MAXIM INSTITUTE

DISCUSSION PAPER

ALGORITHMIC GOVERNMENT: OF THE PEOPLE, FOR THE PEOPLE, BY AI?

DR PAUL HENDERSON
THOMAS SCRIMGEOUR

First published in May 2026 by Maxim Institute
PO Box 49 074, Roskill South, Auckland 1445, New Zealand
Ph +64 9 627 3261 | Fax +64 9 627 3264 | www.maxim.org.nz

978-1-0670414-8-9 | PDF
978-1-0670414-9-6 | Hard copy

Maxim Institute is an independent think tank, working to promote the dignity of every person in New Zealand, by standing for freedom, justice, compassion, and hope.

At the core of our work is an understanding that ideas will shape society over time. So, we're engaged in long-term research and analysis of the policies and ideas that inform our nation's social, political, and cultural practices.

To find out more, head to www.maxim.org.nz.

MAXIM INSTITUTE

ALGORITHMIC GOVERNMENT: OF THE PEOPLE, FOR THE PEOPLE, BY AI?

DR PAUL HENDERSON AND THOMAS SCRIMGEOUR

This is Maxim Institute's sixth discussion paper on artificial intelligence (AI). Its focus is on the use of AI by government(s) and the controls around it. Moreover, by implication, and in a slightly futuristic vein, it asks how much human participation in governing by public servants and MPs is needed, and if this will decrease in the near future as machine intelligence matches human capabilities. Will government be efficiently managed and executed by AI, and would such a scenario benefit New Zealanders? And further, are we moving inexorably to a post-democratic society characterised by fewer and fewer elected representatives, a small public service, and a more informed, able, and better run government?

About the authors

Dr Paul Henderson, Senior Research Fellow. Paul's background is in education. Working from DPMC, under John Key, he was secretary to a cross-party committee tasked to review school choice and educational reform. He holds degrees from Aberdeen, Cambridge, and Laidlaw College, and a PhD from Australian Catholic University. His research focuses on human and artificial intelligence.

Thomas Scrimgeour, Researcher. Thomas is a public policy researcher at Maxim with a background in the trades and community sector. His research interests include social cohesion, New Zealand nationhood, and how communities adapt to a rapidly changing world.

Table of contents

The paper in summary...

Artificial intelligence (AI) in the public sector is shifting from isolated experiments to core administrative infrastructure. By late 2025, governments in comparable jurisdictions—including the UK, Canada, and Australia—have steadily moved toward system-level AI implementation, AI-integrated public registries, and workforce-wide AI adoption to address fiscal constraints and rapidly rising information volumes.

This report analyses the transition through three normative frames:

1. DEI-oriented governance: focused on redress and participation.
2. Competitive national utility: focused on productivity, responsiveness, and state capability.
3. Conservative readings: focused on ordered liberty,* institutional continuity, and rule-of-law accountability.

Core argument

Aotearoa New Zealand should pursue a strategy of competitive national utility disciplined by conservative governance. Constitutional norms, understood in the broadest sense, must be taken seriously. In practice, this means that the Treaty of Waitangi operates as a standing constitutional constraint. Such an approach prioritises high-yield administrative automation (document handling, triage, etc.) to increase state efficiency, while imposing institutional guardrails to prevent post-democratic drift—a scenario where substantive decision-making authority quietly migrates from humans to opaque AI systems.

Key findings

- Value pool: the immediate and near-term utility of AI lies in internal operations and service delivery, not in autonomous policymaking or governing.
- Infrastructure shift: Western nations are rapidly establishing AI-enabled systems and frameworks (architectures) to integrate AI into public administration, service delivery, and decision-making.
- Risk: without prescribed governance, reliance on vendor-controlled models of AI integration creates strategic fragility and sovereignty risks.

* The principle that individual freedom must be balanced with, and is sustainable only within, the framework of social order, law, and the common good—rather than an unrestrained licence that risks descending into chaos.

Recommendations

The report recommends distinct operational mechanisms to implement AI strategy:

- Bounded autonomy: implement "cognitive forcing functions" in decision systems, requiring AI-powered tools used for important choices (for instance, in government services or policies) to include built-in features that make users pause and think carefully before accepting the AI's suggestion. Humans are to record their independent judgement before viewing an AI recommendation to prevent automation bias. AI reasoning is to be transparent for purposes of review to ensure ongoing AI safety.
- Procurement as governance: mandate "model portability" (the ability to transfer, export, or migrate an AI model—including its weights, architecture, and associated data—to alternative providers, platforms, or in-house systems without prohibitive technical, legal, or financial barriers) and "inspection rights" in contracts to prevent vendor lock-in and ensure that public servants and those governing can independently validate "black box" systems.
- Workforce design: collaborate with aligned countries and explore opportunities to second critical AI-capability staff in international contexts. Broad AI literacy development should shift focus from simple tool training to environmental design that supports critical vigilance—staying alert and thoughtfully checking AI outputs—and independence, encouraging employees to make their own judgement, keep final authority, and not become passive overseers of automated systems.
- Treaty operationalisation: institute obligatory "Māori data audits" for all systems to determine if they affect Māori rights or interests, triggering specific partnership requirements.

1. INTRODUCTION

1.1 State of play

Artificial intelligence (AI) is no longer a peripheral experiment in government; it is becoming part of the ordinary machinery of administration and policy in Europe, North America, the Middle East, and Asia. Public-sector AI already covers a wide range of functions—"predictive policing, cyber-defence, adjudication of government benefits and chatbots for answering citizen questions," as well as tools for "filling out, searching, and drafting documents."[1] Mirroring the uptake of generative AI by the public and businesses,[2] governments have moved from isolated pilots toward issuing guidance and experimenting with generative AI (GenAI) in routine public administration.[3]

A useful high-level account of why governments are doing this appears in the OECD's public governance work. The OECD states plainly that "the use of AI by the public sector can increase productivity, responsiveness of public services, and strengthen the accountability of governments."[4] Such is the core promise that drives a "competitive national utility" policy approach: AI is treated as "a capability multiplier that allows the state to do more with limited people and budgets, especially where workloads are document-heavy and repetitive."[5]

Fresh empirical mapping reinforces the findings that the majority of current AI deployments sit in internal operations and service delivery rather than direct policymaking. A recent EU analysis of 250 cases found AI was "used mainly to support improving public service delivery," followed by "enhancing internal management," with relatively few cases assisting policy decision-making.[6] This pattern matters because it indicates where the near-term value pool is likely to be found (process throughput and service responsiveness) and where governance failures will probably be felt first (administrative pipelines that affect large numbers of people), with New Zealanders and their neighbours across the ditch haunted by memories of institutional IT failures from the past.[7]

At the same time, and of more significance, public-sector AI is evolving from scattered projects into system-level *"architectures."* The General Services Administration (GSA) AI Guide—written as a practical resource for US federal leaders—expresses the operational motivation directly: agencies face "increasing complexities… and increasingly large volumes of data," and are "*past the point* where human cognitive abilities can directly process and make sense of all this information."[8] Rather disconcertingly, it concludes: "The federal government needs AI," because "with thousands or millions of pages of documents, we could never even try to hire enough staff to read through them all."[9]

The same logic has increasingly extended beyond the US into a broader trend toward "AI-enabled states," where digital public infrastructure, data-for-policy, algorithmic decision processes, and GovTech service layers are now interacting as an integrated system.[10] These developments do not occur in a vacuum. AI in government is being introduced amid long-running institutional trajectories—bureaucratic rationalisation, digital government reform, and managerial performance pressure—and in an international environment increasingly described in competitive and survivalist terms.

The 2025 bi-weekly updates Maxim Institute has run highlight AI adoption, and in some instances sharpen this competitive picture. A White House executive order has sought to pre-empt state AI regulations as "obstruction" to a unified competitive national approach.[11] In the UK, government has framed new partnerships with frontier model developers (OpenAI, Anthropic, Google) as part of national renewal and growth, including the planned application of "Gemini for Government" to reduce bureaucracy and automate routine tasks.[12] Canada, emphasising AI's utility, has placed public transparency and planning at the centre of its approach through a public AI register covering 400 uses of AI across federal institutions.[13] Australia's public service is now openly curating practical examples via its GovAI use case library and highlighting staff productivity and consistent service response as core benefits.[14] Singapore, for its part, is pairing whole-of-government capability-building with formal risk management instruments such as the Monetary Authority of Singapore's consultation on AI risk guidelines for the financial sector.[15] Aotearoa New Zealand's own digital infrastructure choices are moving in the same direction: the Govt.nz app is being positioned as a unified "front door" for government services and will include a digital wallet for credentials, with further features planned.[16]

Table 1. Summary of late-2025 and early 2026 developments in government AI adoption and governance

Jurisdiction	Development (late-2025 and following)	Policy position and relevance
United States	Federal move to pre-empt state AI regulation	Illustrates competitive national utility framing: centralisation justified as necessary to "win" and avoid regulatory fragmentation (White House 2025)
United Kingdom	Partnership with Google DeepMind ("Gemini for Government")	Illustrates a government–frontier-model compact, linking AI adoption to state modernisation and economic growth (UK Government 2025)
Canada	Public AI Register	Provides a concrete transparency and planning instrument; useful comparator for NZ inventories and registries (Treasury Board of Canada Secretariat 2025)
Australia	GovAI Use Case Library and Australian Public Service (APS)-wide AI capability uplift	Shows how a Westminster system is publishing practical use cases to move beyond pilots to normalise safe use (APS Academy 2025)
Singapore	Monetary Authority of Singapore (MAS) consultation on AI risk guidelines (financial sector)	Shows sectoral supervision for AI risk management, including lifecycle controls and inventories (MAS 2025)
Aotearoa New Zealand	Govt.nz app: digital wallet (credentials) and future capability layer	Illustrates the digital public infrastructure layer on which citizen-facing agentic AI can be deployed (Govt.nz 2025)

1.2 Problem statement

Discussion about AI in government is fragmented. "AI governance" is frequently treated as a single technical problem, when it is in fact a bundle of contested choices about prior intellectual and philosophical operating frameworks, values, institutions, accountability, and power. Even the object of AI governance is unclear: is it governing generative AI modelling (predictive, ranking, classification, generative, etc.), data, organisational practices, procurement pathways, or the relation and distribution of authority between humans and systems?

At the same time, governments are both regulators and users of AI.[17] That dual role creates an additional layer of complexity: the state must set rules for others while also governing its own adoption, including its own procurement choices and risk tolerances.

More fundamentally, the use of AI in and by government is normatively contested: there is no broad consensus on whether, how, or to what extent governments should integrate AI into their core functions. Thus, this report works within three distinct frames that commonly underwrite policy and practice:

- DEI-oriented governance: AI as a tool for redress, group-level fairness constraints, and participatory legitimacy.
- Competitive national utility: AI as an instrument for capability, productivity, resilience, and strategic advantage.
- Conservative readings: AI assessed by its compatibility with ordered liberty, institutional continuity, and rule-of-law accountability.

The report also addresses a further issue in a forward-looking and arresting scenario: whether incremental expansions of AI autonomy could result in a post-democratic settlement, where government is formally democratic but substantively shaped by algorithmic systems that humans supervise rather than truly control.

In Aotearoa New Zealand, these frames intersect with Te Tiriti o Waitangi as a core constitutional context. This introduces Treaty obligations—often expressed as partnership, participation, and protection—and standing constraints on what counts as legitimate data use, model deployment, and automation in domains where Māori rights and interests are engaged. The report does not treat

these standing constraints as a subset of diversity, equity, and inclusion (DEI) or as an optional add-on. This report recognises that Treaty obligations are highly contested and not the only constitutional constraint on AI deployment. However, they are salient and illustrate a fundamental challenge presented by emerging technologies: in many respects, social institutions and norms are downstream from material, technological developments.

In short, rapidly advancing technology raises existential questions for many of the political and social structures we take for granted.

1.3 Research questions

The report addresses four primary questions:

1. Framing: How do different normative frames—DEI-oriented, competitive-utility, and conservative—interpret AI in and by government, and what do they imply for policy design?
2. Practice: How are governments in Aotearoa New Zealand and the reference countries (Singapore, UK, Australia, Canada, and the US) actually using AI, and what are the dominant domains of deployment?
3. Governance and limits: What governance mechanisms are being used to manage risk, accountability, and legitimacy, and what do recent developments show about the drift from AI as decision support toward more agentic systems?
4. Aotearoa New Zealand implications: Given Aotearoa New Zealand's scale, institutional character, and Treaty obligations, what approach best preserves the upside of AI for state capability while avoiding institutional erosion and post-democratic drift?

1.4 Contribution and approach

Methodologically, the report is comparative and document-based. It combines (i) ideal-typical framing; (ii) structured comparison across six reference countries; (iii) synthesis of academic and policy literature; and (iv) the incorporation of bi-weekly scan findings from the last six months in an attempt to keep the empirical picture current (though in this field, innovation and change are very fast-moving).

It applies cross-cutting evaluative criteria: risk governance and assurance, administrative capacity, rule-of-law compatibility, and political economy. It also raises questions about the constitutional implications of AI: New Zealand's constitutional framework, inherited via British colonisation and embedded via Te Tiriti o Waitangi, is a standing constraint.

1.5 Structure of the report

Section 2 sets out more fully the methods and approach of this research report, including the three frames, the country comparison template, the cross-cutting evaluative criteria, and the Treaty/te ao Māori constitutional context for Aotearoa New Zealand. *Section 3* examines DEI-oriented AI governance. *Section 4* explores AI as an instrument of competitive national utility, foregrounding the merits of productivity, capability, and resilience while treating risks as constraints. *Section 5* presents conservative readings of AI in government, emphasising how institutional design can stabilise AI adoption and preserve accountability. *Section 6* turns to post-democratic trajectories, analysing how incremental expansions of AI autonomy could displace substantive human agency. *Section 7* provides a comparative synthesis across Aotearoa New Zealand and the reference countries. *Section 8* sets out recommendations for Aotearoa New Zealand. *Section 9* offers a concluding note.

2. METHODS AND APPROACH

This report is a comparative, conceptual, and normative analysis of (1) AI *in* government (public-sector adoption and use) and (2) AI *for* governing (the strategies, standards, and regulatory arrangements that condition how AI is used). It does not present new fieldwork or original quantitative modelling. Instead, it synthesises empirical studies, policy documents, and conceptual work to clarify the current frames through which governments are adopting and governing AI, and to draw implications for Aotearoa New Zealand's choices.

Methodologically, it combines (1) ideal-typical framing, (2) structured, focused comparison across a small set of reference countries, (3) document-based analysis of strategies, legislation, standards, and official reports (including bi-weekly scan updates), and (4) cross-cutting evaluative criteria and governing considerations focused on risk governance, administrative capacity, rule of law, political economy, and—specifically for Aotearoa New Zealand—Treaty obligations as a standing constraint.

2.1 Ideal-typical frames

The core of the approach is ideal-typical framing: the report uses three distinct frames as analytical tools to review real-world policies, strategies, and deployments. These frames are not empirical categories into which each country fits neatly; rather, they capture recurring rationales and priorities that appear in mixture across differing jurisdictions and institutions.

The framing is necessary because AI governance is not purely a technical set of decisions: global governance debates themselves are "key sites of regulatory contestation," and even the basics—"what exactly the 'it' is that is being governed, how, by who, and why"—are often unclear.[18] The report uses three referents:

- DEI-oriented governance: AI is assessed primarily through fairness, non-discrimination, representation, and participation, requiring explicit fairness constraints and impact assessments in areas where many people can be detrimentally affected.
- Competitive national utility: AI is treated as a strategic resource for state capability, productivity, resilience, and strategic advantage under conditions of complexity and rivalry.

Table 2. Ideal-typical frames for AI in and by government

Dimension	DEI-oriented frame	Competitive national utility frame	Conservative frame
Primary question	Who is disadvantaged and how can AI correct this?	How can AI increase state capability, productivity, and resilience?	How can AI be used without eroding law, institutions, and human responsibility?
Purpose of state	Equity, redress, group-level fairness	National power, efficiency, security, prosperity, and growth	Ordered liberty under stable law and institutions
Main AI aim	Bias correction, equitable outcomes, participation	Throughput, quality, foresight, enforcement capacity	Accountability, contestability, continuity, bounded automation
Typical focal values	Fairness, inclusion, representation	Efficiency, effectiveness, strategic advantage	Legality, the strength of the institutions of civil society, responsibility, prudence
Merits (if well executed)	Makes racial, gender, sexual orientation, age, disability, and socioeconomic bias visible; forces explicit goals	Clear productivity gains; stronger risk detection; better analysis	Stable adoption; clearer accountability; protection against drift

- Conservative readings: AI is judged by its compatibility with ordered liberty, institutional continuity, human responsibility, and rule-of-law accountability.

To keep the frames relevant rather than irritatingly rhetorical, Table 2 sets out their distinguishing features and recurring tensions.

2.2 Case selection and time frame

Aotearoa New Zealand is the report's primary case study. It is our nation. It is small with limited administrative scale, strong accountability expectations, and a highly distinctive constitutional context shaped by Te Tiriti o Waitangi.

The reference countries—Singapore, the United Kingdom, Australia, Canada, and the United States—are selected because they are active adopters of public-sector AI, publish their governance instruments (that can be analysed as texts), and represent different regulatory cultures and institutional designs.

