Beyond the Black Box

Building Transparent and Accountable AI in Government

maximuscanada.ca

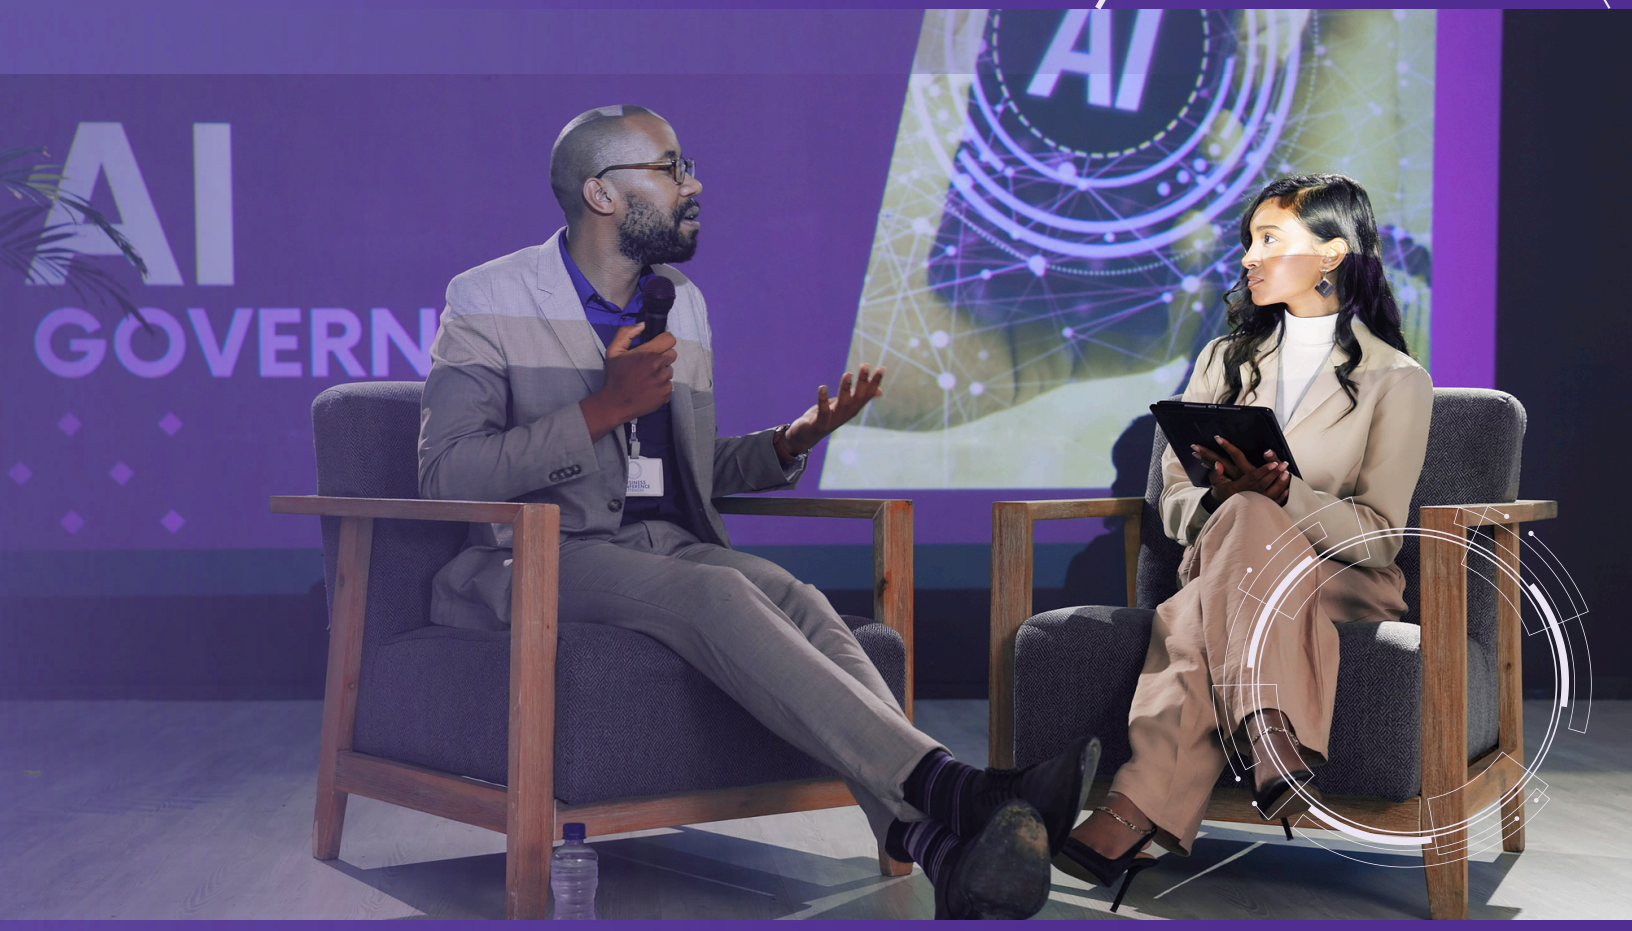

maximus

Moving Canadians, Technology, and Government Forward

maximus

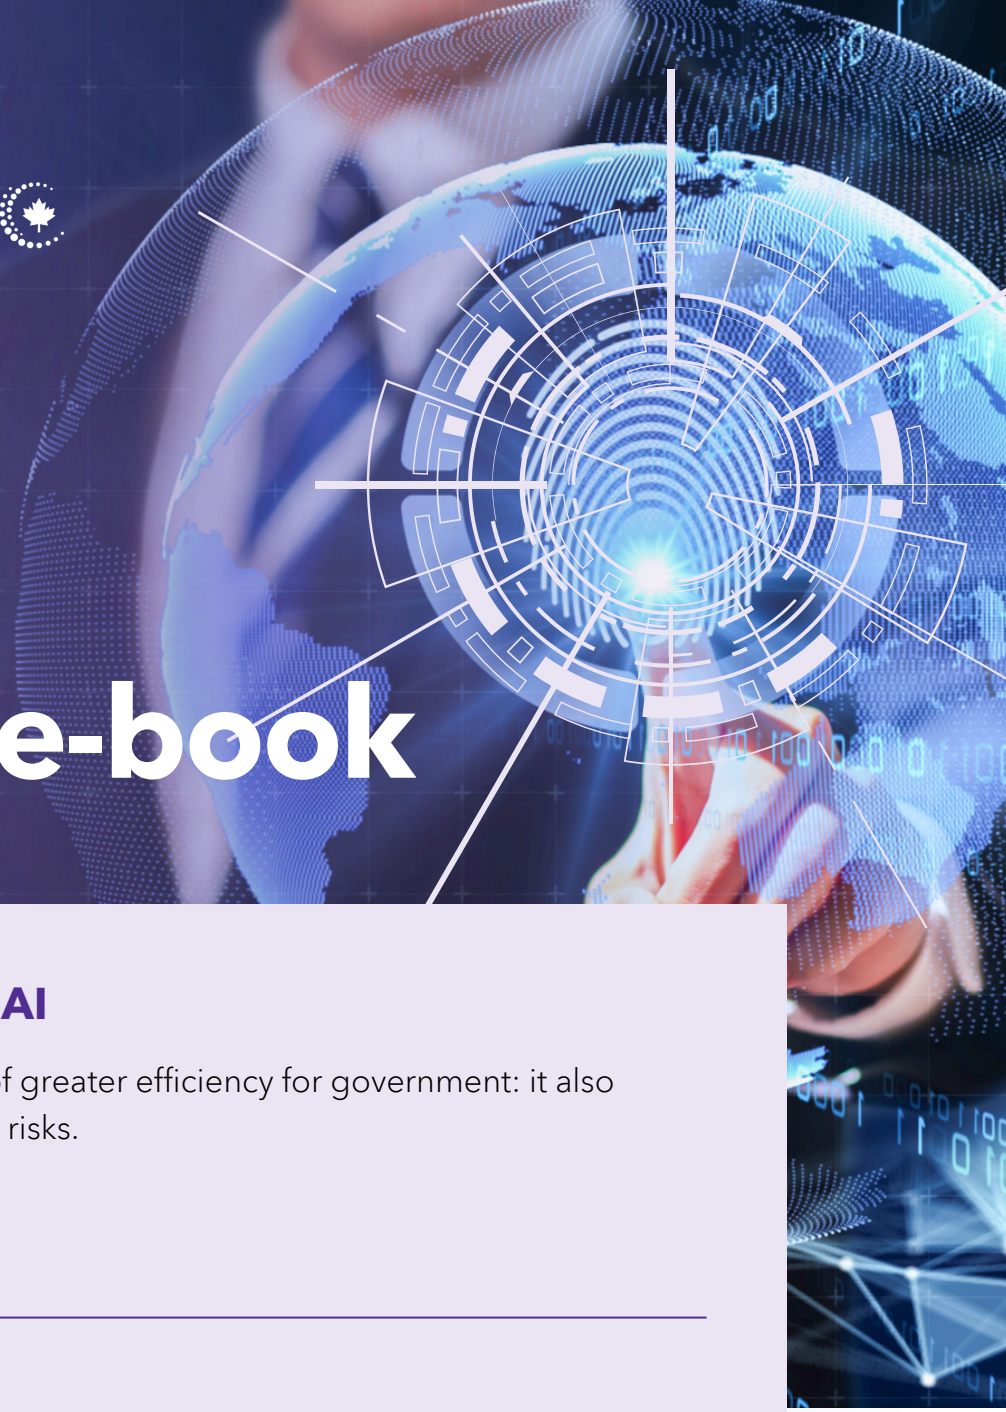

In this e-book

The Promise of AI

AI holds the promise of greater efficiency for government: it also introduces new ethical risks.

Let's Explore In-depth

This e-book will help you mitigate moral hazards in AI, and transform it into a transparent, accountable tool governed in the public interest.

Introduction

Artificial Intelligence is rapidly transforming the way governments design, deliver, and evaluate public services. From automating benefits processing to supporting decision-making and enhancing citizen interactions through web portals, AI is being embedded in the core infrastructure of modern public administration. This evolution holds the promise of greater efficiency, responsiveness, and scalability; however, it also introduces complex ethical and governance risks that are not yet fully understood or effectively managed.

Among these emerging risks is the phenomenon of Moral Hazard, a situation in which those making decisions do not bear the full consequences of their actions, leading to distorted incentives and, potentially, harmful outcomes. In the context of government use of AI, moral hazard can manifest when public agencies, developers, or vendors deploy automated systems that affect people's lives while avoiding responsibility for errors, biases, or unintended consequences. Because AI decisions are often obscure, data-driven, and seen as "neutral," there is a risk that accountability may become diluted or fully denied.

We identify structural conditions and use cases where moral hazard is most likely to arise in public-sector AI deployments, with a focus on high-impact domains such as public insurance, healthcare, program delivery, and citizen access to services. These areas are where automation is increasingly influencing who receives help, how decisions are made, and what recourse is available when things go wrong.

By examining where AI is being used, we aim to highlight the need for stronger safeguards, clearer accountability frameworks, and a more ethical approach to AI governance in the public sector. Governments must ensure that the benefits of AI are not pursued at the expense of democratic values, public trust, or human dignity, and strategically address the tenets of responsible AI use.

Understanding Moral Hazard in the AI Context

Moral hazard is a concept originating in insurance and economics, describing situations where individuals or institutions engage in riskier behaviour because they do not bear the full consequences of that risk. When protections or incentives are structured in a way that weakens or removes accountability, the likelihood of harmful or negligent behaviour increases.

In traditional contexts, moral hazard is often discussed in relation to financial bailouts, insurance policies, or principal-agent relationships. For example, a driver with full-coverage insurance might be less cautious, knowing that damages are financially covered. Similarly, a financial institution expecting a government bailout may engage in high-risk lending, assuming the state will absorb the fallout.

