

Opportunities and risks



This issue covers:

- how AI can and cannot support regulatory work
- considerations for using AI safely

Read this alongside other issues in the bite-sized AI guidance series

AI matters for regulation because it can help regulators act earlier, target effort where it matters most, and make better use of limited resources. If used well, AI can support regulatory practice by helping regulators to:

- detect patterns of non-compliance and emerging risk earlier
- prioritise monitoring, inspections, and interventions based on risk
- support decision-making with timely, evidence-based insights
- identify patterns and anomalies across large datasets that may not be visible through manual analysis
- reduce administrative burden for both regulators and regulated parties.

AI may change the balance of some roles over time, potentially reducing time spent on certain manual tasks while increasing the importance of oversight, assessment, and interpretation.

AI can change workflows, roles, and how decisions get made. Without genuine organisational buy-in, even well-designed AI systems can fail to deliver value, or create new problems rather than solving existing ones ([see the Organisational integration is key issue](#)).

If your regulatory foundations are strong, AI can help amplify good systems and practices by reducing administrative effort, improving targeting, and supporting faster, better-informed decisions. If those foundations are weak, AI will not fix them. It is more likely to expose and amplify existing problems.

By "foundations", we mean the basics of good regulatory practice: clear delegations, sound record-keeping, reliable information management, consistent decision reasoning, and effective oversight. Where these are unclear or inconsistent, introducing AI increases risk. Decisions become harder to explain, errors easier to scale, and accountability harder to trace. AI is also highly dependent on the quality of underlying data. If inspection, licensing, or compliance datasets are incomplete, inconsistent, or poorly structured, AI systems will tend to reinforce those weaknesses rather than produce reliable insights.

Worker capabilities

AI can help by scanning large volumes of information, identifying patterns, or highlighting risks. It does this by recognising statistical patterns in the data

it was trained on, not by understanding context, applying discretion, or weighing legal and ethical considerations. That distinction matters. An AI system can produce an output that looks authoritative and well-reasoned while being based on patterns that do not apply to the situation at hand. Practitioners who do not understand how AI produces its outputs are poorly placed to recognise when this is happening.

The more people rely on AI to do their thinking, the less they exercise their own judgement. Research in cognitive science, including the 2025 MIT study *Your Brain on ChatGPT*, suggests that offloading reasoning to intelligent systems over time can erode the analytical capability, recall, and sound judgement that good decision-making depends on. The tools get better while the people using them become less confident in their own assessments.

This matters acutely in a regulatory context. Regulation is not mechanical. It requires weighing evidence, applying legal judgement, exercising discretion, and being accountable for decisions that affect people's rights and obligations. These are capabilities that need to be actively maintained.

When AI outputs are accepted without adequate challenge, institutional knowledge quietly erodes, domain expertise gets substituted by outputs that sound confident but may be incomplete or wrong, and human reasoning becomes reactive rather than deliberate. This is sometimes called automation bias, the tendency to defer to machine outputs without applying independent judgement.

It is a well-documented risk in any setting where people work alongside intelligent systems. For regulators, where public trust depends on confidence that decisions are made by people who understand what they are deciding and why, automation bias carries real consequences for organisational credibility and social licence to operate.

AI should be making regulatory thinking sharper, not substituting for it. Where regulators cannot explain how an AI output informed a decision, or why it was accepted or rejected, that is a signal that the balance has shifted in the wrong direction.

Research on AI adoption consistently shows that realising value from generative AI requires significant investment in people, culture, and workflow redesign. MIT's Project NANDA report, *The GenAI Divide: State of AI in Business 2025*, suggests that despite significant global investment, 95% of organisations studied saw no measurable return from AI initiatives. The core issue was not the technology itself, but the absence of deliberate integration, workflow redesign, and organisational capability to support it **(see the *Empower your people issue*).**

Ethical use

Ethical use of AI requires ongoing review and adaptation as systems, risks, and contexts change.

When AI is used in regulatory contexts, fairness and privacy are not separate concerns. Both go to the heart of protecting people from unintended harm, especially where decisions affect rights,

access to services, or regulatory outcomes [\(see the AI in Regulation issue\)](#).

Bias can enter AI systems in many ways: through training data, model design, feedback loops, or how outputs are interpreted by people. Regulators need to be alert to these risks and treat fairness as an ongoing responsibility.

Privacy risks also increase when AI draws on large or sensitive datasets. Regulators must apply strong data governance and ensure personal information is handled lawfully, securely, and proportionately.

As a practical safeguard, early AI use should favour public or low-risk data sources and tasks where errors can be detected and corrected easily [\(see the Grow as you learn issue\)](#). Where AI-supported processes affect the public, flexibility is important so people are not disadvantaged by differences in digital access, literacy, or circumstance.

Regulators operate under public scrutiny. A structured approach to AI adoption supports confidence by demonstrating that decisions remain accountable, reviewable, and aligned with the public interest.

Te Tiriti o Waitangi in AI-enabled regulatory practice

Regulators have existing responsibilities under Te Tiriti o Waitangi to recognise

Māori rights and interests, support equitable outcomes, and maintain trust and confidence.

The use of AI in regulatory practice does not change these responsibilities. It can, however, amplify impacts. Decisions supported by AI may affect how Māori experience regulation, including access to services, compliance outcomes, and regulatory oversight. For regulatory leaders, this means considering Te Tiriti implications early when exploring and adopting AI, in a way that is proportionate to the impact of the use case.

Where Māori data is involved, it should be treated as taonga. Māori data refers to data about Māori people, their communities, land, culture, and resources. Māori data sovereignty holds that Māori should have authority over how that data is collected and used, consistent with Te Tiriti o Waitangi [\(see the Further reading issue](#) for resources on Māori data and Māori data sovereignty).

When considering AI use in regulatory settings, regulators should identify and engage with the Māori individuals, groups, or organisations who are affected by, or have mana in relation to, the regulatory activity. The appropriate stakeholders will vary depending on the nature of the regulatory function and the type of impact.

To get more practical steps for how regulatory leaders can lead AI innovation with confidence, check out the full guidance: [Responsible AI in Action](#).