The time frame is broadly 2022–2026. This captures the post-2022 generative-AI inflection and the observable shift from pilots to platforms, registers, and workforce-wide programmes. It also aligns with the emergence of AI-enabled state architecture frameworks, which argue that AI transformation is not simply "enhancement to existing systems" but that importantly it reshapes relationships between infrastructure, data, decision authority, and service delivery.[19]

2.3 Structured, focused comparison template

The report employs a structured, focused comparison template: the same questions are applied to each country to keep comparisons consistent. For each jurisdiction, the report asks:

- Aims: What explicit goals are attached to AI in government—equity, efficiency, security, innovation, trust?
- Domains: Where is AI actually used (internal operations, service delivery, case processing/eligibility, policymaking support, risk monitoring, oversight)?
- Instruments: What governance mechanisms exist—laws, strategies, standards, guidelines, ethics frameworks, procurement rules, registers, evaluation mechanisms, international agreements?
- Capacity: What infrastructure, data governance, workforce capability, institutional support, and procurement frameworks exist to make widespread adoption feasible?
- Dominant frame mix: Which of the three frames described above best captures the dominant orientation of a government, and where are the tensions?
- Trajectory risk: What is the direction of travel—AI as support (with algorithms) or partial authority (by algorithms)—and what guardrails exist?

To keep "domains" concrete, the report also draws on public-sector functional taxonomies that map where AI sits in government. For example, the Algorithmic State Architecture (ASA) positions transformation as a coherent system of interdependent layers—Digital Public Infrastructure, Data-for-Policy, Algorithmic Government/Governance, and GovTech—rather than parallel developments.[20]

Table 3. Functional domains for mapping AI use in government

Functional domain	Typical AI use (illustrative)	Why it matters (evaluative focus)
Internal operations	Document summarisation/classification, analysis, routing, drafting	Productivity gains; auditability; staff capability
Service delivery	Citizen chatbots, forms assistance, appointment systems (online platforms enabling citizens to book, modify, or cancel slots for in-person services)	Trust, accessibility, errors in public-facing outputs (responses from chatbots—textual answers to queries on benefits, taxes or regulations—summaries of information, drafted letters, personalised recommendations, etc.)
Case processing/ eligibility	Triage, risk flags, evidence extraction	Fairness, contestability, autonomy thresholds (the limits on an AI system's independence in decision-making, calibrated to the risk posed to human agency, safety or fundamental rights)
Risk monitoring and enforcement	Fraud/anomaly detection, compliance alerts	Integrity, proportionality (that the deployment, oversight, and intrusiveness of AI systems be calibrated strictly to the risks and benefits involved), false positives
Policymaking support	Consultation mining (extracting insights from large volumes of public consultation responses), scenario modelling	Transparency of reasoning; democratic legitimacy
Oversight and audit	Model monitoring, decision-pattern analysis	Rule-of-law safeguards; institutional capacity

2.4 Data sources: scholarship, policy texts, and bi-weekly scan updates

The report draws on three types of sources:

- Academic and conceptual literature on AI governance and public administration, including systematic reviews, conceptual frameworks, and empirical studies of adoption and human–AI interaction in decision-making.
- Policy and governance instruments, including national AI strategies, public service AI frameworks, technical standards, guidance on public generative-AI tools, AI registers, safety institute materials, procurement instruments, and international declarations.
- Bi-weekly scan updates from reputable sources (government releases, major media, high-quality policy analysis, and institutional publications).

The scan approach has been used in this instance to discover new material that adds value beyond the static literature—especially the move from pilots to platforms, registries, and workforce-wide programmes. The late-2025 and early 2026 updates are treated not as news trivia but as evidence of structural shifts (see Table 1).

2.5 Cross-cutting evaluative criteria and governing considerations

Ideal-typical frames discipline this research paper, but it evaluates each country's overall position using cross-cutting evaluative criteria. These also function as governing considerations when we formulate recommendations:

- Risk governance and assurance: how risks are identified, mitigated, monitored, and audited across a lifecycle; how autonomy is scaled to stakes; and what evaluation processes exist.

- Administrative capacity: data quality, infrastructure, skills, change management, procurement and vendor management, and the ability of oversight institutions to closely scrutinise AI use.
- Rule of law: legality, transparency, contestability, reasons-giving, and remedy; whether governance "by algorithms" undermines conventional accountability.
- Political economy: vendor dependence, platform chokepoints, incentives for symbolic adoption,[21] regulatory fragmentation, and centralisation versus local autonomy.

These criteria also reflect the basic insight that AI in government is not only a technical artefact but an institutional practice, embedded in procurement relationships and governance practices. Thus, Veale, Matus, and Gorwa caution that AI ethics principles can become "industry self-regulation" and urge analysts to keep asking "Who benefits?"[22]

2.6 Limitations and normative character

This report is document-based; it does not include new interviews, confidential procurement information, or internal performance metrics for specific systems. The bi-weekly scans capture major and credible developments but do not claim to offer comprehensive analyses across all national agencies and subnational bodies. Ideal-typical framing necessarily simplifies the intricate mixture of values and incentives that shape policy, but that said, it still offers comparative clarity.

The report is normatively explicit. It reports DEI-oriented governance objectively, but by preference places greater weight on competitive national utility and conservative readings when developing recommendations for Aotearoa New Zealand.

2.7 Te Tiriti o Waitangi and te ao Māori as constitutional context and standing constraints

For Aotearoa New Zealand, Te Tiriti o Waitangi and te ao Māori are treated as constitutional contexts and constraints that modify how all three frames should be applied. Treaty obligations introduce questions that resist generic fairness metrics or conventional privacy compliance: partnership, participation, protection, Māori authority over data, and the taonga status of Māori data and knowledge.

Such a claim is not speculative or evasive. The nation's own digital government guidance states: "We honour te Tiriti — the public service digital workforce understands Māori data is a taonga and is protected under te Tiriti o Waitangi, including upholding the rights and interests of Māori in the collection, ownership and application of their own data."[23] Moreover, Māori data sovereignty principles define Māori Data Sovereignty as the inherent rights and interests Māori have in relation to "the collection, ownership, and application of Māori data," and Māori Data Governance as the principles and structures through which Māori exercise control.[24]

Operationally, this means that whenever public-sector AI systems are trained on, infer from, or materially affect Māori data or Māori communities—particularly in domains such as justice, welfare, health, education, environmental regulation, and Treaty-relevant resource management—the evaluation must consider Treaty obligations and Māori governance expectations.

Māori data sovereignty claims are historically derived constitutional rights under Te Tiriti o Waitangi. They should not be confused with equity concerns grounded in a DEI frame. Comparisons with international examples of indigenous data governance frameworks may prove generative; however, they cannot provide normative grounds for rights claimed under the Treaty.[25]

3. DEI-ORIENTED AI GOVERNANCE

The first influential framing, DEI, casts AI in government as an instrument for *redress and recognition*. The central worry is that administrative systems already produce uneven burdens, and that automation will scale those burdens unless fairness and participation are treated as governing considerations from the start. In public-sector settings, the concern is not only model accuracy but the interaction of AI with institutional power, because "AI decision-making is a political necessity as it entails exercising power over citizens."[26]

DEI-oriented governance typically appears as part of "responsible AI" practice. One widely cited definition frames responsible AI governance as "a set of practices for developing, deploying, and monitoring AI applications… over the entire lifecycle."[27] In DEI-oriented policy and oversight, this definition matters because it signals that fairness is not an after-the-fact audit but a pressing, continuous requirement at every level of design, deployment, monitoring, and corrective action.

3.1 Core commitments and recurring principles

DEI-oriented governance generally insists on three commitments:

First, fairness and non-discrimination must be operationalised, not merely affirmed. Responsible AI frameworks commonly list "democratic participation, equity, [and] diversity inclusion" among core principles.[28] Here, mandating concrete evaluation methods becomes the standard: verifying that the information used to teach a system reflects the whole population; monitoring for results that might unfairly harm specific groups—an essential safeguard; and establishing rigorous inspection steps before using technology in life-altering situations.

Second, participation is treated as a legitimacy requirement. In the generative AI context, the governance literature explicitly stresses the need for "innovative and inclusive approaches" so that development aligns with "societal values," while also highlighting the shortcomings of technocratic, utilitarian governance and the need for broader stakeholder participation.[29] In practical terms, DEI governance asks: who is affected, who is consulted, and who has standing to contest the system's design and effects.

Third, the frame is attentive to "power asymmetries."[30] Public agencies often rely on vendors and inherited procurement arrangements that restrict transparency and constrain agency leverage. In one empirical study of social-services agencies, decision-makers reported that procurement contracts can "disincentivise private entities from being transparent" about model development and evaluation.[31] Analysts hold that this is a DEI-relevant point because opaque systems disproportionately harm those least able to contest decisions.

3.2 Typical governance mechanisms

In practice, DEI-oriented governance tends to rely on a repeatable set of mechanisms that can be understood as governing considerations that constrain adoption. Table 4 summarises the most common instruments and the specific problems they are meant to address.

Table 4. DEI-oriented governance mechanisms in public-sector AI

Governance mechanism	What it tries to prevent	Practical implementation signal
Bias/disparate impact testing	Systemic discrimination and unequal burdens	Pre-deployment testing and post-deployment monitoring for group outcome differences
Data representativeness checks	"Automating inequality" via biased historical data	Documentation of data provenance and sampling gaps; targeted data improvements
Stakeholder participation processes	Legitimacy deficit; "participation-washing"	Early-stage engagement with affected communities and feedback loops after deployment
Transparency and documentation	Opaque decision grounds and weak contestability	Public-facing system summaries; internal documentation; explainability thresholds
Redress and escalation pathways	Harm without remedy	Human review, appeals, audit triggers, incident response

This list overlaps substantially with general "responsible AI" governance. The key DEI distinction is what counts as failure: not only technical malfunction, but biased burden distribution, exclusion, and unaccountable power.[32]

3.3 Participation and implementation reality

While DEI discourse emphasises participation, implementation evidence shows how difficult it is to realise. In Kawakami et al.'s study, agencies reported "barriers and disincentives" to involving impacted communities and described leaders feeling "ill-equipped" to decide "how to mitigate power imbalances" in community involvement.[33] Participants also described the lack of available standardised routines and internal support systems—the mechanisms that would help achieve the goal of "true participation."[34]

This matters for DEI governance because it highlights a practical risk: agencies can default to consultation "theatre" when they lack time, money, and internal expertise for real participation. The same study's "power-conscious" recommendations explicitly warn against "participation-washing" and argue for leadership training so managers and staff can create conditions for meaningful collaboration.[35] In other words, participation is not a checkbox; it is an organisational capability with real costs.

3.4 Human factors: bias does not necessarily disappear when AI operates as "decision support"

DEI framing often assumes that using AI as a decision aid can reduce human bias. However, human–AI interaction research shows that bias can persist or reappear through how people use model outputs. In the public decision-support context, "automation bias refers to undue deference to automated systems," and is manifest as "the use of automation as a heuristic replacement for vigilant information seeking and processing."[36]

This creates a risk that even a fairness-audited model can produce unfair outcomes when public servants defer uncritically under workload pressure. The same work also identifies a second risk: selective use of algorithmic advice when it matches stereotypes. The authors explain that decision-makers may adopt algorithmic advice "selectively, when it matches pre-existing stereotypes" about decision subjects.[37] This directly bears on DEI governance because it implies that fairness is not only a property of the model, but of the full sociotechnical system—interfaces, workload, personnel, training, and institutional incentives.

3.5 Trade-offs and limits

DEI governance tends to broaden the set of objectives agencies must satisfy, which creates predictable trade-offs. Papagiannidis et al. note that "the diversity and breadth of responsible AI principles" make it challenging to implement practices that cover "a variety of goals,"

and they warn that a "principles-first approach may prove counterproductive."[38] For a small public service, this is not a theoretical objection; it is an administrative constraint that can turn governance into paperwork unless agencies invest in the skills and tooling required for securing effective work.

There is also a deeper tension between DEI objectives and broader public values. Hjaltalin and Sigurdarson observe that "While efficiency and service delivery dominate the discourse, citizen engagement remains underemphasized."[39] In this report, that tension is treated as real: DEI considerations function as genuine constraints and tests, while the primary recommendations are weighted toward competitive utility and conservative governing considerations.

3.6 DEI, indigenous rights, and Te Tiriti o Waitangi

In Aotearoa New Zealand, DEI-oriented governance overlaps with but does not exhaust the governing considerations introduced by Te Tiriti o Waitangi. Often, public discussion poorly distinguishes between the two. DEI frameworks typically focus on fairness between groups and on participation in decision-making; those concerns remain relevant, especially given the exposure of Māori to certain state systems (e.g., justice and social support). Where DEI commitments are strongly emphasised, Te Tiriti, or Māori interests generally, are listed among a host of other equity concerns. These may include other racial inequity or environmental and class concerns.

However, Treaty obligations introduce standing constraints that go beyond generic fairness metrics: partnership, participation, and protection in the Crown–Māori relationship, and questions about Māori authority over data and taonga. Accordingly, this report treats Treaty obligations and te ao Māori not as a subset of DEI, but as constitutional context that modifies how all three frames apply in Aotearoa New Zealand. DEI remains a reasonable set of evaluative criteria—especially for monitoring disparate impacts and participation failures. Treaty obligations add distinct requirements about legitimacy, governance, and authority, which are addressed explicitly in the Aotearoa New Zealand analysis and recommendations.

4. AI AS AN INSTRUMENT OF COMPETITIVE NATIONAL UTILITY

A second framing treats AI in government as an instrument of *competitive national utility*—a means of increasing a nation's capability, productivity, and resilience under conditions of complexity and rivalry. The OECD states this ambition in plain terms: "the use of AI by the public sector can increase productivity, responsiveness of public services, and strengthen the accountability of governments."[40]

This framing is not primarily about innovation and novelty; it is about capacity. Governments face rising expectations, persistent backlogs, and information volumes that human teams cannot process fast enough, and AI is presented as a way to recover administrative throughput with a smaller workforce or at the very least without proportionate increases in headcount.[41]

4.1 The competitive-utility value proposition

At its best, the competitive-utility narrative is a three-part claim: AI improves

- *productivity* (more output per unit of labour)
- *responsiveness* (faster, more tailored service delivery), and
- *accountability* (better detection of anomalies, fraud, and governance failures).

The OECD explains the productivity mechanism as turning data into operational advantage: AI systems can "gain more granular insights into user needs and identify patterns" in order to target interventions more effectively, while internal operations improve through "automating complex but repetitive administrative processes" to "free up the time of skilled civil servants" and support reliable service delivery.[42] In this framing, AI is not merely "automation"; it is organisational leverage: shifting staff time away from clerical bottlenecks toward tasks that require judgement and interaction.

US federal guidance advances the same argument in operational language. The GSA AI Guide notes that agencies face "increasing complexities and interdependencies" and "increasingly large volumes of data," concluding that "we are past the point where human cognitive abilities can directly process and make sense of all this information."[43] It then states the striking thesis sentence without hedging: "The federal government needs AI."[44] The competitive-utility frame treats this as a structural reality: modern government is information-intensive, and the state either acquires machine intelligence support to process that information or accepts persistent administrative underperformance and rising costs.

4.2 Case processing as the first large, bankable value pool

The strongest near-term competitive-utility case appears in *case processing* and related administrative pipelines: eligibility checks, file summarisation, routing, standard correspondence, and compliance workflows. Boston Consulting Group (BCG) frames the promise in maximalist terms: "Generative AI is the ideal resource to help reshape core government processing functions."[45] They then quantify the benefit: applied to areas like case processing, governments "can save up to 35% of budget costs in impacted areas over the next ten years."[46]

Those estimates are not a proof that every ministry or agency will achieve such savings, but they demonstrate how the competitive-utility narrative is routinely expressed: AI is treated as a budget-and-throughput lever large enough to matter at national scale. BCG also provides a crucial implementation insight that fits the governing considerations around administrative capacity: most work is not modelling; it is organisational change. The report states that in many AI transformations "70% of the efforts are focused on people and process improvements."[47] This aligns with the practical observation that public-sector AI projects fail to scale less often because the model is "bad" than because workflows, incentives, data quality, and staff training do not change. For Aotearoa New Zealand, this is a decisive point: competitive utility will not be realised by procurement alone; it depends on process redesign and the management and quality of staff.

4.3 How much can be supported by GenAI, and why it differs by domain

The Alan Turing Institute provides a task-based estimate that gives competitive-utility claims empirical shape: "approximately 41% of public sector time is spent on activities that could be supported through the use of

generative AI," ranging from 49% in Education to 33% in Healthcare.[48] This is arguably a conservative estimate given current advances in machine intelligence, but, in relation to this report, it does not constitute a claim about job replacement; it is an observation on where time is spent—drafting, summarising, record-updating, scheduling, and information retrieval. It implies that even partial support across those tasks will generate substantial capacity gains.

The same study also argues that to "scale at speed" is not always wise. It warns that a "blanket, one-size-fits-all approach to adoption... would likely lead to unintended consequences," producing missed opportunities in some domains and disorganised uptake in others.[49] For a competitive-utility strategy, this implies portfolio discipline: select high-yield workflows, pilot rapidly, scale where gains persist under real-world constraints (security, privacy, low error tolerance), and avoid trying to automate everything at once.