In the context of Artificial Intelligence, moral hazard takes on new dimensions, particularly because AI introduces layers of automation and complexity that further obscure lines of accountability. When governments adopt AI systems to make or support decisions about people's eligibility for services, allocation of resources, or even risk profiling, the consequences of those decisions may be felt by citizens, not the institutions, developers, or decision-makers deploying the systems.

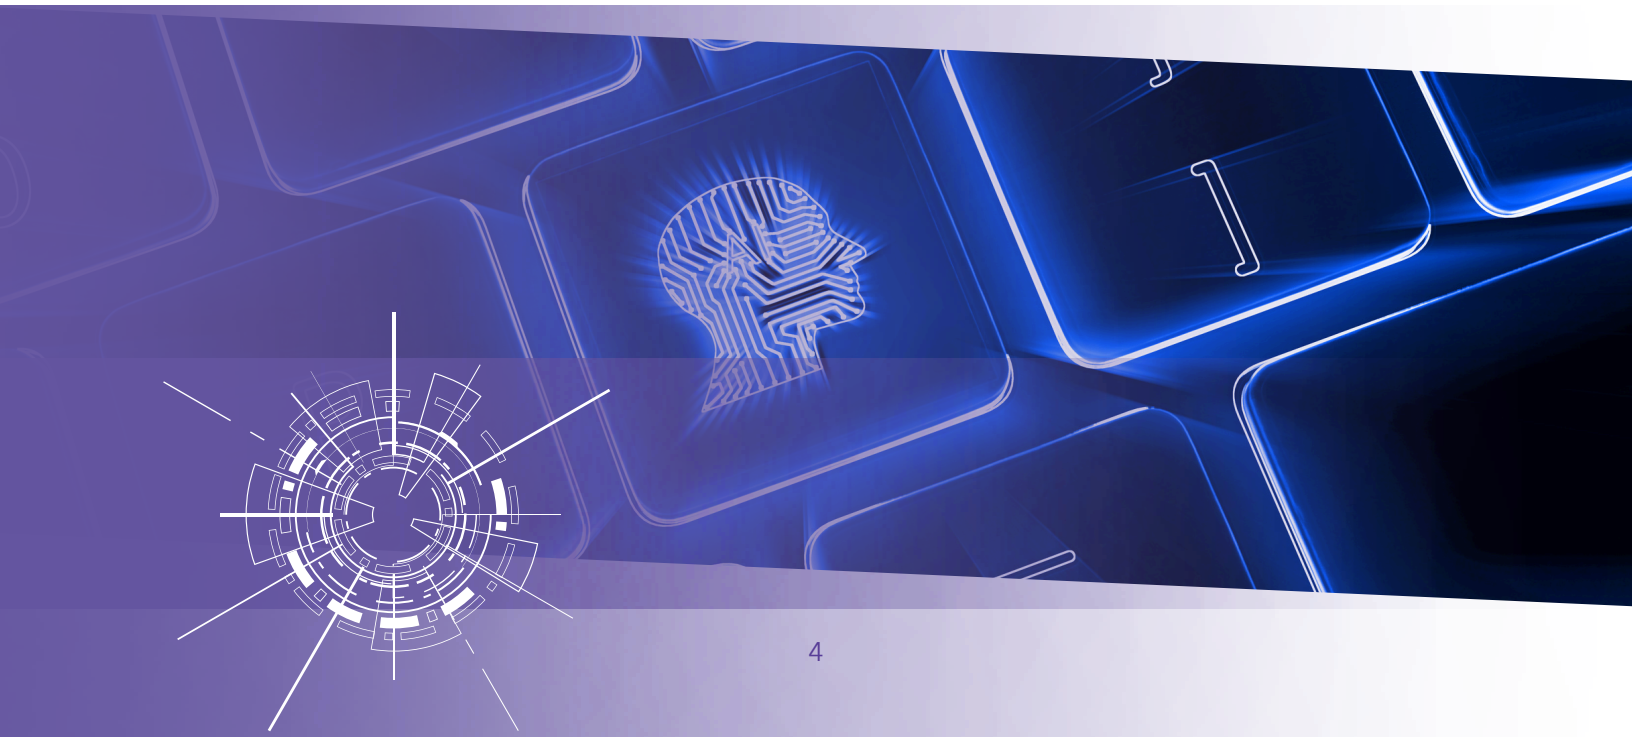

AI exacerbates moral hazard in several ways:

- Opacity of Decision-Making: Many AI systems, especially those using machine learning, function as “black boxes” — their decision logic is not easily interpretable, even by the developers themselves. This makes it difficult to assign responsibility when harm occurs.
- Delegation of Authority: When frontline workers or program administrators rely on AI outputs to guide or even make decisions, there is a tendency to defer to the system rather than exercise professional judgment. This delegation can reduce critical oversight and due diligence.
- Diffused Accountability Across Vendors and Agencies: Governments often procure AI systems from third-party vendors, creating a complex chain of responsibility. If the system causes harm, it may be unclear whether the blame lies with the vendor, the government agency, or the individual end-user.
- Incentive Misalignment: Vendors may be incentivized to deliver performance metrics (e.g., fraud detection rates, cost reductions) without equal concern for fairness, inclusion, or accuracy. Governments, under pressure to reduce costs or modernize services quickly, may accept these trade-offs without sufficient scrutiny.
- Knowledge Management: When knowledge management governance is weak and redundant, outdated, and trivial (ROT). If ROT is left unchecked, this can result in organizations' risk of polluting their knowledge ecosystems with low-quality information. This digital clutter not only degrades searchability and compliance but also directly undermines the reliability of AI systems and human decision-making. AI agents and knowledge consumers alike depend on the integrity of the underlying content to generate accurate, contextually appropriate responses.

Importantly, moral hazard in public-sector AI doesn't necessarily stem from malice or neglect. It often arises from institutional structures, procurement models, and decision-making cultures that do not adequately anticipate or manage the downstream effects of automation. Left unaddressed, these dynamics can result in real harms, particularly for vulnerable populations who may not have the means or knowledge to contest algorithmic decisions.

