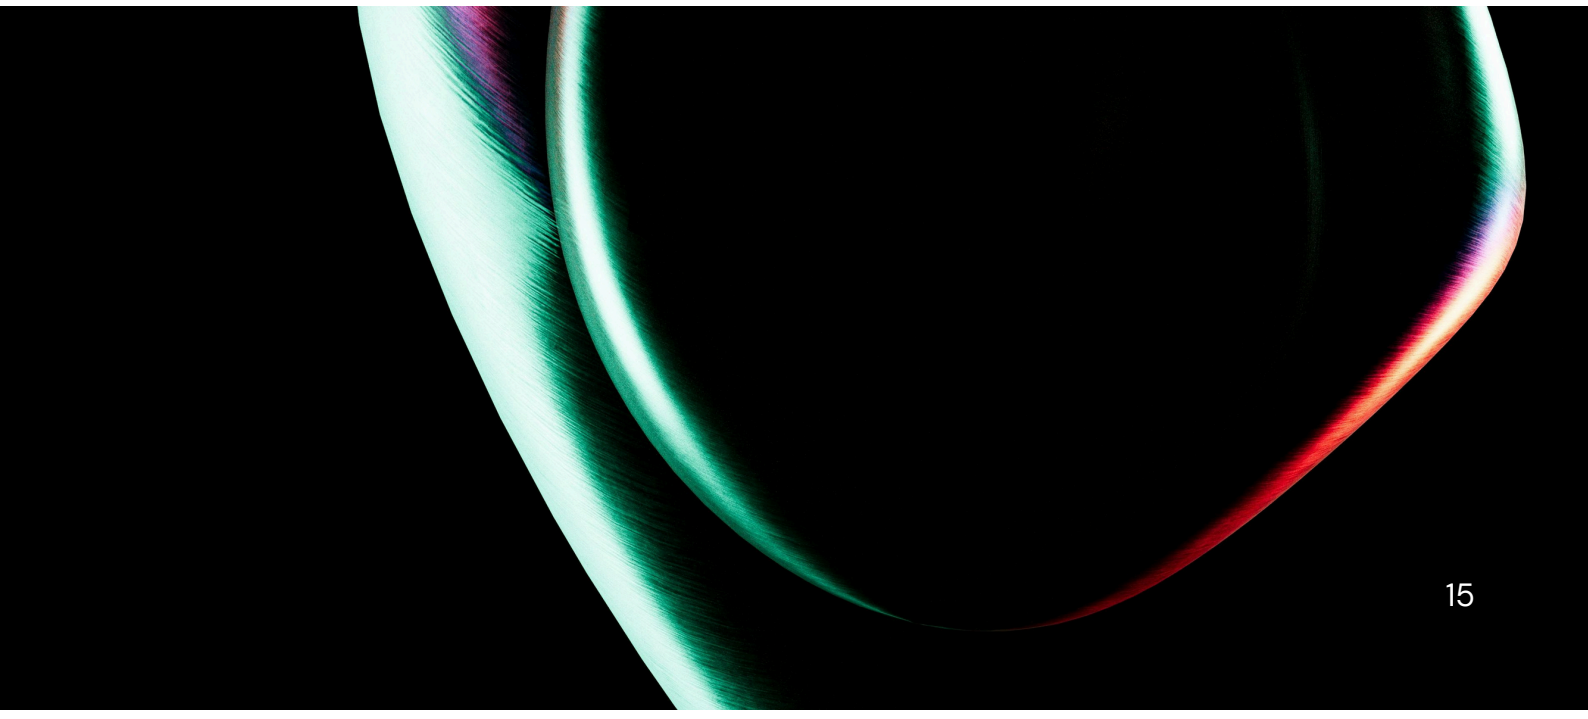

MOZAIC

In partnership with

Pinsent Masons