4.4 From pilots to capability: architecture and integration

Competitive national utility—the capacity for a government to provide reliable, large-scale digital services—is not achieved by isolated projects; it depends on a coherent architecture that ties infrastructure, data, algorithmic processes, and service delivery together. Establishing such a system ensures that technology serves the public interest consistently rather than functioning as a series of disconnected tools.

Engin et al. propose an "Algorithmic State Architecture (ASA), a novel four-layer framework" that conceptualises how "Digital Public Infrastructure, Data-for-Policy, Algorithmic Government/Governance, and GovTech interact as an integrated system in AI-enabled states."[50] They then make the key systems claim: each layer "builds upon and enables the capabilities of others, creating a coherent whole that is greater than the sum of its parts."[51] To understand this framework, one may view it as a building: the foundation consists of digital identity and secure channels, while the upper floors represent the specific applications citizens use to access benefits or information.

The logic of competitive utility expresses itself architecturally: without foundational infrastructure and disciplined data governance, algorithmic deployments will remain stuck at pilot level or degrade under operational stress. An algorithmic state architecture approach also highlights a common failure mode—a pattern of predictable errors—where governments treat technological elements as separate domains. When ministries and agencies operate with distinct governance frameworks and weak integration, the result is often "implementation failures, governance gaps, and missed opportunities."[52] Such fragmentation prevents a government from using information effectively to meet the needs of its citizens.

4.5 Governing considerations for scaling competitive utility

Competitive utility's method is not "adopt AI at any cost"; it is more in the style of "capture high-benefit uses while preventing predictable failure modes."[53] Here, two governing considerations are decisive: vendor dependency and assurance capacity. On vendor dependency, the GSA AI Guide warns that even when commercial products exist, relying too heavily on external parties creates "a major operating risk for an agency's core functions."[54] It argues government agencies should "focus on building their own sustainable institutional AI capability" and should not "overly rely on external parties such as vendors/contractors."[55] This is a type of competitive-utility reasoning that is not just moralising against private enterprise; rather, it posits that a state which cannot evaluate, govern, and, where necessary, replace vendor systems is *strategically* fragile.

On assurance capacity, the OECD's stance, for instance, is explicitly upside-weighted: "understanding, promoting, and enabling the positive aspects of using AI, rather than only preventing the negative ones, will remain a priority," because risk-only framing can deter "high-benefit, low-risk uses of AI to improve public services."[56] Competitive national utility therefore requires a practical assurance stance: not blanket bans, but risk-tiered governance that enables scaling where stakes are low and guardrails effective.

To make these claims concrete, Table 5 summarises the competitive-utility value chain—benefits, enabling conditions, governing considerations, and typical evidence sources.

Table 5. Competitive national utility: benefits, enabling conditions, and governing considerations

Element	What it means in practice
Productivity gains	Faster document handling, reduced backlogs, lower clerical workload
Service responsiveness	Faster citizen interaction, improved routing/triage, proactive support
Accountability uplift	Better anomaly detection, integrity checks, audit support
Enabling conditions	Data-for-policy maturity; secure service channels; staff training
Governing considerations	Vendor dependence; lifecycle assurance; autonomy proportional to stakes

4.6 Late-2025 developments that reinforce the competitive-utility trajectory

The most recent bi-weekly updates strengthen the competitive-utility story in two ways: (i) governments are shifting from pilots to platforms, and (ii) they are explicitly centralising or coordinating adoption to avoid fragmentation. In the UK, government–frontier-model partnerships are being framed as tools for modernising public services and reducing bureaucracy ("Gemini for Government"), reinforcing the idea that AI capability is now part of a state upgrading strategy.[57]

In Canada, the public AI Register is explicitly tied to planning and efficiency: it makes visible what exists across institutions, enabling reuse and reducing duplication.[58] In Australia, GovAI's public use case library is designed to move agencies from scattered experimentation to shared solutions and workforce capability uplift.[59] In Aotearoa New Zealand, the Govt.nz app roadmap signals that citizen-facing AI will likely be layered onto a unified digital public infrastructure channel, shifting AI from "tools inside government" to "interfaces between state and citizen."[60] These are not marginal developments. They indicate a structural transition that competitive-utility strategies require: shared platforms, inventories, and workforce-wide normalisation.

Moreover, recent US developments show how the competitive-utility posture is now being translated into hard procurement controls for large language models. In December 2025, the Office of Management and Budget issued Memorandum M-26-04, requiring agencies to update procurement policies so that contracts for LLMs include requirements aligned with "Unbiased AI Principles," and to establish user reporting processes for violating outputs, with a compliance deadline of 11 March 2026.[61] This illustrates a key dynamic for public-sector AI: governments do not only scale adoption; they increasingly attempt to define acceptable model behaviour through procurement instruments that sit upstream of technical implementation.[62]

4.7 Implications for later sections

The evidence supports a clear conclusion: competitive national utility is now a mature policy rationale with (i) plausible high-yield use cases, (ii) empirical task-exposure estimates, and (iii) observable institutional scaling. The remaining question is governance design: which governing considerations preserve the upside while limiting drift into opacity, risky autonomy, or strategic fragility. That question is addressed next through the conservative frame, and then stress-tested through a post-democratic scenario, before being translated into Aotearoa New Zealand-specific recommendations that incorporate Treaty obligations and te ao Māori as constitutional context and standing constraints.

5. CONSERVATIVE READINGS OF AI IN GOVERNMENT

A third framing evaluates AI in government through the governing considerations of *ordered liberty, institutional continuity, and human responsibility*. It begins from a hard premise: administrative AI is not "just software," because "AI decision-making... entails exercising power over citizens."[63] On this view, the central question is not whether AI can improve outputs, but whether it can be integrated without eroding the conditions of lawful and accountable government, and institutional relations: reasons-giving, contestability, remedy, identifiable responsibility, and the "little platoons" or institutions of civil society.[64]

This frame is not anti-AI. Its advantage is that it can make AI adoption *durable*: it shifts the focus from pilots and hype to enforceable governance, institutional capability, and explicit boundaries on automation. Such is especially valuable in small states like Aotearoa New Zealand, where reputational failure can freeze adoption for years and where vendor dependency can quietly become a structural constraint on sovereignty.

5.1 From principles to enforceable governance

A recurring conservative diagnosis is that principle statements do not govern by themselves. In a literature review on AI governance, Birkstedt et al. note that "principle-based ethics provide limited assurance that the principles are met in practice," because principles often focus on the "what" rather than the "how," and therefore "must also be enforceable through governance."[65]

The conservative value of this point is practical: where DEI and "trustworthy AI" manifest as lists, the conservative response is institutionalisation—standards, auditability, decision rights, and accountability structures that survive political cycles. This also clarifies why conservative governance is compatible with scaling. It does not treat risk management as an obstruction; it treats it as the condition for continued adoption. In practice, it means building the same disciplines that high-reliability organisations require: clear responsibility, controlled change, monitoring, and incident response—particularly where systems drift over time or behave differently in edge cases.

5.2 Human oversight, bounded autonomy, and democratic control

Conservative readings are particularly concerned with the distribution of agency between humans and systems. Engin et al.'s algorithmic framework describes the Process Layer (Algorithmic Government) as implementing "AI-enabled decision-making and algorithmic automation... while maintaining appropriate human oversight."[66] They add that success depends on "well-designed transfer mechanisms" so systems can hand over to human operators when they hit edge cases.[67] This is exactly the workings of a conservative framework in operational form: automation is permissible where humans remain responsible and can intervene meaningfully, and where the institutions of society are thus strengthened and preserved.

Algorithmic state architecture also embeds a broader constitutional point: it elevates accountability as a design principle, stating that it "ensures the preservation of human oversight and democratic control throughout the system."[68] In its recommendations, it makes the governing consideration explicit: "Maintain democratic oversight: Ensure that AI-enabled transformation strengthens rather than undermines democratic accountability."[69] This provides a clear conservative criterion for evaluating high-stakes uses: does the system increase a state's capacity while leaving authority visible and contestable, or does it smuggle decisions made into opaque pipelines?

5.3 The human factor: decision-makers as safeguards, and the limits of that hope

A common reform claim is that "humans in the loop" solves the legitimacy problem. Conservative analysis is sceptical of that slogan, because it is easy to keep formal human sign-off while eroding substantive human judgement. Alon-Barkat and Busuioc note that in the "rise of algorithmic governance," human decision-makers are treated as "important safeguards," but they argue it becomes "critical" to discover whether "our cognitive limits allow us to act as effective decisional mediators."[70]

Their work defines automation bias as "a well-documented human propensity to automatically defer to automated systems" despite warning signals or contradictory information.[71] For conservative governance, the implication is not "ban decision support," but strengthen governing considerations that protect judgement: training, slowed decision tempos in high-stakes contexts, structured override requirements, and auditing of when staff follow or reject algorithmic advice. It also reinforces why contestability matters: if an affected person cannot challenge the grounds of a decision, then "human oversight" becomes a paper shield.

The caveat on this is whether humans are, or will be in the near future, capable of following the processes of algorithmic reasoning and be able to better its judgements. While the concern to detect and deflect bias is laudable, it might prove impossible.

5.4 Institutional capability and vendor dependence as governance problems

Importantly, conservative readings treat procurement and capability as governance questions, not back-office details. The US General Services Administration's AI Guide warns that even if an agency can find an adequate commercial off-the-shelf product, "using it would be a *major operating risk* for an agency's core functions to rely so much on external parties."[72] This is not ideological posturing; it is a sovereignty and continuity argument. When core state functions rely on vendor-controlled systems, a nation's ability to change, contest, or disengage with those same systems becomes structurally constrained.

The same conservative logic also applies to the "phantomisation" of government AI: shared platforms and enterprise tools can yield productivity gains, but they concentrate power and create chokepoints. Conservative governance therefore favours diversification, strong internal evaluation capability, and contractual rights that preserve transparency and auditability, especially where models are updated frequently and where operational dependencies deepen over time.

5.5 Safety, drift, and continuous assurance

A conservative approach to governance is also attentive to time and systems drift. A strong statement of operational requirements appears in FTI Consulting's presentation on "AI Governance in Practice." It contends that "Generative AI models drift. The only way... to know when/how they are drifting is to continuously test, monitor and audit the AI applications... every second of every minute of every day."[73] The prose is intentionally emphatic, but the underlying governing consideration is sound: static compliance cannot manage dynamic systems, especially those deployed at scale.[74]

This is where conservative governance aligns with adaptive approaches. In a paper on generative AI and adaptive governance, Araz Taeihagh argues governance must "match and mimic the iterative development, speed, and collaboration patterns" of generative AI, and that policy organisations should be "ambidextrous," able to pursue contradictory aims such as close monitoring and significant innovation.[75] In conservative terms, this is not "move fast and break things"; it is build feedback loops and controlled experimentation so that governments and public services can implement and use AI without being surprised by emergent failure.

5.6 Conservative governing considerations in one view

Table 6 consolidates the conservative governing considerations developed above into a practical checklist. The point is not to multiply bureaucracy; it is to make AI adoption legible and controllable so that scaling is politically and institutionally sustainable.

Table 6. Conservative governing considerations for AI in government

Governing consideration	What it prevents	Concrete mechanisms and instruments
Non-delegable responsibility	"Nobody responsible" outcomes	Human sign-off for rights-affecting decisions; named decision owners
Bounded autonomy	Quiet drift from "with" to "by" algorithms	Cognitive forcing functions; autonomy tiers; transfer-to-human mechanisms
Democratic accountability	Opaque, uncontestable rule by systems	Registries; reasons-giving; public scrutiny of high-impact uses
Contestability and remedy	Harm without recourse	Appeals; independent review; audit trails; explainability thresholds
Institutional capability	Vendor capture; fragile dependence	Model portability clauses; inspection rights; diversification
Continuous assurance	Drift, silent degradation	Monitoring, testing, auditing; incident response; post-deployment review

5.7 Implications for Aotearoa New Zealand (including Treaty obligations as standing constraints)

For Aotearoa New Zealand, the conservative frame carries a particular force because it aligns with institutional realities: a small public service, high trust expectations, and strong rule-of-law norms. It also aligns with the report's Treaty framing. Treaty obligations and te ao Māori as constitutional contexts operate as standing constraints on what constitutes legitimate "utility" gains—especially where AI systems are trained on, infer from, or materially affect Māori data or Māori communities.

The conservative approach is therefore not merely a generic "Western" caution; it is an argument for preserving the integrity of Aotearoa New Zealand's constitutional and institutional order, which includes Treaty obligations that shape legitimacy in high-impact domains. The practical result is an approach marked by disciplined adoption: capture high-benefit uses in internal operations and service delivery while enforcing strict boundaries, transparency, and contestability in rights-affecting and Treaty-sensitive areas.

6. AI REPLACING PUBLIC SERVANTS AND POLITICIANS: TOWARD A POST-DEMOCRATIC ORDER?

This section addresses a plausible trajectory rather than a normative programme: the gradual shift from AI as administrative support to AI as de facto governing infrastructure. The public sector's predicament is well-captured by Misuraca et al., who argue that the nation-state faces a "tragic double bind": obligations to protect citizens from algorithmic harms collide with the temptation to increase its own efficiency—"to govern algorithms, while governing by algorithms."[76] The risk is not only that the state adopts AI, but that the state's mode of governing changes subtly: substantive human judgement is displaced by systems that shape decisions, priorities, and citizen interaction at scale.

6.1 Why automation pressure intensifies

The drivers are familiar: fiscal constraint due to structural deficits and an ageing population, rising demand, shortages of specialised labour, ease of use, and political pressure to reduce backlogs. A competitive-utility logic insistently presents AI as the best response to information overload and administrative bottlenecks.[77] But the post-democratic question arises because, as AI tools move from drafting and triage into recommendation engines and agentic workflows, the line between "support" and "governance" becomes easier to cross in practice.

6.2 Autonomy shifts: from "with" to "by" algorithms

A clear way to describe post-democratic drift is through autonomy modes. Commentators distinguish "government / governance by algorithms" from "government / governance with algorithms," and they define the "by algorithms" mode as a shift where "significant decision-making agency" moves to algorithmic systems and humans "assume supervisory roles… reviewing algorithmic recommendations rather than generating solutions," which "challenges conventional accountability structures."[78]

The risk of automated governance remains incremental: administrative bodies may continue to claim that "humans are in the loop" even as the core logic of decision-making migrates into model pipelines. When the substantive reasoning behind a policy or action shifts to an algorithm, the human role risks becoming performative rather than substantive; consequently, the oversight mechanism functions as a facade for automated logic.

Engin et al. contend that maintaining control requires intentionality rather than assumption. One of their primary design principles for public-sector AI insists on "the preservation of human oversight and democratic control throughout the system." Such a standard addresses the fundamental challenges of modern governance: technical autonomy must be strictly bounded and audited—preventing "black box" outcomes—while authority must remain subject to public contestation, and institutional accountability must be explicitly defined across every technical and administrative level of the system.

Table 7 provides a practical autonomy spectrum and indicates the main guardrails needed at each step.

Table 7. Autonomy spectrum in government and required guardrails

Mode	What it looks like	Main post-democratic risk	Minimum guardrails (governing considerations)
Government *with* algorithms	Decision support: AI drafts/ summarises; humans decide	Automation bias; rubber-stamping	Training; mandatory human reasoning notes for high-stakes decisions; audit of overrides
Government *by* algorithms	AI allocates, prioritises, recommends at scale; humans supervise	Accountability diffusion; policy by pipeline	Explicit autonomy tiers; transfer-to-human mechanisms; public registries; contestability
Government *of* algorithms	Meta-governance: align models to missions/values; enforce transparency and remedy	Formal governance without real leverage over vendors	Audit rights; documentation standards; liability provisions; independent oversight capacity
Government *through* platforms	AI embedded into core infrastructure (identity, service portals, secure messaging)	Citizen-state relationship mediated by model outputs	Authentication and logging; visible disclaimers; redress pathways; output monitoring

6.3 Human psychology: why "human-in-the-loop" can fail

Post-democratic drift is not only structural; it can be cognitive. Alon-Barkat and Busuioc argue that in "the rise of algorithmic governance," humans must be treated as "important safeguards," and that it will be critical to test whether human "cognitive limits allow us to act as effective decisional mediators."[79] They define automation bias as "a well-documented human propensity to automatically defer to automated systems," even in the presence of warning signals or contradictory evidence.[80] Such a response may be a mark of trust—an acknowledgement that the machine intelligence is seen as highly "able" and reliable. But the risk is that the user passes responsibility for the work commissioned over to an AI with no further thought, or more seriously, the unspoken conviction that it will do a better job anyway.

If agencies scale decision-support tools without investing in training, pacing, and institutional norms for challenge and override, human oversight will likely become formal rather than substantive. The implication for governance is direct: post-democratic risk is heightened when decision-makers are rushed, overloaded, or structurally discouraged from dissenting from machine recommendations. That risk increases further when AI advice is perceived as "objective," a dynamic that intersects with the classic bureaucratic claim to neutrality and that, when scaled, changes citizens' experience of "government."

The same risk runs right through each layer of government from elected representatives to public service graduate employees. As each becomes more familiar with AI and the advantages it offers in a wide range of operations, it could easily become second nature to use it continuously. Like the use of electricity, it will be taken for granted and ubiquitous. But the challenge remains: who will actually be deciding on the adoption of recommendations and acting upon them? Or will this be passed to a machine intelligence, too?