Understanding and identifying moral hazard in the AI lifecycle is crucial for developing governance structures that prioritize human dignity, uphold accountability, and maintain democratic legitimacy in public services.

Government Use Cases at Risk of AI Moral Hazard

The integration of AI into government functions holds the promise of more efficient, data-driven public services. However, the introduction of automation into high-stakes public systems also introduces significant risk, especially when decisions are made or influenced by systems whose developers, vendors, or administrators are not fully accountable for the consequences. Nowhere is this more evident than in key areas such as public insurance, healthcare, social program delivery, and access to government services, where moral hazard can quietly emerge beneath the surface of efficiency.

In public insurance systems, such as those managing employment insurance, workers' compensation, and publicly funded auto insurance, AI is frequently deployed to expedite claim adjudication and detect fraudulent patterns. While this may improve operational efficiency, it introduces the risk of automated decision-making processes wrongly denying legitimate claims. AI tools used to flag anomalies or high-risk applications are often built on historical datasets that may embed institutional biases or statistical assumptions that do not reflect current realities. Because vendors and procurement teams are typically focused on minimizing fraud or increasing throughput, there is a tendency to underinvest in validating the equity or fairness of these systems.

As a result, wrongful denials may occur with limited avenues for redress, and those designing or implementing the AI face little to no consequence for errors. The burden instead falls on claimants, often in vulnerable positions, to navigate complex appeal systems while trust in the legitimacy of public insurance erodes.

In the context of public healthcare, AI tools are now being used to guide diagnosis, triage patient populations, and support administrative decision-making. Governments and health authorities utilize predictive models to identify patients at risk of specific conditions or to allocate treatment resources more effectively. While these innovations are touted as transformative, their deployment introduces subtle forms of moral hazard. Health professionals may become overly dependent on AI-generated recommendations, particularly in under-resourced systems, while third-party developers who create these tools may not bear responsibility when those recommendations prove inaccurate.

The opacity of many machine learning systems complicates oversight; errors or biases may remain hidden until they result in measurable harm. Moreover, if these tools are not adequately tested on diverse populations, their outcomes may disproportionately disadvantage patients who are already experiencing systemic inequities in healthcare.

Government program delivery, especially those tied to social assistance, disability supports, or housing benefits, is another area where AI is being used to assess eligibility, prioritize cases, and streamline application reviews. In these systems, moral hazard emerges when algorithms are treated as objective arbiters, while the structural biases embedded within them go unchecked. For example, risk-scoring systems intended to identify fraudulent or “high-need” applicants may instead reinforce discriminatory patterns, excluding qualified individuals based on flawed correlations in the training data. In practice, frontline workers often defer to these algorithmic scores, particularly when under pressure to reduce workloads or meet policy targets. As vendors are incentivized to meet performance benchmarks, like throughput or cost savings, without being held accountable for social harms, individuals may be denied services critical to their well-being with little understanding of how or why the decision was made.

Even at the interface level, AI plays a growing role in shaping how citizens access government services. Chatbots, recommendation engines, and virtual assistants are being deployed to help people navigate complex systems such as immigration processes, tax inquiries, and health portals. While often marketed as tools for expanding access, these AI systems carry their own form of moral hazard. Agencies may overestimate the capacity of AI to replace human service agents, leading to the withdrawal of traditional support channels. The companies that design these interfaces may not be required to address failures experienced by users with disabilities, low literacy, or limited digital access. When misinformation or routing errors occur, such as being misdirected by a chatbot or receiving incomplete instructions, the human cost is borne by the user, not by the system designers or public officials who approved the technology.

Across all these use cases, the pattern is clear: AI systems shift decision-making power away from people while also blurring lines of accountability. Government actors, under pressure to modernize, may defer ethical questions to technical experts or vendors if automation is inherently believed to produce better outcomes. Vendors, in turn, may focus on optimizing for efficiency metrics rather than social impact. The result is a systemic moral hazard in which those with the power to shape outcomes are not the ones who bear the costs when things go wrong. Identifying these dynamics early is critical to building AI governance structures that ensure both technical reliability and ethical legitimacy.

Structural Drivers of AI Moral Hazard in Government

The emergence of moral hazard in government use of artificial intelligence is not simply the result of poor system design or individual negligence. Rather, it is deeply rooted in structural dynamics that shape how public institutions adopt, procure, and govern AI systems. These dynamics often reinforce the detachment of decision-making authority from responsibility, allowing harm to occur without clear accountability. Understanding these drivers is essential for addressing the systemic nature of AI-related moral hazard.

A key structural driver is the asymmetry of technical knowledge between AI vendors and public agencies. Most government departments lack the in-house expertise necessary to evaluate, audit, or challenge the technical assumptions embedded in AI systems. As a result, they become dependent on external vendors, whose priorities may be shaped more by commercial incentives than public interest. This dependence creates a risk environment in which systems are purchased or deployed without sufficient scrutiny of their social, legal, and ethical implications.

Compounding this challenge is the opacity of public procurement processes. AI tools are often introduced through general-purpose contracts or innovation initiatives that lack transparent evaluation criteria or impact assessments. When systems are acquired through opaque channels or under innovation exemptions, considerations for the

responsible use of AI are not addressed, making it difficult for civil society, regulators, or even internal oversight bodies to examine their design, data, or performance. Procurement officers may prioritize cost savings, speed, or digital modernization over long-term impacts on equity or rights, especially when there are no clear metrics to assess social risk and the responsible use of AI.

Another driver is the fragmentation of accountability across actors and institutions. In many cases, the chain of responsibility for decisions made with or by AI systems spans multiple entities, including a public agency, one or more private vendors, and possibly external consultants or cloud service providers. This diffusion makes it difficult to determine who is liable when a system produces harmful outcomes. For instance, if an automated decision tool used in welfare eligibility causes wrongful denial, the agency may blame the algorithm, the vendor may blame the input data, and frontline staff may claim they were simply following system guidance. The result is a governance vacuum in which affected individuals face enormous barriers to seeking redress.

Public sector AI is also shaped by performance incentives that may undervalue ethical or social outcomes. Government departments are often rewarded for increasing efficiency, reducing fraud, or achieving digital transformation goals, all of which are measurable through key performance indicators (KPIs). However, few performance regimes include metrics related to fairness, transparency, or citizen experience. This creates a misalignment between what is rewarded institutionally and what is necessary for accountable, human-centred governance, and the delivery of high-quality and equitable citizen experiences. Vendors, in turn, often tailor their products to meet procurement KPIs rather than align with long-term public values.

Ultimately, the absence of enforceable regulatory frameworks enables AI systems to be deployed in high-stakes public settings without formal oversight. While some jurisdictions have issued guidelines or voluntary standards for ethical AI, few have codified these into binding law, particularly in the context of public administration. In the absence of statutory obligations to assess bias, explainability, or human oversight, many agencies proceed with deployment under the assumption that automation is a neutral or low-risk innovation. Taken together, these structural drivers illustrate how moral hazard in government AI use is less about rogue algorithms and more about institutional cultures, market dynamics, and governance gaps. Addressing these root causes will require reforms in procurement, capacity building, public accountability, and legal frameworks, not only to prevent harm but to build trustworthy public systems in the age of automation.

Consequences of Ignoring AI Moral Hazard

The failure to address moral hazard in government AI systems has wide-ranging and compounding effects that are legal, social, institutional, and ethical in nature. When responsibility is evaded, the systems themselves not only risk delivering harm but also undermine the credibility and integrity of public administration.

One of the most immediate consequences is the erosion of public trust. When citizens interact with opaque or error-prone automated systems and have no clear channel for appeal or recourse, they may come to view government as impersonal, unaccountable, or even hostile. This effect is particularly pronounced among vulnerable populations that rely most heavily on public systems and are least equipped to challenge automated decisions. The perception that AI decisions are made "without a human in the loop" can amplify existing alienation from government institutions.

The normalization of moral hazard in AI also contributes to institutional fragility. Government departments that delegate critical decisions to AI tools without retaining ownership or oversight become less resilient in the face of error, policy change, or public backlash. When harm occurs, officials may struggle to respond effectively, as internal accountability mechanisms are unclear or underdeveloped. Over time, this weakens the rule of law and the government body's ability to maintain procedural fairness.

Moreover, unchecked moral hazard can exacerbate existing social inequalities. Historical bias embedded in datasets or model logic can be perpetuated by AI systems deployed in education, housing, healthcare, and justice. When responsibility for these outcomes is diffused, governments may delay or avoid interventions to correct inequities. This dynamic not only compounds injustice but can also lead to disproportionate legal exposure, especially as more jurisdictions introduce algorithmic transparency or human rights requirements.

Ultimately, ignoring moral hazard undermines the very foundation of democratic governance. AI systems that operate without public scrutiny or participatory design processes risk being perceived as instruments of bureaucratic control rather than tools of public service. The quiet transfer of decision authority from elected representatives and civil servants to automated systems and private contractors diminishes democratic legitimacy and may provoke long-term resistance, noncompliance, or litigation.

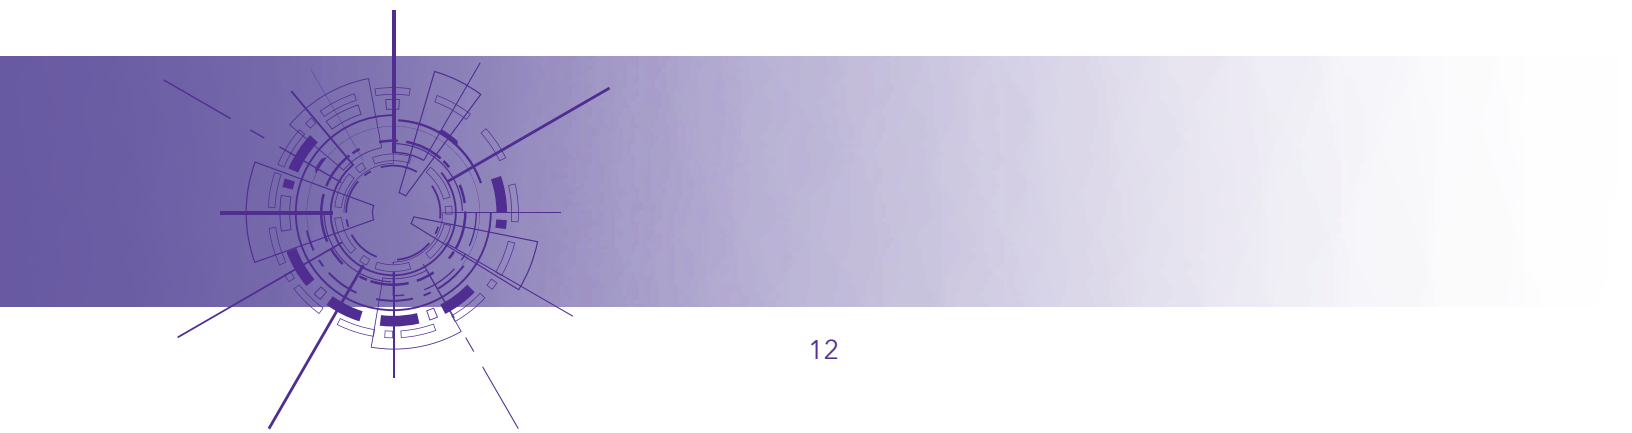

What Can Be done to Address Moral Hazard

Policy and Governance Recommendations

Addressing AI moral hazard in the public sector requires structural reforms and proactive governance mechanisms that anchor responsibility, ensure transparency, and foster trust. Governments must shift from simply deploying AI tools to governing them meaningfully, grounded in public interest, ethical design, and enforceable accountability.