6.4 Procurement opacity and vendor lock-in as hidden sovereignty risk

Post-democratic dynamics also arise through procurement and platform dependencies. Kawakami et al. report that department and agency leaders felt constrained by procurement contracts and by vendor reluctance to explain systems: "We cannot share that information with you because it's proprietary."[81]

The post-democratic risk here is not so much that elected officials disappear, but that critical state functions become dependent on opaque systems that government cannot fully inspect, explain, or replace. This risk is now exacerbated by the shift from single-purpose tools to whole-of-government platforms and large vendor partnerships.[82] The UK's late-2025 partnership with Google DeepMind, including exploration of "Gemini for Government," illustrates the strategic scale of such compacts. Canada's public AI Register, by contrast, illustrates a transparency countermeasure that can help states see and manage the portfolio of AI dependencies.[83] These are two different ways of responding to the same environment: one emphasises rapid capability via partnership; the other emphasises visibility and governance as a precondition for scale.

6.5 "Policy optimisation" and governance-by-experiment

A more advanced post-democratic scenario is governance-by-experiment: AI systems optimising policy through continuous feedback loops that treat populations as data-generating environments. Mökander and Schroeder's tax-policy thought experiment makes this vivid: repeated large-scale experimentation could "shape the behaviour and preferences of the population in line with" a policy, and "shaping the preferences of citizens… would mean control through (AI-generated) knowledge… [but] it would mean treating the population as experimental subjects."[84] Perhaps this is akin to current practices where government consults, develops and introduces law and policy, monitors their effects, and modifies them accordingly. But it is one thing to have human consultation, argumentation, authoring, implementation, and feedback, and quite another to have all this done by a machine intelligence in the blink of an eye.

This said, such a form of governance is not primarily about replacing politicians with robots; it is about relocating political choice into model design (goal functions, constraints, levers), where public, democratic contestation is harder.

Paradoxically, the same authors note a positive aspect that cuts the other way: policy optimisation can force explicitness. "The advantage of AI-driven policy optimisation is that it requires that normative ends are made explicit and formalised, thereby subjecting them to public scrutiny and debate."[85] A well-designed governance regimen can harness that advantage—explicit goal-setting and constraint definition—without accepting the experimental-subject dynamic.

All this assumes a high level of human competency and benign, controllable developments with AI.[86] Yet it remains a key bridge into recommendations: the goal is not to ban optimisation thinking, but to ensure it remains publicly accountable and institutionally bounded.

6.6 Implications for Aotearoa New Zealand and Treaty obligations as standing constraints

For Aotearoa New Zealand, post-democratic risk is not only a matter of efficiency and autonomy. Treaty obligations operate as standing constraints on what counts as legitimate AI-driven governance, particularly where Māori data is used, where Māori communities are unduly affected, or where Treaty-sensitive domains (e.g., environmental governance involving iwi/hapū rights and interests) are implicated.

A post-democratic settlement in Aotearoa New Zealand would not merely weaken democratic accountability; it would also risk weakening the Crown's ability to honour Treaty obligations in practice, especially if decision systems become impossible to see into or vendor-controlled. Accordingly, the remainder of the report treats post-democratic drift as a stress test for strategy. Competitive national utility justifies substantial AI adoption; conservative governance specifies the guardrails; and Treaty obligations impose additional constraints on data architecture, co-governance expectations, and domain boundaries.

The practical aim is to keep AI adoption in the "with algorithms" mode for key decision-making, to build "of algorithms" governance (registries, auditability, contestability), and to treat "by algorithms" deployments as exceptional and bounded by strict autonomy and oversight regimes.

7. COMPARATIVE SYNTHESIS: AOTEAROA NEW ZEALAND AND REFERENCE COUNTRIES

This section compares Aotearoa New Zealand, Singapore, the United Kingdom, Australia, Canada, and the United States across the three frames (DEI-oriented, competitive national utility, conservative readings), using the report's evaluative criteria (risk governance, administrative capacity, rule-of-law compatibility, and political economy) and Aotearoa New Zealand's Treaty obligations as standing constraints.

The goal is not encyclopaedic country description; it is to make visible how different jurisdictions are converging (pilots to platforms) and diverging (standards-led vs vendor-compacts vs registers), and to identify what is most strategically relevant for Aotearoa New Zealand. Across the reference set, there is a strong empirical pattern: public-sector AI is currently used mainly to improve service delivery and internal operations, and far less to directly drive policymaking. The EU landscaping of 250 cases found AI "used mainly to support improving public service delivery," then internal management, with policy decision-making relatively limited.[87] This aligns with the competitive-utility account: the early value pool is in document-heavy administrative pipelines and service channels, not in replacing democratic institutions.

7.1 Aotearoa New Zealand

In late 2025 the Government moved to build a unified citizen-facing digital channel as a platform for future AI-mediated interaction. The official Govt.nz app page now states that "In upcoming releases we will deliver an AI assistant, digital credentials via the wallet and services from government." This is an architectural choice: it places AI and automation on top of a consolidating digital public infrastructure layer, rather than leaving AI confined to internal agency workflows.

This citizen-facing channel is one facet of a larger shift towards a digital target state, which "sets the direction for improved government digital services. The vision is to transform the public sector operating model, moving towards a customer-oriented approach."[88]

Aotearoa New Zealand's policy stance is also distinctive because of the way it attends to Treaty obligations and te ao Māori. The Public Service AI Framework makes this explicit: it recognises Māori data as taonga and includes commitments that AI use should be compatible with democracy and the rule of law.[89] In practical terms, this means Aotearoa New Zealand's governance choices—data linking, vendor contracts, automation in justice/welfare/health, and the design of citizen-facing assistants—cannot be assessed solely through generic DEI or efficiency criteria.

7.2 Singapore

Singapore remains the clearest and arguably best example in this comparison set of a high-discipline, integrated "AI-enabled state" stance, combining capability-building with technical guardrails and sectoral supervision. The current public narrative emphasises augmentation rather than replacement, with the message that AI should serve as a "co-pilot, not an autopilot."[90] This framing is aligned with conservative governing considerations: bounded autonomy, strong governance, and explicit accountability, while still pursuing competitive utility at scale.

A key recent development is the Monetary Authority of Singapore's consultation on proposed AI risk management guidelines for financial institutions. It explicitly positions AI risk governance as lifecycle governance for all AI systems used by financial institutions.[91] For Aotearoa New Zealand, Singapore functions less as a simple "model" and more as a reference point for how a government can pair national capability-building with sectoral risk governance and technical guardrails.

7.3 United Kingdom

The UK is increasingly pursuing a blended model: competitive-utility adoption framed as modernisation and growth, coupled with a strong emphasis on safety, evaluation, and governance legitimacy. The recent partnership with Google DeepMind illustrates this: government frames AI as a driver of "national renewal and growth," including the exploration of "Gemini for Government" to reduce bureaucracy and automate routine tasks in the public sector.[92] This is competitive utility expressed as a whole-of-state project: modernising public services, accelerating science, and advancing national security.

The UK is also moving from isolated pilots toward a centrally coordinated "toolchain" approach. A January 2026 announcement by digital government leadership

describes an AI Engineering Lab and a planned early-2026 rollout of licences for leading AI tools, explicitly including GitHub Copilot, Gemini, Amazon Q/Kiro, and Claude, with accompanying technical setup and training support.[93] Here is a specific instance of the "pilots to platform portfolio" shift: rather than each department procuring tools *ad hoc*, government is attempting to curate and support a multi-tool environment for civil servants at scale.

At the same time, the UK is positioning itself as a coordinator of global evaluation regimes. In December 2025, the International Network of AI Safety Institutes transitioned into the "International Network for Advanced AI Measurement, Evaluation and Science," with the UK as network coordinator.[94] Thus the UK provides an important comparative point for Aotearoa New Zealand's evaluative criteria: a shift from abstract safety principles to shared scientific methods of measurement and evaluation.

7.4 Australia

Australia is notable for a standards-forward approach that aims to enable adoption at scale while keeping governance enforceable. This is consistent with the conservative impulse to institutionalise guardrails rather than rely on principles alone. The public-facing GovAI ecosystem provides unusually concrete evidence of how AI is being embedded in day-to-day government work: GovAI is described as "foundational AI for the APS," designed to uplift capability, improve productivity, and provide a secure environment for experimentation.[95] APS Academy highlights concrete use cases (e.g., automated transcription and structured analytics for archives; summarisation of parliamentary material; policy-aligned chatbots), making Australia unusually transparent about what "utility" looks like in practice (APS Academy 2025).

Australia's scaling mechanism is now clearer in institutional terms. GovAI describes AIDE (Artificial Intelligence Delivery and Enablement) as working collaboratively with GovAI to implement key aspects of the APS AI Plan, with GovAI serving as a "secure, centralised technology foundation" providing infrastructure, model options, development tools, and a growing library of use cases, alongside development and trial of a general-purpose GovAI Chat tool planned for 2026.[96] This provides a concrete comparator for Aotearoa New Zealand: whole-of-government AI capability is being built not only through guidance, but through dedicated enablement functions and shared technical foundations.

For Aotearoa New Zealand, Australia provides a near-peer comparator in institutional form (Westminster traditions, public trust concerns), and a useful example of how to make AI adoption measurable and governable through standards and shared services rather than only procurement deals.

7.5 Canada

Canada's most distinctive recent move is a transparency-and-planning instrument: the public AI Register. The Register "includes input from 42 institutions and features over 400 systems" at stages from research and proof-of-concept to fully deployed tools supporting operations and service delivery.[97] This is not only a public transparency tool; it is a managerial instrument for avoiding duplication and guiding portfolio governance. It aligns closely with the report's conservative governing considerations (contestability, visibility, auditability) while still supporting competitive-utility adoption.

Canada's AI Register is now explained in methodological terms: the government describes the Register as a "minimum viable product" assembled from existing information sources, including Algorithmic Impact Assessments, Access to Information requests, Parliamentary Question responses, Personal Information Banks, and the federal Service Inventory.[98] For Aotearoa New Zealand, Canada's *modus operandi* provides a practical template for how to bootstrap a register without waiting for perfect instrumentation: begin with a minimum viable product assembled from existing disclosure pathways, then iterate toward fuller coverage and stronger metadata requirements.

Canada's broader policy context also increasingly ties AI to productivity and "sovereign" capability. Budget 2025 messaging highlights "made-in-Canada sovereign AI tools" and an Office of Digital Transformation to lead adoption across government.[99] For Aotearoa New Zealand, Canada is a high-value comparator because it shows how a medium-sized democracy can combine safety institutions, registries, and procurement strategies without abandoning liberal-democratic governance norms.

7.6 United States

The US remains the most scale-intensive environment in the comparison set, but the late-2025 updates sharpen two themes relevant to this report: centralisation and politicisation. First, the December 2025 executive order

"Ensuring a National Policy Framework for Artificial Intelligence" explicitly aims to eliminate "state law obstruction" of national AI policy.[100] Reuters reports that the order ties compliance to broadband funding and contemplates federal litigation against restrictive state AI laws (Reuters 2025).[101] This is a competitive-utility move: centralisation justified as necessary to avoid fragmentation and to keep pace with rival nations.

Second, US procurement governance is now explicitly imposing "ideological neutrality" requirements on large language models sold to federal agencies, building on the July 2025 executive order on "Preventing Woke AI in the Federal Government."[102] Whether one agrees with the policy or not, it is analytically significant: it embeds an overtly philosophical and political criterion into technical procurement, and it illustrates how AI governance in government can become a proxy arena for broader cultural and political conflict.

At the time of writing this report, the US Department of Defense and the Pentagon voided their contracts with Anthropic, the producer of the LLM, Claude. Reportedly, "the Pentagon needs to be able to utilise models free from usage policy constraints that may limit lawful military applications.... Anthropic, meanwhile, doesn't want its technology used for operations including domestic surveillance and autonomous lethal activities."[103] In consequence of its policy stance, Anthropic lost its business with government, and may be classified as a supply chain risk.[104]

For Aotearoa New Zealand, this is a cautionary comparator: choosing procurement criteria can lead to a conflict of interest between government contractors and vendors, and possibly increase the likelihood of privacy breaches, while also eroding public trust and distorting meaningful evaluation of the technologies that have been adopted. The difficulty lies in the fact that the private sector has the most up-to-date and powerful forms of machine intelligence and that governments, when they can afford it, want it.

A second point for consideration is that both current and previous US administrations have encouraged a light regulatory touch with AI companies, especially in comparison with Europe.[105] This carries the advantages of creating the conditions for dynamic and remarkable innovation, but also the concomitant risks—ranging from machine hallucinations, ineptitude, and deliberate false reporting to those risks associated with the pending realisation of an artificial general intelligence that may reason and act outside the bounds set for it. The US and, to a lesser extent, the UK governments want to adopt, centralise, and use the best technology available generated through competition within the private sector, but face the challenges that come with doing so—be they individual company policy positions, opacity around proprietary rights, or the adoption of models that do not deliver what they have promised.

7.7 Cross-cutting observations

Several cross-cutting observations follow.

- **Convergence on administrative AI first.** Across jurisdictions, AI is used mainly for internal operations and service delivery. This is consistent with the EU landscaping evidence (comprehensive research on how AI is being used) and with the task-exposure pattern in public sector work.[106]
- **Shift from pilots to platforms and registers.** Australia's GovAI, Canada's AI Register, and Aotearoa New Zealand's unified Govt.nz channel all point to the same structural shift: AI is becoming a whole-of-government capability rather than an unevenly timed agency-by-agency experiment (GovAI 2025; Treasury Board 2025; Govt.nz 2025).
- **Evaluation is becoming an institution.** The UK-led transition of the AI Safety Institute network toward measurement and evaluation reflects a move from principles to scientific governance practice (UK Government 2025). This aligns with conservative governing considerations: controlled experimentation, monitoring, the ability to justify decisions, adoption, and change.
- **Vendor compacts are now a governance tool.** The UK–DeepMind MoU and partnership announcements show a model of state modernisation via frontier-model partnerships, raising questions about shared vision, dependency, accountability, and procurement leverage (UK Government 2025; DeepMind 2025).
- **Treaty obligations make Aotearoa New Zealand's governance problem distinct.** Aotearoa New Zealand cannot treat "equity" as a generic DEI setting; Treaty obligations and te ao Māori as constitutional contexts require additional governance arrangements—especially around Māori data as taonga, data governance, and co-decision expectations in Treaty-sensitive domains (NZ Digital Government 2025).

7.8 Comparative snapshot and evaluative matrix

To consolidate the comparison, Table 8 summarises the dominant frame mix and key governance instruments across the six jurisdictions, while Table 9 provides a high-level evaluative matrix for the three frames (used as analytic tools rather than as moral scorecards).

Table 8. Comparative snapshot: dominant frames, governance stance, and recent instruments

Jurisdiction	Dominant frame mix	Current "pilot to platform" signal	Notable late-2025 instrument(s)
Aotearoa New Zealand	Mixed (utility and conservative constraints and Treaty obligations; DEI as constraint)	Unified citizen channel and AI guidance	Public Service AI Framework; Govt.nz app roadmap
Singapore	Utility and conservative-style governance discipline	Blueprint and sectoral guidelines	MAS AI risk consultation; co-pilot framing
United Kingdom	Utility and evaluation emphasis; rights language	Partnerships and evaluation network	DeepMind partnership and MoU; international evaluation network
Australia	Utility and standards-forward conservative discipline	GovAI shared service	GovAI platform and use case library
Canada	Utility and safety/transparency posture	Public AI Register	AI Register; Budget 2025 sovereign tools
United States	Strong utility and centralisation; politicised procurement constraint	Federal pre-emption of state rules	Executive order on national AI framework; "bias test" procurement direction

Table 9. Evaluative matrix: frames against core criteria

Evaluative criterion	DEI-oriented frame	Competitive national utility frame	Conservative frame
Productivity and throughput	Mixed (generally adds compliance load)	High (central aim)	Medium (accepts friction for legitimacy)
Fairness and disparate impact	High (explicit focus)	Indirect (if built-in)	Indirect but present through equal law, oversight
Rule-of-law compatibility	Variable (depends on transparency and legal clarity)	At risk if opacity is tolerated	High (legality, contestability, remedy)
Administrative feasibility	High demand for audits/ consultation	High demand for infrastructure/capability	Moderate demand (oversight and boundaries)
Political economy risk	Co-opted rhetorically	Risk of vendor lock-in and government by opaque AI systems	Risk of status quo bias
Post-democratic drift	Medium (metric governance)	High if autonomy scales unchecked	Lower if boundaries and accountability are enforced

8. RECOMMENDATIONS FOR AOTEAROA NEW ZEALAND

Aotearoa New Zealand is already in broad experimentation with AI in government, but the next phase requires a coherent strategy that preserves the advantages of AI for state capability while preventing predictable failure modes: uncontrolled drift in autonomy, weak contestability, vendor lock-in, and loss of legitimacy in Treaty-sensitive domains. The recommendations below are organised as governing considerations and implementation steps. They are anchored in the competitive national utility and conservative frames, with DEI treated as a constraint and Te Tiriti o Waitangi / te ao Māori treated as constitutional contexts and standing constraints.