A foundational step is to establish clear lines of accountability. This includes defining who is responsible for decisions made or supported by AI, whether it's a government official, an agency, or a vendor. Where AI is used in decision-making, public bodies must retain formal ownership and liability for outcomes. Contractual frameworks should require vendors to disclose model assumptions, limitations, and auditability features, and to share responsibility for any harms that may arise when appropriate.

Transparency must also be institutionalized through mandatory impact assessments and algorithmic registries. Governments should be required to publish information on the AI systems in use, the decisions they influence, and the evaluations conducted to assess bias, fairness, and privacy risks. The EU's AI Act, Canada's Directive on Automated Decision-Making, and New Zealand's Algorithm Charter offer early examples of how transparency mandates can improve oversight.

Third-party auditing and oversight are essential to maintaining public accountability. Independent auditors, ombuds services, or civil society watchdogs should be empowered to inspect AI systems, investigate harm, and review data governance practices. These institutions should be protected from political interference and have the authority to mandate corrective action where necessary.

Procurement processes must evolve to incentivize ethical design. Requests for proposals (RFPs) should include criteria for explainability, accessibility, and harm mitigation, in addition to cost and efficiency. Governments should favour open standards, source code access, and performance guarantees that align with public values.

Finally, public participation must be integrated into the AI policy lifecycle. Citizens, particularly those most impacted by automation, should be consulted in the design, deployment, and evaluation of AI systems. Co-design practices and participatory governance help to ensure that AI tools reflect the needs and experiences of diverse communities, rather than reinforcing institutional blind spots.

International Perspectives and Best Practices

While many jurisdictions are still in the early stages of AI governance, several countries and regions have taken proactive steps to mitigate the risks of moral hazard through legal, regulatory, and policy innovations.

The European Union's AI Act represents the most ambitious attempt to classify and regulate AI systems based on risk. High-risk systems, such as those used in public service delivery, will be subject to stringent transparency, documentation, and human oversight requirements. Although enforcement mechanisms are still evolving, the Act creates a baseline for liability and compliance across member states.

Canada's Directive on Automated Decision-Making provides a framework for assessing and managing risks associated with the use of AI in federal government services. The directive mandates Algorithmic Impact Assessments (AIAs), tiered governance measures based on system risk level, and requires human-in-the-loop review for higher-risk systems. While not a legislative instrument, it has become a benchmark for public sector AI use.

New Zealand's Algorithm Charter takes a values-based approach, encouraging public agencies to adopt principles of transparency, partnership with Māori communities, and responsible data use. Although the Charter is non-binding, it fosters a cultural shift toward more conscientious automation and sets expectations for ethical stewardship. The United Kingdom's Office for AI and Centre for Data Ethics and Innovation have advanced initiatives in ethical AI procurement and public engagement. Pilot programs involving "data trusts" and responsible innovation sandboxes have provided useful models for balancing innovation with accountability.

Across jurisdictions, a common theme is emerging: transparency, legal accountability, and public oversight are essential conditions for the development of ethical AI. As national frameworks develop, multilateral coordination will also be critical, particularly in standardizing definitions of harm, ensuring cross-border legal remedies, and preventing regulatory arbitrage.

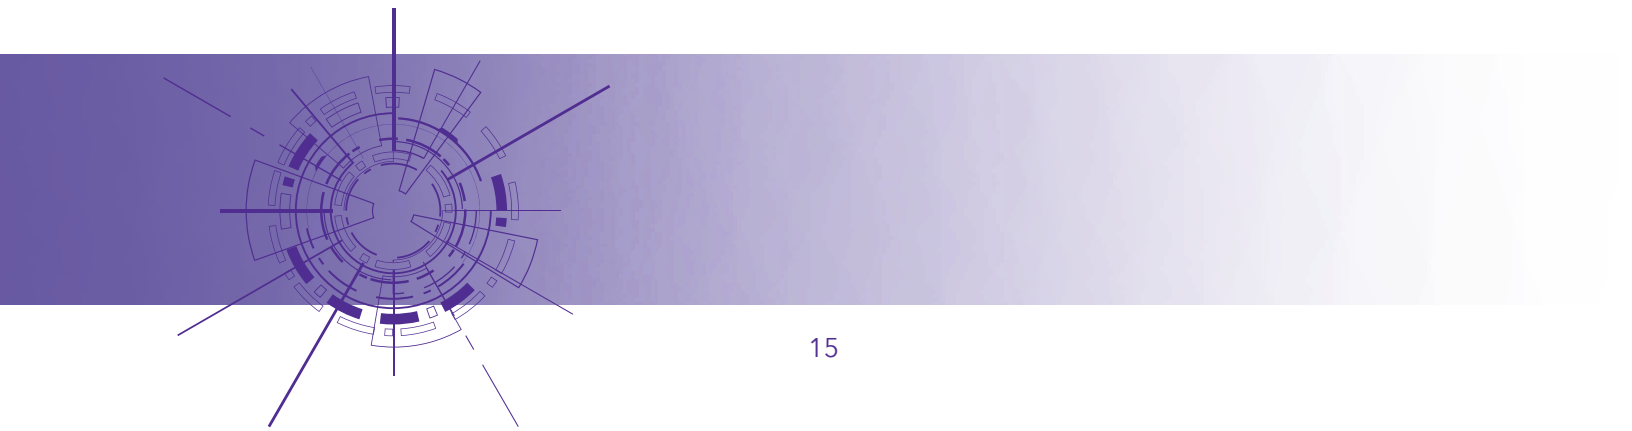

Conclusion

AI has the potential to transform public service delivery for the better, enhancing efficiency, expanding reach, and tailoring support to those in need. However, with this promise comes the responsibility to deploy these systems in a manner that preserves democratic accountability, ensures fairness, and avoids harm. Moral hazard, if left unaddressed, poses a structural threat to these goals.

We have outlined how moral hazard arises when governments adopt AI systems without fully owning the consequences of their decisions. Whether in public insurance, healthcare, or social service delivery, automation can obscure responsibility, displace human judgment, and allow vendors to shape critical systems without bearing the social costs.

To mitigate this, governments must build institutional capacity, strengthen legal frameworks, demand transparency, and embed accountability mechanisms at every stage of the AI lifecycle. Most importantly, they must view automation not as a means of reducing human involvement, but as a public infrastructure that must be governed with the same care, scrutiny, and legitimacy expected of any other function of democratic governance.

AI must not be a black box imposed on society. It must be a public tool, auditable, just, and governed in the public interest.

Action Plan to Avoid Moral Hazard

Action Plan to Avoid Moral Hazard

Item	Action	How
Establish Clear Accountability and Liability Frameworks	Define who is responsible for decisions made or supported by AI (Government officials, agencies, or vendors). Ensure public bodies retain formal ownership and liability for outcomes.	Update procurement contracts to require vendors to disclose model assumptions, limitations, and auditability features. Share responsibility for harms when appropriate.
Mandate Transparency and Impact Assessments	Require mandatory impact assessments and algorithmic registries for all AI systems in use.	Publish information on AI systems, their decision domains, and evaluations for bias, fairness, and privacy risks. Use frameworks like the EU's AI Act, Canada's Directive on Automated Decision-Making, and provincial guidelines (e.g., Ontario's Bill 194, BC's Responsible-Use Principles)
Implement Human-in-the-Loop Reviews	Ensure that all adverse or high-impact decisions made by AI are subject to human review before finalization.	Assign owners for algorithmic decisions, require "algorithm cards" for each system (summarizing logic, risks, and oversight), and update operational protocols to include human checks
Strengthen Procurement Clauses for Ethical AI	Revise Requests for Proposals (RFPs) to include criteria for explainability, accessibility, and harm mitigation, not just cost and efficiency.	Favor open standards, source code access, and performance guarantees that align with public values. Require vendors to participate in third-party audits and share results
Empower Independent Oversight and Auditing	Establish or empower independent auditors, ombuds services, or civil society watchdogs to inspect AI systems and investigate harm.	Protect these institutions from political interference and give them authority to mandate corrective action
Build Technical and Ethical Capacity in Government	Invest in training for public servants on AI risks, ethics, and governance.	Develop internal expertise to evaluate, audit, and challenge technical assumptions. Encourage cross-jurisdictional learning and best practice sharing
Engage Public Participation and Co-Design	Consult citizens, especially those most impacted by automation, in the design, deployment, and evaluation of AI systems.	Use participatory governance and co-design practices to ensure AI tools reflect diverse needs and experiences
Enforce Regulatory and Legal Standards	Codify guidelines and voluntary standards into binding law, especially for high-risk AI systems.	Use international models (EU AI Act, Canada's Directive, New Zealand's Algorithm Charter) and adapt them to local contexts. Monitor compliance and update regulations as technology evolves
Monitor and Manage Knowledge Integrity	Actively manage knowledge ecosystems to avoid redundant, outdated, and trivial information that can pollute AI systems.	Regularly audit and clean organizational data, ensure high-quality content for AI training, and maintain robust knowledge management governance

Works Consulted

1. OECD (2021). State of the Art in the Use of Emerging Technologies in the Public Sector.
2. Ada Lovelace Institute (2022). Examining the Black Box: Tools for Meaningful Accountability in Automated Decision-Making.
3. Pauly, M. V. (1968). The Economics of Moral Hazard. American Economic Review, 58(3), 531–537.
4. Eubanks, V. (2018). Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor.
5. Selbst, A. D., & Barocas, S. (2018). The Intuitive Appeal of Explainable Machines. Fordham Law Review.
6. Government of Canada (2023). Directive on Automated Decision-Making. Treasury Board of Canada Secretariat. Retrieved from canada.ca.
7. AI Now Institute (2022). AI Now 2022 Report.
8. European Commission (2024). AI Act Final Text: Regulation on Artificial Intelligence.
9. Arrow, K. J. (1963). Uncertainty and the Welfare Economics of Medical Care. American Economic Review.
10. Holmstrom, B. (1979). Moral Hazard and Observability. The Bell Journal of Economics.
11. Burrell, J. (2016). How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms. Big Data & Society.
12. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation. International Data Privacy Law.
13. OECD (2021). The Role of Procurement in Responsible AI Adoption by the Public Sector.
14. AI Now Institute (2022). Algorithmic Accountability and Public Sector Use of AI: Challenges and Recommendations.
15. Government Accountability Office (2021). Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities. U.S. GAO.
16. Privacy International (2020). Welfare Surveillance Systems: A Global Problem. Retrieved from privacyinternational.org.
17. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447–453.
18. Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness and Machine Learning: Limitations and Opportunities. Retrieved from fairmlbook.org.
19. Mozilla Foundation (2022). AI and Accessibility in Public Services: Challenges and Opportunities for Inclusion. Retrieved from foundation.mozilla.org.
20. Ada Lovelace Institute (2022). Rethinking Data and AI Accountability in the Public Sector: Proposals for Meaningful Oversight. Retrieved from adalovelaceinstitute.org.
21. OECD (2021). The Role of Public Procurement in the Adoption of Responsible AI. OECD Working Papers on Public Governance.
22. Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the Draft EU Artificial Intelligence Act. Computer Law Review International, 22(4), 97–112.
23. Government of Canada (2020). Directive on Automated Decision-Making. Treasury Board of Canada Secretariat. Retrieved from canada.ca.
24. O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown Publishing.
25. United Nations Human Rights Council (2021). The Right to Privacy in the Digital Age: Report of the UN High Commissioner for Human Rights.
26. Kroll, J. A., et al. (2017). Accountable Algorithms. University of Pennsylvania Law Review, 165(3), 633–706.
27. Ada Lovelace Institute (2022). The Data Divide: Public Attitudes Toward the Use of AI in Government Services.
28. AI Now Institute (2020). Confronting Black Boxes: A Shadow Report on Algorithmic Accountability in the Public Sector.
29. OECD (2021). The Role of Public Procurement in the Adoption of Responsible AI. OECD Working Papers on Public Governance, No. 45.
30. Mozilla Foundation (2021). Foundations of Trust: Why Participation Matters in Public Sector AI. Retrieved from foundation.mozilla.org.
31. European Commission (2024). Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (AI Act).
32. Government of Canada (2023). Algorithmic Impact Assessment Tool. Retrieved from canada.ca.
33. Stats NZ & Department of Internal Affairs (2020). New Zealand Government Algorithm Charter for Aotearoa New Zealand. Retrieved from data.govt.nz.
34. Doctors got worse at cancer screening after using AI helpers (2025). The Neuron, https://www.theneuron.ai/explainer-articles/doctors-got-worse-at-cancer-screening-after-using-ai-helpers
35. Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study (2025), The Lancet Gastroenterology and Hepatology, https://www.thelancet.com/journals/langas/article/PIIS2468-1253(25)00133-5/abstract