8.1 Set an explicit "capability and accountability" purpose statement

Aotearoa New Zealand should state plainly that public-sector AI is pursued for capability: higher productivity, faster responsiveness, and stronger oversight and enforcement, while remaining bounded by law and accountability. OECD's formulation is a useful benchmark: AI "can increase productivity, responsiveness of public services, and strengthen the accountability of governments."[107]

This purpose statement should be used as a decision gate in procurement and programme design: if a proposal cannot show (i) concrete capability gains and (ii) enforceable accountability pathways, it does not proceed. A common failure mode is treating governance as a defensive appendix. OECD warns that governments should prioritise "enabling the positive aspects of using AI" because a risk-only posture can deter "high-benefit, low-risk uses of AI to improve public services."[108] In Aotearoa New Zealand, that means making capability the driver, while defining rule-of-law, Treaty obligations, and safety as non-negotiable limits that shape how capability is pursued.

8.2 Build an AI-enabled state architecture incrementally, not by accretion

Aotearoa New Zealand should adopt an explicit architectural roadmap rather than letting AI grow by accretion. The ASA framework offers a strong conceptual scaffold: AI-enabled states depend on coherent interaction between Digital Public Infrastructure, Data-for-Policy, Algorithmic Government/Governance, and GovTech.[109] It also diagnoses the core failure mode: treating technical domains as separate produces "governance gaps" and "missed opportunities" where the technology never reaches optimisation.[110]

The immediate practical implication for Aotearoa New Zealand is to treat identity, secure citizen channels, data standards, model governance, and service tooling as a coordinated programme rather than dispersed agency projects. This is directly relevant to Aotearoa New Zealand's current trajectory: the Govt.nz app is explicitly designed as a unified service channel, with planned "AI assistant" and a digital wallet for credentials. That "front door" will become the default interface through which many citizens experience AI-mediated government. Therefore, logging, authentication, redress pathways, and clear disclosure must be designed now, before conversational agents become routine.

8.3 Prioritise high-leverage use cases and avoid "mission creep" into high-stakes domains

Aotearoa New Zealand should prioritise the domain where competitive utility is most credible: internal operations, service delivery, and case processing. BCG's report frames the opportunity and scale: "Generative AI is the ideal resource to help reshape core government processing functions"[111] and claims up to 35% budget cost savings in impacted areas over ten years.[112] The implementation point is just as important: "70% of the efforts" are in "people and process improvements."[113]

New Zealand should therefore target workflows where AI support can be embedded into processes with clear handoffs, rather than pursuing high-profile "AI transformation" programmes without operational redesign. The Turing Institute's task-exposure estimate offers a disciplined way to avoid overreach: about 41% of public-sector time could be supported by GenAI, but the distribution varies by sector, and one-size adoption risks "unintended consequences."[114] New Zealand should treat this as an adoption principle: start where support is high and error tolerance is moderate, and scale only when assurance, training, and auditability keep pace.

Table 10. Priority NZ use cases and governing considerations

Domain / use case	Example	Primary frame(s)	Governing considerations
Case processing triage	Classify and route applications	Competitive utility; DEI as constraint	Bias monitoring; cognitive forcing (blind review); human sign-off for adverse outcomes
Document-heavy regulation	Summarise and extract evidence	Competitive utility; conservative constraint	Audit trails; reasons-giving; strict data handling
Fraud/anomaly detection	Pattern detection on claims/returns	Competitive utility; conservative upside	False-positive control; proportional interventions; oversight review
Citizen information	Chatbots for FAQs/status	Competitive utility	Clear disclosure; monitoring; escalation to humans
Policymaking support	Summarise submissions, draft briefs	Competitive utility; conservative constraint	Provenance of sources; human responsibility; versioning and logging
Treaty-sensitive domains	Any AI affecting Māori rights/interests	Treaty obligations as standing constraints	Mandatory Māori data audit; Māori governance expectations; no autonomy creep

8.4 Create a public-sector AI register and use it as a governance instrument

Aotearoa New Zealand should establish a public-sector AI register modelled on Canada's approach. Canada's Treasury Board launched a federal AI register that includes "input from 42 institutions and features over 400 systems," spanning research, pilots, and deployed tools.[115] The utility of a register is not merely transparency; it is portfolio governance: duplication reduction, reuse, prioritisation, and the ability to see autonomy creep.

A New Zealand register should include at minimum: purpose, domain, data sources, model type, autonomy level, risk tier, procurement/vendor dependency, monitoring regime, and redress pathways. It should also mark whether the system is Treaty-sensitive and what governance arrangements exist for Māori data. This will make Treaty obligations operational and not just rhetorical.

8.5 Adopt "bounded autonomy" via technical friction and cognitive forcing

Aotearoa New Zealand should formally adopt an autonomy stance that treats AI as an augmentation to, rather than a replacement for, human judgement in rights-affecting decisions. The risk otherwise is that AI shifts agency to systems and "challenges conventional accountability structures."[116] To prevent this drift, agencies should implement cognitive forcing functions—interface designs that require the human operator to input their own assessment of the evidence *before* the AI recommendation is revealed.

This mechanism counters "automation bias," described as the "undue deference" to automated systems even when contradictory evidence exists.[117] By enforcing a "blind review" step in high-stakes workflows, government ensures that the human in the loop remains an active decision-maker rather than a passive reviewer. Furthermore, Aotearoa New Zealand should establish explicit autonomy tiers: mandatory transfer-to-human mechanisms must trigger automatically when systems encounter edge cases or low-confidence intervals.

8.6 Build continuous assurance capacity and incident response

Aotearoa New Zealand should treat monitoring and audit as the condition for scaling AI, especially for generative AI. The FTI governance report states bluntly: "Generative AI models drift," and the only way to know how is to "continuously test, monitor and audit" systems "every second… every day."[118] This supports a conservative governing consideration: compliance cannot be static when systems and contexts change.

Aotearoa New Zealand should therefore require: (i) pre-deployment evaluation; (ii) post-deployment monitoring; (iii) incident reporting; and (iv) independent audit capacity for high-risk systems. The Monetary Authority of Singapore's AI risk guidelines for financial institutions show how a regulator can formalise lifecycle expectations for AI risk management in a critical sector.[119] New Zealand should consider whether analogous sectoral guidance is needed for domains like financial services, health, and critical infrastructure.

8.7 Treat procurement as a governance instrument

Aotearoa New Zealand faces a dilemma. We are, by virtue of our size, consumers rather than producers of this technology. We cannot choose what capability is built, only how we use and regulate it. Therefore, procurement becomes a constitutional function, not merely a purchasing decision, given the role that AI tools may play in governing decisions in the foreseeable future. Reliance on external parties creates "a major operating risk" for core functions if the state cannot independently inspect or modify the tools it uses.[120] To mitigate this, agencies should mandate three specific provisions in contracts with frontier-model providers:

1. **Model portability:** Contracts must ensure the state retains the right to export fine-tuned weights and training data to alternative platforms, preventing slippage into structural vendor lock-in.
2. **Inspection rights:** The state must reserve the right to subject "black box" systems to independent validation using negotiated and controlled test sets, rather than relying solely on vendor-supplied performance metrics.
3. **Liability and indemnity:** Contracts must clarify that vendors retain liability for system failures attributable to design defects, preventing the outsourcing of risk alongside the outsourcing of computation. This approach ensures that partnerships, such as those seen in the UK, do not permanently constrain Aotearoa New Zealand's sovereign ability to govern its own administrative architecture.

8.8 Workforce capability: environmental design and critical vigilance

Aotearoa New Zealand should pursue workforce-wide AI literacy and further invest in specialist capability.

Government agencies should embed key staff members in relevant sister government agencies in aligned countries with a specific emphasis on AI governance. While reports such as this are important, much of the knowledge and wisdom required for a capable public service will be tacit, and can only be learned hands on.

Broad-based AI literacy must focus on environmental design rather than simple tool proficiency. Research indicates that decision-makers often use automation as a "heuristic replacement for vigilant information seeking," leading to rubber-stamping of machine outputs.[121] Training alone is an insufficient control for this propensity.

Instead, agencies must design the decision environment to support independence. This involves:

1. **Interface architecture:** Designing workflows where AI outputs are presented as supporting evidence rather than conclusions, and where the provenance of that evidence is clearly visible.
2. **Critical norms:** Establishing an institutional expectation that divergence from an AI recommendation is a valid and often necessary exercise of judgement, rather than an error requiring justification.
3. **Management capability:** Equipping leaders to recognise "false inclusion," where community consultation or human oversight is performed without genuine influence on the outcome.[122]

Finally, externally procured AI training must be independent. AI product providers have a vested interest in providing "training" that steers organisations towards particular

platforms. Care must be taken to ensure that upskilling workforce capability does not become procurement by other means.

8.9 Te Tiriti o Waitangi as an operational constraint

Aotearoa New Zealand should integrate Treaty obligations into public-sector AI governance as operational rules, not just aspirational values. Guidance already establishes that "Māori data is a taonga" and that the Crown must uphold Māori rights in the "collection, ownership and application" of that data.[123] To ensure this in practice, agencies should implement a mandatory Māori data audit for any proposed system.

This audit must determine if a system: (i) is trained on Māori data; (ii) infers attributes regarding Māori identity; or (iii) determines access to resources in domains where Māori are disproportionately affected (e.g., justice, health). If a system meets these criteria, it triggers specific governance requirements, including partnership mechanisms that go beyond standard consultation. This ensures that data "sovereignty" and "protection" are treated as governing considerations in the system architecture itself.

8.10 Implementation sequencing: start with high-benefit, low-risk; scale only with assurance

Aotearoa New Zealand should explicitly adopt an implementation sequence that captures competitive utility without drifting into post-democratic automation. It should start with internal operations and service support where error tolerance is higher and harms can be contained; then expand to risk monitoring and compliance; and treat any movement toward autonomous decision-making in high levels of governance and rights-affecting domains as exceptional, bounded, and subject to higher thresholds of transparency and remedy.

9. CONCLUSION AND FUTURE DIRECTIONS

This report has argued that AI in and by government cannot be understood—or governed—adequately without making explicit the frames that structure adoption and constrain legitimacy. AI is already embedded in public administration. Its most common uses are not in autonomous agency but practical throughput gains: document handling, triage, case processing, risk monitoring, and citizen-facing information support. That said, agentic and generative AI are now steadily seeing uptake across the private sector and the pressure on public servants and government employees to adopt it, in some cases in the privacy of their own homes, will increase.

The strategic question for Aotearoa New Zealand is therefore not whether AI will enter state mechanisms—this is already happening—but how to harvest the capability gains without degrading accountability, legality, and trust.

There is a second set of related questions, beyond the scope of this report. At what point would leaders in government hand authority for decision-making and action over to machine intelligences or a machine intelligence (as in an artificial general intelligence)? A crisis, such as an infrastructure cyber-attack? In which areas of government would this occur—defence, transport and communications, health? What criteria would be settled upon for permitting agentic AI to "take control"—including considerations for kill switches and containment measures? And how will humans relate to autonomous (agentic) AI—or, more importantly, how might it relate to them—friend, slave, enemy? These questions, seemingly absurd now, will need to be addressed within five years as AI becomes more and more a part of our lives.

9.1 What the three frames clarify

In the meantime, the three ideal-typical frames used in this report identify real features of the governance problem. *DEI-oriented governance* is strongest at revealing distributional harms and participation failures. It insists that administrative AI does not merely "predict"; it governs, and that if bias or exclusion are embedded, they scale. It also contends that bias is not only a model property but an interaction effect: decision-support systems can still produce unfair outcomes through automation bias and selective adherence. However, DEI governance also exposes an implementation constraint that matters for small public services: governance ideals can become "principles-first" lists that are removed from reality, hard, and costly to operationalise at scale.

Competitive national utility provides the strongest account of why governments are moving fast, and why they should do so. The OECD's proposition that AI can increase "productivity, responsiveness… and accountability" is not simply rhetorical; it maps onto observable deployments and measurable task exposure.[124] The strongest near-term value pool is administrative case processing and internal operations, where large volumes and repeated steps make support tools materially valuable. But competitive utility also has predictable failure modes: vendor lock-in, fragmented governance, autonomy creep, and the selection of procurement criteria. It will remain a very powerful frame for AI because of the nature of competition in industry and between nations,[125] and because of the desire for ease and efficiency.[126]

Conservative readings do not deny utility; they specify the governing considerations that make utility durable: enforceable accountability, bounded autonomy, contestability, and remedy, and internal institutional capability rather than total dependence on vendors. This is why the conservative frame aligns naturally with a "platform era" of public-sector AI: it treats procurement, monitoring, drift, and audit as core governance rather than compliance afterthought. The operational demand is continuous: "Generative AI models drift," and the only way to check this trend is to "continuously test, monitor and audit" deployed applications.[127] A conservative frame aims at preserving the integrity of New Zealand's institutions and those things that matter to its people; and at acknowledging that while change is inevitable, it should be prudently managed. Given the magnitude of change AI will bring to government and society, its value is significant.

9.2 What the late-2025 and early 2026 updates show

Recent bi-weekly updates reinforce the intuition that governments are moving from pilots to *platforms, registers, and whole-of-workforce capability programmes*. Canada's public AI Register changes the transparency baseline by enumerating hundreds of systems and their status, shifting governance from scattered inventories to portfolio visibility. Australia's GovAI programme and published use case library

show a deliberate attempt to normalise safe adoption through shared environments and reusable patterns. The UK–DeepMind partnership illustrates a different mode of scaling: capability via frontier-model compacts, paired with the evolution of the AI safety institute network into an evaluation and measurement network.

The recent US pre-emption order and procurement "bias test" requirements show that AI governance can also become a vehicle for centralisation. Singapore's Government AI Blueprint and MAS consultation show how a high-discipline state can combine whole-of-government capability building with sectoral risk supervision.

For Aotearoa New Zealand, the key takeaway is structural: the next stage of adoption is not simply "more pilots," but will likely be the creation of governance infrastructure—registers, standards, audit capacity, and citizen-facing digital public infrastructure on which AI will sit. The Govt.nz app roadmap signals that Aotearoa New Zealand is already moving toward the infrastructure-first path: AI will increasingly mediate citizen-state interaction through shared channels. That increases the premium on early design choices about logging, disclosure, escalation, and redress.

9.3 Core synthesis: capability disciplined by accountability

The report's central synthesis is that Aotearoa New Zealand's best route is *competitive national utility disciplined by conservative governing considerations*, with DEI functioning as a constraint, and Treaty obligations as standing constraints. The objective is not to slow AI adoption for its own sake, but to ensure that high-benefit uses scale while rights-affecting uses remain bounded, contestable, and institutionally responsible.

Two practical claims follow:

- **The near-term benefit is real and large, especially in internal operations and case processing.** OECD describes the mechanism clearly—automation of complex but repetitive processes, freeing skilled civil servants for higher level work, and improving reliability. This can and should be pursued.
- **The sustainability of that benefit depends on governance infrastructure:** public-sector AI registers, autonomy tiers, continuous assurance, audit rights, internal capability, and redress. Without these, scaling will trigger legitimacy failures and political backlash that freeze adoption.

9.4 Future directions

This report also identifies areas where further work is needed.

- **Aotearoa New Zealand portfolio visibility.** Establishing a public-sector AI register and maintaining it as a living governance instrument should be a near-term priority. Canada provides a working comparator for scale and structure.
- **Capability for oversight institutions.** Ombudsman, Office of the Auditor-General, courts, and privacy/human-rights bodies need technical support to scrutinise AI systems at scale. AI should strengthen oversight capacity, not bypass it.
- **Autonomy thresholds and agentic AI.** Governments are now shifting from "what GenAI can write" to "what AI agents can do." Aotearoa New Zealand should pre-define autonomy tiers and domain constraints before agentic tools become normal.
- **Treaty-consistent data governance.** Māori data as taonga must be treated as an operational governance rule, not a statement of values. This includes governance of data sharing, training corpora, and use cases in Treaty-sensitive domains.
- **Measurement and evaluation.** The trend toward international evaluation and measurement networks should be watched and joined where appropriate. Evaluation standards will increasingly shape procurement, interoperability, and risk governance.

9.5 Closing statement

AI is already part of government. The question is whether it will remain a tool that strengthens the state's capacity while preserving accountable, lawful rule—or whether it will become a pathway for opaque centralisation, the incremental erosion of human responsibility, and algorithmic government that inexorably moves towards a post-democratic society when AI makes choices as to how society should best be served. The competitive-utility case for adoption is strong, particularly in administrative workflows; but the conservative case for enforceable guardrails is equally strong if adoption is to endure in a beneficial manner.

In Aotearoa New Zealand, those commitments must be nested within Treaty obligations and te ao Māori as constitutional context. The task ahead is therefore not to choose between capability and legitimacy, but to institutionalise capability in a way that remains recognisably legitimate.

BIBLIOGRAPHY

Government documents

Aotearoa, Te Kāwanatanga o, and New Zealand Government. *Public Service AI Framework*. (2026). https://docref.digital.govt.nz/nz/generative-ai-guidance-gcdo/public-service-ai-framework/2025/en/. Accessed March 2026.

Artificial Intelligence Delivery and Enablement (AIDE), and Australian Public Service Commission. *Artificial Intelligence Delivery and Enablement*. (2026). https://www.govai.gov.au/aide. Accessed March 2026.

Australian Government. *APS AI Use Case Library*. 2026. https://www.govai.gov.au/explore/use-cases. Accessed March 2026.

Australian Government. Australian Public Service Academy. *Exploring AI Innovations: Highlights from the GovAI 'AI Use Case Library'*. 2025. https://www.apsacademy.gov.au/news/exploring-ai-innovations-highlights-govai-ai-use-case-library. Accessed March 2026.