Beyond the Black Box

Building Transparent and Accountable AI in Government

maximuscanada.ca

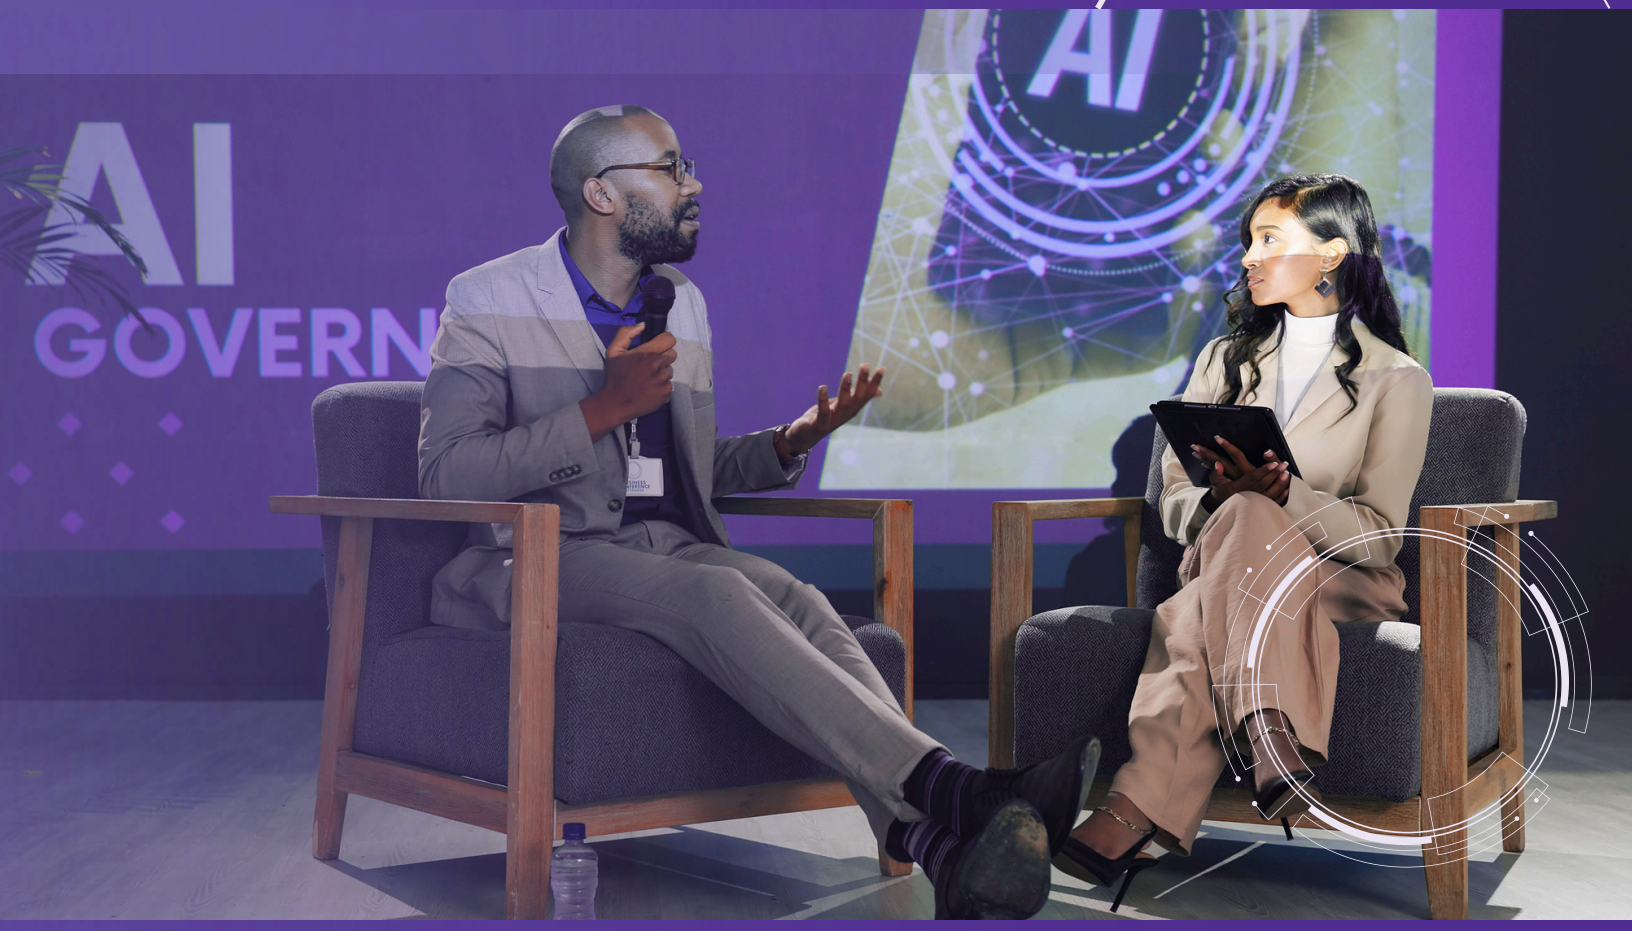

maximus

Moving Canadians, Technology, and Government Forward