California State Government. *California AI Risk Assessment Report*. Sacramento: California State Government, 2025.

Congressional Research Service. *Algorithmic Accountability Act: Implementation Challenges and Opportunities*. Washington, DC: Congressional Research Service, 2024.

European Commission. *EU AI Act Implementation Guide for Public Sector*. Brussels: European Commission, 2024.

European Commission. Nemitz, Paul. *Constitutional Democracy and Technology in the Age of Artificial Intelligence*. Belgium: European Commission, 2018.

Executive Office of the President, and Office of Management and Budget. "Increasing Public Trust in Artificial Intelligence through Unbiased AI Principles." December 11, 2025. https://www.whitehouse.gov/wp-content/uploads/2025/12/M-26-04-Increasing-Public-Trust-in-Artificial-Intelligence-Through-Unbiased-AI-Principles-1.pdf. Accessed March 2026.

Government of Canada. *AI Register*. (2025). https://open.canada.ca/data/en/dataset/fcbc0200-79ba-4fa4-94a6-00e32facea6b. Accessed March 2026.

Government of Canada. *Government of Canada AI Register: Progress on AI in Government*. (2026). https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/progress.html. Accessed March 2026.

Government of Canada. Treasury Board. *Register of AI Uses in Federal Government*. 2025. https://www.canada.ca/en/treasury-board-secretariat/news/2025/11/canada-launches-first-register-of-ai-uses-in-federal-government.html. Accessed March 2026.

Govt.nz App. 2025. https://www.govt.nz/about/the-govt-nz-app/. Accessed March 2026.

GovTech Singapore. "A Co-Pilot, Not a Replacement: How AI Will Augment Human Labour in the Workforce." 2025. https://www.tech.gov.sg/technews/how-ai-will-augment-human-labour-in-the-workforce. Accessed March 2026.

National AI Advisory Committee. *National AI Advisory Committee Year 1 Report*. Washington, DC: Department of Commerce, 2023.

National Science Technology Council. *AI Research and Development Progress Report 2020–2024*. Washington, DC: Executive Office of the President, 2024.

New York City Government. *New York City Automated Decision Systems Task Force Report*. New York: NYC Government, 2024.

New Zealand Government | Department of Internal Affairs. "Digital Capability Framework 2025: Upholding Māori Rights and Interests in Data." (2025). https://www.digital.govt.nz/. Accessed March 2026.

Office of Management and Budget. *OMB Memorandum M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence*. Washington, DC: Executive Office of the President, 2024.

Royal Commission into the Robodebt Scheme. Canberra: Department of the Prime Minister and Cabinet, 2023. https://robodebt.royalcommission.gov.au/publications/report. Accessed March 2026.

Singapore, Monetary Authority of. "Consultation Paper on Guidelines on Artificial Intelligence Risk Management." 2025. https://www.mas.gov.sg/-/media/mas-media-library/publications/consultations/bd/2025/final_consultation_paper_on_guidelines_on_ai_risk_management_forrelease.pdf. Accessed March 2026.

Stats NZ. *Algorithm Assessment Report*. Wellington: Stats NZ, 2018. https://data.govt.nz/use-data/analyse-data/government-algorithm-transparency.

The White House. Executive Order. *Ensuring a National Policy Framework for Artificial Intelligence*. 2025. https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy. Accessed March 2026.

U.S. Department of Defense. *Department of Defense AI Strategy Update*. Washington, DC: Department of Defense, 2024.

U.S. Department of Justice. *AI in Criminal Justice: Federal Agency Practices and Oversight*. Washington, DC: Department of Justice, 2024.

U.S. Government Accountability Office. *Artificial Intelligence - Use and Oversight in Financial Services*. Washington, DC: Government Accountability Office, 2025.

U.S. Government Accountability Office. *Artificial Intelligence - Agencies Are Implementing Management and Personnel Requirements*. Washington, DC: Government Accountability Office, 2024.

U.S. Government Accountability Office. *Artificial Intelligence - GAO's Work to Leverage Technology and Ensure Responsible Use*. Washington, DC: Government Accountability Office, 2024.

U.S. Government Accountability Office. *Generative AI Use and Management at Federal Agencies*. Washington, DC: GAO, 2025. https://www.gao.gov/products/gao-25-107653.

UK Department for Science, Innovation and Technology. *UK Government AI White Paper Implementation Update*. London: HM Government, 2024.

UK Gov | Government Digital and Data. "Join AI Engineering Lab and Help Shape the Future of Digital Government." (January 12, 2026). https://cddo.blog.gov.uk/2026/01/12/join-ai-engineering-lab-and-help-shape-the-future-of-digital-government/. Accessed March 2026.

UK Gov. "AI to Accelerate National Renewal and Growth as Google DeepMind Backs UK Tech and Science Sectors." November 2025. https://www.gov.uk/government/news/ai-to-accelerate-national-renewal-and-growth-as-google-deepmind-backs-uk-tech-and-science-sectors. Accessed March 2026.

United States Department of Health and Human Services. *Healthcare AI Governance: Federal Policy and Implementation*. Washington, DC: HHS, 2024.

United States General Services Administration. AI Center of Excellence. *AI Guide for Government*. Washington, DC: General Services Administration, 2024. https://coe.gsa.gov/coe/ai-guide-for-government/print-all/index.html. Accessed March 2026.

Journal articles

Alon-Barkat, Saar, and Madalina Busuioc. "Human–AI Interactions in Public Sector Decision Making: 'Automation Bias' and 'Selective Adherence' to Algorithmic Advice." *Journal of Public Administration Research and Theory* 33, no. 1 (2022): 153–69. https://doi.org/10.1093/jopart/muac007.

Batool, Amna, Didar Zowghi, and Muneera Bano. "AI Governance: A Systematic Literature Review." *AI and Ethics* (2024). https://link.springer.com/article/10.1007/s43681-024-00653-w.

Birkstedt, Teemu, Matti Minkkinen, Anushree Tandon, and Matti Mäntymäki. "AI Governance: Themes, Knowledge Gaps and Future Agendas." *Internet Research* (2022). https://doi.org/10.1108/intr-01-2022-0042.

Bradley, Anusha. "The Rise of Cybercrime and NZ's Fragmented Response." *New Zealand In Depth*, 2021. https://www.rnz.co.nz/news/in-depth/444735/the-rise-of-cybercrime-and-nz-s-fragmented-response.

Buhmann, Alexander, and Christian Fieseler. "Deep Learning Meets Deep Democracy: Deliberative Governance and Responsible Innovation in Artificial Intelligence." *Business Ethics Quarterly* (2022): 1–34. https://doi.org/10.1017/beq.2021.42.

Burgess, Paul. "Algorithmic Augmentation of Democracy: Considering Whether Technology Can Enhance the Concepts of Democracy and the Rule of Law through Four Hypotheticals." *AI & Society* 37 (2022): 97–112. https://doi.org/10.1007/s00146-021-01170-8.

Carroll, Stephanie Russo, Ibrahim Garba, Oscar L. Figueroa-Rodríguez, Jarita Holbrook, Raymond Lovett, Simeon Materechera, Mark Parsons, et al. "The CARE Principles for Indigenous Data Governance." *Data Science Journal* 19 (2020): 43. https://doi.org/10.5334/dsj-2020-043.

Cho, Eunsun. "The Social Credit System: Not Just Another Chinese Idiosyncrasy." *Princeton Journal of Public and International Affairs* (May 2020). https://jpia.princeton.edu/news/social-credit-system-not-just-another-chinese-idiosyncrasy.

Christodoulou, Eleni, and Kalypso Lordanou. "Democracy under Attack: Challenges of Addressing Ethical Issues of AI and Big Data for More Democratic Digital Media and Societies." *Frontiers in Political Science* 3 (2021). https://doi.org/10.3389/fpos.2021.682945.

Coeckelbergh, Mark. "Democracy, Epistemic Agency, and AI: Political Epistemology in Times of Artificial Intelligence." *AI and Ethics* (2022): 1–10. https://doi.org/10.1007/s43681-022-00239-4.

Couldry, Nick, and Ulises A. Mejias. "Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject." *Television & New Media* 20, no. 4 (2019): 336–49. https://doi.org/10.1177/1527476418796632.

Cutter, Chip. "Here's Where AI Is Tearing through Corporate America." *WSJ*, 2026. https://www.wsj.com/tech/ai/ai-software-business-stock-market-4b17b432?mod=saved_content. Accessed March 2026.

Donahoe, Eileen, and Megan MacDuffee Metzger. "Artificial Intelligence and Human Rights." *Journal of Democracy* 30, no. 2 (2019): 115–26. https://doi.org/10.1353/jod.2019.0029.

Engin, Zeynep and Philip Treleaven. "Algorithmic Government: Automating Public Services and Supporting Civil Servants in Using Data Science Technologies." *The Computer Journal* 62, no. 3 (2019): 448–60. https://doi.org/10.1093/comjnl/bxy082.

Feldstein, Steven. "The Road to Digital Unfreedom: How Artificial Intelligence Is Reshaping Repression." *Journal of Democracy* 30, no. 1 (2019): 40–52. https://doi.org/10.1353/jod.2019.0003.

Filgueiras, Fernando. "The Politics of AI: Democracy and Authoritarianism in Developing Countries." *Journal of Information Technology & Politics* 19, no. 4 (2022): 449–64. https://doi.org/10.1080/19331681.2021.2016543.

Gupta, Suman. "Procedural Democracy and Electronic Agents." *AI Magazine* 42, no. 1 (2021): 30–38.

Hagey, Keach, Amrith Ramkumar, Deborah Acosta, and Vera Bergengruen. "'Woke' AI Feud Escalates between Pentagon and Anthropic." *Wall Street Journal*, February 18, 2026. https://www.wsj.com/politics/national-security/woke-ai-spat-escalates-between-pentagon-and-anthropic-433b7c5c. Accessed March 2026.

Hubert, Kent F., Kim N. Awa, and Darya L. Zabelina. "The Current State of Artificial Intelligence Generative Language Models Is More Creative Than Humans on Divergent Thinking Tasks." *Scientific Reports* 14 (2024): 3440. https://doi.org/10.1038/s41598-024-53303-w.

Janssen, Marijn, Martijn Hartog, Ricardo Matheus, Aaron Yi Ding, and George Kuk. "Will Algorithms Blind People? The Effect of Explainable AI and Decision-Makers' Experience on AI-Supported Decision-Making in Government." *Social Science Computer Review* 40, no. 2 (2022): 478–93.

Keall, C. "Ransomware Attacks: Privacy Commissioner Plans Investigation as Justice, Health Hit." *NZ Herald*, December 6, 2022.

König, Pascal, and Georg Wenzelburger. "Opportunity for Renewal or Disruptive Force: How Artificial Intelligence Alters Democratic Politics." *Government Information Quarterly* 37 (2020). https://doi.org/10.1016/j.giq.2020.101489.

König, Pascal. "Citizen Conceptions of Democracy and Support for Artificial Intelligence in Government and Politics." *European Journal of Political Research* 61, no. 4 (2022): 1–21. https://doi.org/10.1111/1475-6765.12570.

Manheim, Karl, and Lyric Kaplan. "Artificial Intelligence: Risks to Privacy and Democracy." *The Yale Journal of Law and Technology* 21 (2019): 106–88.

Marchant, Gary, and Jordan Buckwald. "Private Standards as Liability Shields: A Pro-Innovation Artificial Intelligence Regulatory Approach for States." *Minnesota Journal of Law, Science & Technology* 27, no. 1 (2026). https://scholarship.law.umn.edu/mjlst/vol27/iss1/. Accessed March 2026.

Mason, Max, John Davidson, and Ayesha de Kretser. "Inside Australia's Most Invasive Data Hack." *Australian Financial Review*, November 9, 2022.

Mergel, I., H. Dickinson, J. Stenvall, and M. Gasco. "Implementing AI in the Public Sector." *Public Management Review* (2023). https://www.tandfonline.com/doi/full/10.1080/14719037.2023.2231950.

Misuraca, Gianluca, and Colin van Noordt. "Artificial Intelligence for the Public Sector: Results of Landscaping the Use of AI in Government across the European Union." *Government Information Quarterly* (2022). https://doi.org/10.1016/j.giq.2022.101714.

Papagiannidis, Emmanouil, Patrick Mikalef, and Kieran Conboy. "Responsible Artificial Intelligence Governance: A Review and Research Framework." *Government Information Quarterly* (2024). https://doi.org/10.1016/j.giq.2024.101929.

Sandhane, Raghav. "Artificial Intelligence in Cyber Security." *Journal of Physics: Conference Series*, 2021.

Schechner, Sam. "When AI Bots Start Bullying Humans, Even Silicon Valley Gets Rattled." *The Wall Street Journal*, February 13, 2026. https://www.wsj.com/tech/ai/when-ai-bots-start-bullying-humans-even-silicon-valley-gets-rattled-0adb04f1. Accessed March 2026.

Schroeder, Ralph. "Aadhaar and the Social Credit System: Personal Data Governance in India and China." *International Journal of Communication* 16 (2022). https://ijoc.org/index.php/ijoc/article/view/19059.

Taeihagh, Araz. "Governance of Generative AI." *Policy & Society* 44, no. 1 (2025). https://doi.org/10.1093/polsoc/puae035.

Tallberg, Jonas, Eva Erman, Markus Furendal, Johannes Geith, Mark Klamberg, and Magnus Lundgren. "Global Governance of Artificial Intelligence: Next Steps for Empirical and Normative Research." *International Studies Review* 25, no. 3 (2023). https://doi.org/10.1093/isr/viad040.

Yampolskiy, Roman V. "On the Controllability of Artificial Intelligence: An Analysis of Limitations." *Journal of Cyber Security and Mobility* 11, no. 3 (2022): 321–404. https://doi.org/10.13052/jcsm2245-1439.1132.

Zaidan, Esmat, and Imad Antoine Ibrahim. "AI Governance in a Complex and Rapidly Changing Regulatory Landscape: A Global Perspective." *Humanities and Social Sciences Communications* (2024). https://doi.org/10.1057/s41599-024-03560-x.

Zuiderwijk, Anneke, Yu-Che Chen, and Fadi Salem. "Implications of the Use of Artificial Intelligence in Public Governance: A Systematic Literature Review and a Research Agenda." *Government Information Quarterly* (2021). https://doi.org/10.1016/j.giq.2021.101637.

Reports and Special Publications

Agarwal, Akshat, Charu Singhal, and Renny Thomas. *AI-Powered Decision Making for the Bank of the Future*. New York: McKinsey & Company, 2021. https://www.mckinsey.com/~/media/mckinsey/industries/financial%20services/our%20insights/ai%20powered%20decision%20making%20for%20the%20bank%20of%20the%20future/ai-powered-decision-making-for-the-bank-of-the-future.pdf.

Benaich, Nathan, and Ian Hogarth. *State of AI Report 2024*. London: State of AI, 2024.

Boston Consulting Group (BCG). *Generative AI for the Public Sector: The Journey to Scale*. BCG, 2024. https://www.bcg.com/publications/2024/gen-ai-journey-to-scale-in-government.

Boston Consulting Group (BCG). Heidi Kim, Richard Sargeant, Miguel Carrasco, Joerg Hildebrandt, Steven Mills, Christopher Daniel, Madeleine Shaw, and Ethan Gromet. *How AI Can Cut through Bureaucracy, Boost Efficiency, and Build Trust in Government*. 2025. https://www.bcg.com/publications/2025/benefits-of-ai-in-government. Accessed March 2026.

Boston Consulting Group (BCG). Hildebrandt, Peter, and Eunyoung Kim. *Benefits of AI in Government*. 2025. https://www.bcg.com/publications/2025/benefits-of-ai-in-government.

Bradshaw, Philip N., and Samantha Howard. *Challenging Truth and Trust: A Global Inventory of Organised Social Media Manipulation*. The Computational Propaganda Project. University of Oxford, 2018.

Center for Democracy and Technology. *Exploring the 2024 Federal AI Inventories: Key Improvements, Trends, and Continued Inconsistencies*. Washington, DC: Center for Democracy and Technology, 2025.

Fagan, M. *AI for the People: Use Cases for Government*. Harvard Kennedy School, 2024. https://www.hks.harvard.edu/sites/default/files/centers/mrcbg/working.papers/M-RCBG%20Working%20Paper%202024-02_AI%20for%20the%20People.pdf.

Fei-Fei Li, Christopher Manning, and Anka Reuel. *The Stanford Emerging Technology Review: A Report on Ten Key Technologies and Their Policy Implications*. Stanford University, 2026.

Foundation, Electronic Frontier. *Deepening Government Use of AI and E-Government Transition in Latin America: 2024 in Review*. San Francisco: Electronic Frontier Foundation, 2025.

Gillespie, N., S. Lockey, C. Curtis, J. Pool, and A. Akbari. *Trust in Artificial Intelligence: A Global Study*. The University of Queensland and KPMG Australia, 2023.

Institute for Law. *Advanced AI Governance: A Literature Review of Problems, Options, and Proposals*. Institute for Law & AI, 2025. https://law-ai.org/advanced-ai-gov-litrev/.

International Association of Privacy Professionals. *AI Governance in Practice Report 2024*. IAPP, 2024. https://iapp.org/resources/article/ai-governance-in-practice-report/.

International Telecommunication Union. *Key Findings on the State of Global AI Governance*. ITU, 2024. https://www.itu.int/hub/2024/07/key-findings-on-the-state-of-global-ai-governance/.

M. Myllylahti and G. Treadwell. *Trust in News in Aotearoa New Zealand 2023*. AUT Research Centre for Journalism, Media and Democracy (JMAD), 2023. https://www.jmadresearch.com/trust-in-news-in-new-zealand.

Meta Ad Library Report: New Zealand. Meta, 2023. Accessed May 26, 2023. https://www.facebook.com/ads/library/report/?source=nav-header.

National Institute of Standards and Technology. *NIST AI Risk Management Framework (AI RMF 1.0)*. Gaithersburg, MD: NIST, 2024.

Olawale, Salaudeen, Anka Reuel, Ahmed Ahmed, Suhana Bedi, Zachary Robertson, Sudharsan Sundar, Ben Domingue, Angelina Wang, and Sanmi Koyejo. *Measurement to Meaning: A Validity-Centered Framework for AI Evaluation*. arXiv, 2025. https://arxiv.org/abs/2505.10573.

Organisation for Economic Co-operation and Development (OECD). *AI in Government Report 2024*. Paris: OECD, 2024. https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en.html. Accessed March 2026.

Organisation for Economic Co-operation and Development (OECD). *Artificial Intelligence in Society*. Paris: OECD, 2021.

Organisation for Economic Co-operation and Development (OECD). *Governing with Artificial Intelligence: Are Governments Ready?* 2024. https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/06/governing-with-artificial-intelligence_foe316f5/26324bc2-en.pdf. Accessed March 2026.

Organisation for Economic Co-operation and Development (OECD). Observatory of Public Sector Innovation. *Artificial Intelligence in the Public Sector*. Paris: OECD, 2023. https://oecd-opsi.org/work-areas/ai/.

Oxford Insights. *2024 Government AI Readiness Index*. 2025. https://oxfordinsights.com/ai-readiness/ai-readiness-index/.

Oxford Insights. *Government AI Readiness Index 2024*. London: Oxford Insights, 2025.

Patel, Oliver. *AI Governance in 2024: A Year in Review*. 2024. https://oliverpatel.substack.com/p/ai-governance-in-2024-a-year-in-review.

Reducing Bias in AI-Based Financial Services. Washington, DC: Brookings Institution, 2020. https://www.brookings.edu/research/reducing-bias-in-ai-based-financial-services/.

Splunk. *AI Governance in 2025: A Full Perspective on Governance for Artificial Intelligence*. 2025. https://www.splunk.com/en_us/blog/learn/ai-governance.html.

Stanford Institute for Human-Centered Artificial Intelligence, AI Index Steering Committee. Maslej Nestor, Loredana Fattorini, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, et al. *The AI Index 2023 Annual Report*. Stanford University, 2023.

Stanford Institute for Human-Centered Artificial Intelligence. "AI Index 2025: State of AI in 10 Charts." 2025. https://hai.stanford.edu/news/ai-index-2025-state-of-ai-in-10-charts.

Stanford Institute for Human-Centered Artificial Intelligence. *The 2024 AI Index Report*. Stanford University, 2024.

Stanford Institute for Human-Centered Artificial Intelligence. *The 2025 AI Index Report*. Stanford University, 2025.

Survey Findings - Cycle 4 Report: Section 7 - Reporting to the Police. New Zealand Ministry of Justice, 2022. https://www.justice.govt.nz/assets/NZCVS-Cycle-4-Core-Report-Section-7-Reporting-to-the-Police-fin.pdf.

Taylor Fry. Algorithm Charter for Aotearoa New Zealand Year 1 Review 2021. https://www.data.govt.nz/assets/data-ethics/algorithm/Algorithm-Charter-Year-1-Review-FINAL.pdf.

FTI Technology. *Privacy and AI Governance Report 2024*. 2024. https://www.ftitechnology.com/resources/white-papers/2024-privacy-and-ai-governance-report.

The Brookings Institution. *A Comprehensive and Distributed Approach to AI Regulation*. Washington, DC: Brookings Institution Press, 2024.

Threat Report. Toronto: Blackberry Intelligence, 2022.

Tully, Tim, Joff Redfern, Deedy Das, and Derek Xiao. *2025: The State of Generative AI in Enterprise*. Menlo, 2025. https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/.

UK Office for National Statistics. *Mapping the Potential: Generative AI and Public Sector Work*. Alan Turing Institute, 2025. https://www.turing.ac.uk/sites/default/files/2025-05/ons_tus_final_report.pdf.

World Economic Forum. *The Future of Jobs*. Geneva: World Economic Forum, 2023.

Books and book chapters

Ala-Pietilä, Pekka, and Nathalie A. Smuha. "A Framework for Global Cooperation on Artificial Intelligence and Its Governance." In *Reflections on Artificial Intelligence for Humanity*, edited by B. Braunschweig and M. Ghallab. Switzerland: Springer, 2021.

Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press, 2014.

Casonato, Carlo. "AI and Constitutionalism: The Challenges Ahead." In *Reflections on Artificial Intelligence for Humanity*, 127–50. Switzerland: Springer, 2021.

Djeffal, Christian. *The Democratization of Artificial Intelligence: Net Politics in the Era of Learning Algorithms*. Edited by Andreas Sudmann. Columbia: Columbia University Press, 2019.

Finlay, Rebecca, and Hideaki Takeda. "Reflections on Decision-Making and Artificial Intelligence." In *Reflections on Artificial Intelligence for Humanity*, edited by B. Braunschweig and M. Ghallab. Cham: Springer, 2021.

Husain, Amir. *The Sentient Machine: The Coming of Age of Artificial Intelligence*. New York: Scribner, 2017.

Kingwell, Martin. "Sentient AI Persons." In *The Oxford Handbook of Ethics*, 1–19. Oxford: Oxford University Press, 2020.

Meese, James, and Sara Bannerman. "Introduction: Governing the Algorithmic Distribution of the News." In *Governing the Algorithmic Distribution of the News*, edited by James Meese and Sara Bannerman. Cham: Palgrave Macmillan, 2022.

Palmer, Gina Granados. "AI Ethics: Four Key Considerations for a Globally Secure Future." In *Artificial Intelligence and Global Security*, 167–76. Emerald Publishing Limited, 2020.

Russell, Stuart. *Artificial Intelligence: A Binary Approach*. Oxford: Oxford University Press, 2020.

Shanks, Pauline, and Casey Thomas Hart. "Artificial Intelligence and Moral Reasoning: Shifting Moral Responsibility in War?" In *Artificial Intelligence and Global Security*, 121–36. Emerald Publishing Limited, 2020.

Simoncini, Andrea, and Erik Longo. "Fundamental Rights and the Rule of Law in the Algorithmic Society." In *Constitutional Challenges in the Algorithmic Society*, edited by Hans-W. Micklitz et al. Cambridge: Cambridge University Press, 2021.

Tegmark, Max. *Life 3.0: Being Human in the Age of Artificial Intelligence*. New York: Penguin, 2018.

Other resources

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" Paper presented at the Conference on Fairness, Accountability, and Transparency, 2021.

Brkhan, Maja. "Artificial Intelligence and Democracy: The Impact of Disinformation, Social Bots and Political Targeting." *Delphi* 2 (2019): 66–71. https://doi.org/10.21552/delphi/2019/2/4.

Buolamwini, Joy, and Timnit Gebru. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." Paper presented at the Proceedings of the 1st Conference on Fairness, Accountability and Transparency, New York, NY, February 2018.

Chaudhry, Uzma. *AI Governance in Practice*. FTI Consulting, 2024. https://assets.contentstack.io/v3/assets/bltd4dd5b2d705252bc/blt99a18c62ab0b92a7/ai_governance_in_practice_report_2024.pdf. Accessed March 2026.

Dave, Karan. *Understanding AI Governance in 2024: The Stakeholder Landscape*. Washington, DC, 2024.

Delacroix, Sylvie, Joelle Pineau, and Jessica Montgomer. "Democratising the Digital Revolution: The Role of Data Governance." *Artificial Intelligence for Humanity* (2021). https://doi.org/10.1007/978-3-030-69128-8_3.

Freifeld, Karen. "White House Pauses Executive Order That Would Seek to Preempt State Laws on AI, Sources Say." Reuters, November 21, 2025.

Grossman, Gary. “Better Than Humans? AI Barrels Towards AGI.” *VentureBeat*, 2022. https://venturebeat.com/ai/better-than-humans-ai-barrels-towards-agi/.

Hashem, Youmna. “Generative AI Could Free up Valuable Public Sector Time.” Alan Turing Institute, 2025. https://www.turing.ac.uk/news/generative-ai-could-free-valuable-public-sector-time. Accessed March 2025.

Kawakami, Kazuto, et al. “Studying up Public Sector AI: How Networks of Power Relations Shape Agency Decisions.” *arXiv* preprint, 2024. https://arxiv.org/abs/2405.12458.

Kukutai, T., Campbell-Kamariera, K., Mead, A., Mikaere, K., Moses, C., Whitehead, J. & Cormack, D. (2023). *Māori data governance model*. Te Kāhui Raraunga.

Kukutai, T., Cassim, S., Clark, V., Jones, N., Mika, J., Morar, R., Muru-Lanning, M., Pouwhare, R., Teague, V., Tuffery Huria, L., Watts, D. & Sterling, R. (2023). *Māori data sovereignty and privacy. Tikanga in Technology discussion paper*. Hamilton: Te Ngira Institute for Population Research.

Kuziemski, Maciej, and Gianluca Misuraca. “AI Governance in the Public Sector: Three Tales from the Frontiers of Automated Decision-Making in Democratic Settings.” *Global Perspectives* (2020). https://pmc.ncbi.nlm.nih.gov/articles/PMC7164913/.

Leslie, David, Christopher Burr, Mhairi Aitken, Josh Cowls, Mike Katell, and Morgan Briggs. *Artificial Intelligence, Human Rights, Democracy, and the Rule of Law: A Primer*. The Alan Turing Institute, 2021.

Mokander, Jakob, and Ralph Schroeder. “Artificial Intelligence, Rationalization, and the Limits of Control in the Public Sector.” *arXiv* preprint, 2024. https://arxiv.org/abs/2407.05336.

Policy and Society. “Governance of Generative AI.” *Policy & Society* (2025). https://academic.oup.com/policyandsociety/article/44/1/1/7997395.

Pollicino, Oreste, and Giovanni De Gregorio. “Constitutional Challenges in the Algorithmic Society.” (2021). https://doi.org/10.1017/9781108914857.

Reuel, Jonathan, and Espen Undheim. “Generative AI Needs Adaptive Governance.” *arXiv* preprint, 2024. https://arxiv.org/abs/2406.04554.

Schippers, Birgit. “Artificial Intelligence and Democratic Politics.” *Political Insight*, March 2020.

Sehl Mellouli, Marijn Janssen, and Adegboyega Ojo. “Introduction to the Issue on Artificial Intelligence in the Public Sector: Risks and Benefits of AI for Governments.” *Digital Government: Research and Practice* (2024). https://dl.acm.org/doi/10.1145/3636550.

Thomas-Evans, Megan. “A Critical Analysis into the Beneficial and Malicious Utilizations of Artificial Intelligence.” *Artificial Intelligence and National Security* (2022): 81–96. https://doi.org/10.1007/978-3-031-06709-9_5.

Wiseman, Jane. *Using Generative AI in Government: Reflections on Public Service*. Harvard University, 2025. https://janewiseman.scholars.harvard.edu/publications/using-generative-ai-government-reflections-public-service.

Wright, Nicholas. *Artificial Intelligence and Democratic Norms: Meeting the Authoritarian Challenge*. Washington, DC: Think Democracy, 2020.

Engin, Zeynep, Jon Crowcroft, David Hand, and Philip Treleaven. “The Algorithmic State Architecture (ASA): An Integrated Framework for AI-Enabled Government.” *arXiv* preprint, 2025. https://arxiv.org/abs/2503.08725.

ENDNOTES

1. Marijn Janssen, Sehl Mellouli, and Adegboyega Ojo, "Introduction to the Issue on Artificial Intelligence in the Public Sector: Risks and Benefits of AI for Governments," Digital Government: Research and Practice (2024): 1, https://dl.acm.org/doi/10.1145/3636550.
2. Tim Tully et al., *2025: The State of Generative AI in Enterprise*, Menlo (December 9, 2025), 155, https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/.
3. Sehl Mellouli, "Artificial Intelligence in the Public Sector," 1; Chip Cutter, "Here's Where AI Is Tearing Through Corporate America," *WSJ* 2026, https://www.wsj.com/tech/ai/ai-software-business-stock-market-4b17b432?mod=saved_content. Accessed March 2026.
4. Organisation for Economic Co-operation and Development (OECD), *Governing with Artificial Intelligence: Are Governments Ready?*, Organisation for Economic Co-operation and Development (2024), 4, https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/06/governing-with-artificial-intelligence_f0e316f5/26324bc2-en.pdf. Accessed March 2026.
5. Organisation for Economic Co-operation and Development (OECD), *Governing with Artificial Intelligence: Are Governments Ready?*, 5.
6. C. Van Noordt and G. Misuraca, "Artificial intelligence for the Public Sector: Results of Landscaping the Use of AI in Government Across the European Union," *Government Information Quarterly* 39, no. 6 (2022): 1.
7. Such as Robodebt in Australia. Royal Commission, *Royal Commission into the Robodebt Scheme* (2023), https://robodebt.royalcommission.gov.au/. Accessed March 2026.
8. AI Center of Excellence United States. General Services Administration, *AI Guide for Government*, General Services Administration (2024), 7, https://coe.gsa.gov/coe/ai-guide-for-government/print-all/index.html. Accessed March 2026.
9. United States. General Services Administration, *AI Guide*, 7.
10. Zeynep Engin et al., "The Algorithmic State Architecture (ASA): An Integrated Framework for AI-Enabled Government," *arXiv preprint* (2025): 9-10, https://arxiv.org/abs/2503.08725.
11. The White House. Executive Order, *Ensuring a National Policy Framework for Artificial Intelligence* (2025). https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy. Accessed March 2026.
12. UK Government, *AI to Accelerate National Renewal and Growth as Google DeepMind Backs UK Tech and Science Sectors* (2025). https://www.gov.uk/government/news/ai-to-accelerate-national-renewal-and-growth-as-google-deepmind-backs-uk-tech-and-science-sectors. Accessed March 2026.
13. Government of Canada. Treasury Board, *Register of AI Uses in Federal Government* (2025). https://www.canada.ca/en/treasury-board-secretariat/news/2025/11/canada-launches-first-register-of-ai-uses-in-federal-government.html. Accessed March 2026.
14. Australian Government. Australian Public Services Academy, *Exploring AI Innovations: Highlights from the GovAI AI Use Case Library* (2025). https://www.apsacademy.gov.au/news/exploring-ai-innovations-highlights-govai-ai-use-case-library. Accessed March 2026; Australian Government, *APS AI Use Case Library* (2026). https://www.govai.gov.au/explore/use-cases. Accessed March 2026.
15. Monetary Authority of Singapore, "Consultation Paper on Guidelines on Artificial Intelligence Risk Management" (2025). https://www.mas.gov.sg/-/media/mas-media-library/publications/consultations/bd/2025/final_consultation_paper_on_guidelines_on_ai_risk_management_forrelease.pdf. Accessed March 2026.
16. New Zealand Government, *Wallet App* (2025). https://www.govt.nz/about/the-govt-nz-app. Accessed March 2026.
17. Organisation for Economic Co-operation and Development (OECD), *AI in Government Report 2024*, Organisation for Economic Co-operation and Development (Paris, 2024), 5, https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en.html. Accessed March 2026.
18. Kira Matus, Michael Veale, and Robert Gorwa, "AI and Global Governance: Modalities, Rationales, Tensions," *Annual Review of Law and Social Science* (2024): 255, https://www.annualreviews.org/content/journals/10.1146/annurev-lawsocsci-020223-040749.
19. Philip Treleaven and Zeynep Engin, "Algorithmic Government: Automating Public Services and Supporting Civil Servants in Using Data Science Technologies," *The Computer Journal* 62, no. 3 (2019): 12, https://academic.oup.com/comjnl/article/62/3/448/5070384.
20. Zeynep Engin et al., "The Algorithmic State Architecture," 3.
21. These arise when political, reputational, or coercive pressures prompt governments to embrace AI initiatives primarily for signalling modernity or alignment with global trends, rather than achieving deep integration. Executive mandates, funding tied to innovation metrics or international benchmarks—such as OECD AI Principles—can drive agencies to launch pilots, publish strategies or procure tools with fanfare, yet without commensurate investment in skills, data infrastructure or cultural change.
22. Michael Veale, "AI and Global Governance: Modalities, Rationales, Tensions," 255.
23. "How to Build Digital Capability," 2025, https://dns.govt.nz/standards-and-guidance/strategy-and-planning/digital-capability-public-serviceworkforce/how-to-build-digital-capability. Accessed March 2026.
24. "Te Mana Raraunga," 2018, https://www.temanararaunga.maori.nz/tutohinga. Accessed March 2026. See also, Kukutai, T., Campbell-Kamariera, K., Mead, A., Mikaere, K., Moses, C., Whitehead, J. & Cormack, D. (2023). *Māori data governance model*. Te Kāhui Raraunga.
25. Stephanie Russo Carroll et al., "The CARE Principles for Indigenous Data Governance," *Data Science Journal* 19 (2020), https://doi.org/10.5334/dsj-2020-043.

26. Sehl Mellouli, "Artificial Intelligence in the Public Sector," 4.
27. Patrick Mikalef, Emmanouil Papagiannidis, and Kieran Conboy, "Responsible Artificial Intelligence Governance: A Review and Research Framework," *Government Information Quarterly* (2024): 6, https://www.sciencedirect.com/science/article/pii/S0963868724000672.
28. Emmanouil Papagiannidis, "Responsible Artificial Intelligence Governance," 6.
29. Araz Taeihagh, "Governance of Generative AI," *Policy & Society* 44, no. 1 (2025): 6, https://academic.oup.com/policyandsociety/article/44/1/1/7997395.
30. Kazuto Kawakami et al., "Studying Up Public Sector AI: How Networks of Power Relations Shape Agency Decisions," *arXiv preprint* (2024): 4, https://arxiv.org/abs/2405.12458.
31. Kawakami et al., "Studying Up Public Sector AI: How Networks of Power Relations Shape Agency Decisions," 4.
32. And note Stanford's analysis of "Limitations and Risks of Current AI Systems", Fei-Fei Li, Christopher Manning, and Anka Reuel, *The Stanford Emerging Technology Review: A Report on Ten Key Technologies and Their Policy Implications*, Stanford University (2026), 28.
33. Kawakami et al., "Studying Up Public Sector AI: How Networks of Power Relations Shape Agency Decisions," 4.
34. Kawakami et al., "Studying Up Public Sector AI," 17.
35. Kawakami et al., "Studying Up Public Sector AI," 17.
36. Madalina Busuioc and Saar Alon-Barkat, "Human–AI Interactions in Public Sector Decision Making: 'Automation Bias' and 'Selective Adherence' to Algorithmic Advice," *Journal of Public Administration Research and Theory* 33, no. 1 (2022): 155, https://academic.oup.com/jpart/article/33/1/153/6524536.
37. Saar Alon-Barkat, "Human–AI Interactions in Public Sector Decision Making," 154.
38. Emmanouil Papagiannidis, «Responsible Artificial Intelligence Governance,» 7.
39. Illugi Torfason Hjaltalin and Hallur Thor Sigurdarson, "The Strategic Use of AI in the Public Sector: A Public Values Analysis of National AI Strategies," 1, https://www.sciencedirect.com/science/article/abs/pii/S0740624X24000066.
40. Organisation for Economic Co-operation and Development (OECD), *AI in Government Report 2024*, 2.
41. Organisation for Economic Co-operation and Development (OECD), *AI in Government Report 2024*, 5.
42. Organisation for Economic Co-operation and Development (OECD), *Governing with Artificial Intelligence: Are Governments Ready?*, 5.
43. United States. General Services Administration, *AI Guide*, 7.
44. United States. General Services Administration, *AI Guide*, 7.
45. Heidi Kim et al., *How AI Can Cut Through Bureaucracy, Boost Efficiency, and Build Trust in Government*, BCG, 2, https://www.bcg.com/publications/2025/benefits-of-ai-in-government. Accessed March 2026.
46. Heidi Kim et al., *How AI Can Cut Through Bureaucracy*, 2.
47. Heidi Kim et al., *How AI Can Cut Through Bureaucracy*, 4.
48. Youmna Hashem, *Generative AI Could Free Up Valuable Public Sector Time*, Alan Turing Institute (2025), 6, https://www.turing.ac.uk/news/generative-ai-could-free-valuable-public-sector-time. Accessed March 2025.
49. Hashem, *Generative AI*, 16.
50. Zeynep Engin et al., "The Algorithmic State Architecture," 2.
51. Zeynep Engin et al., "The Algorithmic State Architecture," 3.
52. Zeynep Engin et al., "The Algorithmic State Architecture," 5.
53. United States. General Services Administration, *AI Guide*, 21.
54. United States. General Services Administration, *AI Guide*, 22.
55. United States. General Services Administration, *AI Guide*, 22.
56. Organisation for Economic Co-operation and Development (OECD), *AI in Government Report 2024*, 23.
57. UK Government, *AI to Accelerate National Renewal and Growth*.
58. Government of Canada. Treasury Board, *Register of AI Uses in Federal Government*.
59. Australian Government, *APS AI Use Case Library*.
60. New Zealand Government, *Wallet App*.
61. Executive Office of the President and Office of Management and Budget, "Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles," (December 11, 2025), https://www.whitehouse.gov/wp-content/uploads/2025/12/M-26-04-Increasing-Public-Trust-in-Artificial-Intelligence-Through-Unbiased-AI-Principles-1.pdf. Accessed March 2026.
62. FedScoop, *OMB Lays Out Requirements for Agencies* (2025), https://fedscoop.com/omb-requirements-woke-ai-federal-agencies. Accessed March 2026.
63. Sehl Mellouli, "Artificial Intelligence in the Public Sector," 4.
64. Edmund Burke, *Reflections on the Revolution in France*, ed. Conor Cruise O'Brien (London: The Folio Society, 2010), §39.

65. Matti Minkkinen, Teemu Birkstedt, Anushree Tandon, and Matti Mäntymäki, "AI governance: themes, knowledge gaps and future agendas," *Internet Research* (2022): 3, https://doi.org/10.1108/intr-01-2022-0042, https://www.emerald.com/insight/content/doi/10.1108/intr-01-2022-0042/full/html.
66. Zeynep Engin et al., "The Algorithmic State Architecture," 15.
67. Zeynep Engin et al., "The Algorithmic State Architecture," 15.
68. Zeynep Engin et al., "The Algorithmic State Architecture," 20.
69. Zeynep Engin et al., "The Algorithmic State Architecture," 24.
70. Saar Alon-Barkat, "Human–AI Interactions in Public Sector Decision Making," 154.
71. Saar Alon-Barkat, "Human–AI Interactions in Public Sector Decision Making," 154.
72. United States. General Services Administration, *AI Guide*, 77.
73. Uzma Chaudhry, *AI Governance in Practice*, FTI Consulting (2024), 60, https://assets.contentstack.io/v3/assets/bltd4dd5b2d705252bc/blt99a18c62ab0b92a7/ai_governance_in_practice_report_2024.pdf. Accessed March 2026.
74. And note the risk of drift in relation to reliability. Fei-Fei Li, Christopher Manning, and Anka Reuel, *The Stanford Emerging Technology Review*, 30.
75. Taeihagh, "Governance of Generative AI," 7.
76. Gianluca Misuraca and Colin van Noordt, "Artificial Intelligence For The Public Sector: Results Of Landscaping The Use Of AI In Government Across The European Union," *Government Information Quarterly* (2022): 4, https://www.sciencedirect.com/science/article/abs/pii/S0740624X22000478.
77. Organisation for Economic Co-operation and Development (OECD), *AI in Government Report 2024*, 2; United States. General Services Administration, *AI Guide*, 7.
78. Zeynep Engin, "Algorithmic Government," 13-14.
79. Saar Alon-Barkat, "Human–AI Interactions in Public Sector Decision Making," 154.
80. Saar Alon-Barkat, "Human–AI Interactions in Public Sector Decision Making," 154.
81. Kawakami et al., "Studying Up Public Sector AI: How Networks of Power Relations Shape Agency Decisions," 14.
82. UK Government, *AI to Accelerate National Renewal and Growth*.
83. Government of Canada. Treasury Board, *Register of AI Uses in Federal Government*.
84. Jakob Mokander and Ralph Schroeder, "Artificial Intelligence, Rationalization, and the Limits of Control in the Public Sector," *arXiv preprint* (2024): 12, https://arxiv.org/abs/2407.05336.
85. Mokander and Schroeder, "Artificial Intelligence, Rationalization, and the Limits of Control in the Public Sector," 15.
86. With the apparent imminent emergence of an artificial general intelligence, it might be worth pausing to consider the significance of such reporting: Sam Schechner, "When AI Bots Start Bullying Humans, Even Silicon Valley Gets Rattled," *The Wall Street Journal*, 13 February 2026, https://www.wsj.com/tech/ai/when-ai-bots-start-bullying-humans-even-silicon-valley-gets-rattled-0adb04f1?mod=saved_content. Accessed March 2026.
87. Gianluca Misuraca and Colin van Noordt, "Artificial Intelligence For The Public Sector," 1.
88. Govt.nz app," 2025, https://www.govt.nz/about/the-govt-nz-app/. Accessed March 2026. GDDA 2026 *Digital Target State*, *Dns.govt.nz*. Available at: https://dns.govt.nz/digital-government/digitising-government-programme-dgp/digital-target-state (Accessed: 19 April 2026).
89. Te Kāwanatanga o Aotearoa and New Zealand Government, *Public Service AI Framework* (2026), 1, https://docref.digital.govt.nz/nz/generative-ai-guidance-gcdo/public-service-ai-framework/2025/en/. Accessed March 2026.
90. GovTech Singapore, *A Co-pilot, Not a Replacement: How AI Will Augment Human Labour In The Workforce* (2025), https://www.tech.gov.sg/technews/how-ai-will-augment-human-labour-in-the-workforce. Accessed March 2026.
91. Singapore, "Consultation Paper on Guidelines on Artificial Intelligence Risk Management."
92. UK Government, *AI to Accelerate National Renewal and Growth*.
93. UK Government and Government Digital and Data, *Join AI Engineering Lab And Help Shape The Future Of Digital Government* (January 12, 2026), https://cddo.blog.gov.uk/2026/01/12/join-ai-engineering-lab-and-help-shape-the-future-of-digital-government/. Accessed March 2026.
94. UK Government, *AI to Accelerate National Renewal and Growth*.
95. Australian Government. Australian Public Services Academy, *Exploring AI Innovations: Highlights from the GovAI AI Use Case Library*.
96. AIDE Artificial Intelligence Delivery and Enablement and Australian Public Service Commission, *Artificial Intelligence Delivery and Enablement* (2026), https://www.govai.gov.au/aide. Accessed March 2026.
97. Government of Canada. Treasury Board, *Register of AI Uses in Federal Government*.
98. Government of Canada, *Government of Canada AI Register: Progress on AI in government* (2026), https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/progress.html. Accessed March 2026.
99. Government of Canada. Treasury Board, *Register of AI Uses in Federal Government*.

100. The White House. Executive Order, *Ensuring a National Policy Framework for Artificial Intelligence*.

101. Gary Marchant and Jordan Buckwald, "Private Standards as Liability Shields: A Pro-Innovation Artificial Intelligence Regulatory Approach for States," *Minnesota Journal of Law, Science & Technology* 27, no. 1 (2026), https://scholarship.law.umn.edu/mjlst/vol27/iss1/. Accessed March 2026; Gary Marchant and Jordan Buckwald, "Private Standards as Liability Shields: A Pro-Innovation Artificial Intelligence Regulatory Approach for States."; S. M. Seaman and P. E. Hernandez, "D&O Liability & Coverage: 2025 Trends, Developments & Decisions," *Mealey's Litigation Report: Artificial Intelligence* 3, no. 4 (2025), https://www.hinshawlaw.com/newsroom-publications-mealeys-report-dando-liability-and-coverage-2025-trends-developments-and-decisions.html. Accessed March 2025; Karen Freifeld, *White House Pauses Executive Order That Would Seek To Preempt State Laws On AI, Sources Say*, Reuters (November 21, 2025).

102. FedScoop, *OMB Lays Out Requirements for Agencies*.

103. Keach Hagey et al., *'Woke' AI Feud Escalates Between Pentagon and Anthropic*, Wall Street Journal (February 18, 2026), https://www.wsj.com/politics/national-security/woke-ai-spat-escalates-between-pentagon-and-anthropic-433b7c5c?mod=hp_lead_pos1. Accessed March 2026.

104. Especially in light of the fact that "AI is expected to have a profound impact on militaries worldwide. Weapons systems, command and control, logistics, acquisition, and training will all seek to leverage AI to operate more effectively and efficiently, at lower cost and with less risk to friendly forces. Trying to overcome decades of institutional inertia, the US Department of Defense is dedicating billions of dollars to institutional reforms and research advances aimed at integrating AI into its warfighting and war preparation strategies. Senior military officials are concerned that failure to adapt to the emerging opportunities and challenges presented by AI would pose significant national security risks, particularly considering that both Russia and China are investing heavily in AI capabilities." Fei-Fei Li, Christopher Manning, and Anka Reuel, *The Stanford Emerging Technology Review*, 33.

105. Thus, "In the United States, the Trump administration revoked the Biden administration's Executive Order (EO) 14110, Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. It subsequently issued a new EO titled Removing Barriers to American Leadership in Artificial Intelligence, aiming to promote AI innovation and leadership by eliminating the restrictions and requirements on AI contained in EO 14110. In August 2025, the administration launched America's AI Action Plan, which outlines a policy road map to "accelerate innovation, build American AI infrastructure, and lead in international diplomacy and security." "This plan faces challenges, especially regarding alignment with concurrent proposals to scale back broader scientific research funding." And, "In the European Union, the most ambitious attempts to regulate AI came into force in August 2024 with the EU Artificial Intelligence Act. The act forbids certain applications of AI, such as individual predictive policing based solely on a person's data profile or tracking of their emotional state in the workplace and educational institutions, unless for medical or safety reasons. Additionally, it imposes a number of requirements on high-risk AI systems and foundation models, addressing transparency and explainability, human oversight, cybersecurity, and robustness." Fei-Fei Li, Christopher Manning, and Anka Reuel, *The Stanford Emerging Technology Review*, 31.

106. Gianluca Misuraca and Colin van Noordt, "Artificial Intelligence For The Public Sector," 1; Hashem, *Generative AI*, 6.

107. Organisation for Economic Co-operation and Development (OECD), Governing with Artificial Intelligence: Are Governments Ready?, 2.

108. Organisation for Economic Co-operation and Development (OECD), Governing with Artificial Intelligence: Are Governments Ready?, 23.

109. Zeynep Engin et al., "The Algorithmic State Architecture," 1.

110. Zeynep Engin et al., "The Algorithmic State Architecture," 12.

111. Heidi Kim et al., *How AI Can Cut Through Bureaucracy*, 1.

112. Heidi Kim et al., *How AI Can Cut Through Bureaucracy*, 2.

113. Heidi Kim et al., *How AI Can Cut Through Bureaucracy*, 4.

114. Hashem, *Generative AI*, 16.

115. Government of Canada, *AI Register* (2025), https://open.canada.ca/data/en/dataset/fcbc0200-79ba-4fa4-94a6-00e32facea6b. Accessed March 2026; Government of Canada, *Government of Canada AI Register: Progress on AI in government*; Government of Canada. Treasury Board, *Register of AI Uses in Federal Government*.

116. Zeynep Engin et al., "The Algorithmic State Architecture," 9; Zeynep Engin, "Algorithmic Government."

117. "Human–AI Interactions in Public Sector Decision Making," 155.

118. FTI Technology, *Privacy and AI Governance Report 2024*, FTI Technology (2024), 60, https://www.ftitechnology.com/resources/white-papers/2024-privacy-and-ai-governance-report.

119. Singapore, "Consultation Paper on Guidelines on Artificial Intelligence Risk Management."

120. United States. General Services Administration, *AI Guide*, 22.

121. Saar Alon-Barkat, "Human–AI Interactions in Public Sector Decision Making," 155.

122. Kawakami et al., "Studying Up Public Sector AI: How Networks of Power Relations Shape Agency Decisions," 4.

123. New Zealand Government and Department of Internal Affairs, "Digital Capability Framework 2025: Upholding Māori Rights and Interests in Data," (2025), https://www.digital.govt.nz/. Accessed March 2026.

124. Hashem, *Generative AI*, 6; Organisation for Economic Co-operation and Development (OECD), *Governing with Artificial Intelligence: Are Governments Ready?*, 2.

125. And note: "The intensifying technological race between the United States and China regarding AI is entering a new phase. While the United States continues to push the technical frontier with increasingly capable models, China is aggressively diffusing existing AI capabilities across every sector from education to manufacturing to governance—aiming to lock in large-scale network advantages at home and abroad. China's open-source model releases, such as DeepSeek, further challenge America's frontier status, accelerate global adoption, and undermine US containment efforts." Fei-Fei Li, Christopher Manning, and Anka Reuel, *The Stanford Emerging Technology Review*, 32.

126. Note: "The global AI market is projected to be worth $244.22 billion in 2025, with North America receiving 33.8% of total AI revenues.[2] The Stanford Institute for Human-Centred Artificial Intelligence (HAI) AI Index 2025 Annual Report found that private investment in all AI start-ups totaled $150.79 billion in 2024, surpassing the previous record high of over $120 billion, in 2021, after two consecutive years of decline. One estimate forecasts that generative AI—which can create novel text, images, video, and audio output and is discussed in more detail later in this chapter—could raise global GDP by $7 trillion and raise productivity growth by 1.5% over a 10-year period if it is adopted widely." Fei-Fei Li, Christopher Manning, and Anka Reuel, *The Stanford Emerging Technology Review*, 24.

127. Technology, *Privacy and AI Governance Report 2024*, 60.

MAXIM
INSTITUTE

PO Box 49 074, Roskill South, Auckland 1445, New Zealand / 49 Cape Horn Road, Hillsborough, Auckland / Ph +64 9 627 3261 / **www.maxim.org.nz**