REGULATORY TRANSFORMATION IN THE AGE OF AI

JAMIE AMARAT SANDHU NOAM KOLT GILLIAN K. HADFIELD

NOVEMBER 2023



0

ABOUT THE AUTHORS

JAMIE AMARAT SANDHU

Policy Researcher, Schwartz Reisman Institute for Technology and Society; Conducted research and analysis and took primary responsibility for writing.

NOAM KOLT

Vanier Scholar and Doctoral Candidate, University of Toronto Faculty of Law; Graduate Affiliate, Schwartz Reisman Institute for Technology and Society; Conducted research and analysis and contributed to writing.

GILLIAN K. HADFIELD¹

Director, Schwartz Reisman Institute for Technology and Society; Schwartz Reisman Chair in Technology and Society; Professor of Law and Strategic Management, University of Toronto; CIFAR AI Chair, Vector Institute for Artificial Intelligence; Senior Policy Advisor, OpenAI; Conceived of and oversaw research and contributed to writing.

*The authors have no conflicts to declare. **The opinions in this brief are those of the authors only.

¹ Many thanks to Monique Crichlow, Maggie Arai, Isaac Gazendam, Uchenna Felicia Ugwu, Alyssa Wong, and Alessia Woolfe for helpful comments, discussions, and research assistance.

LAND ACKNOWLEDGMENT

We wish to acknowledge this land on which CIFAR operates. For thousands of years it has been the traditional territory of many nations including the Mississaugas of the Credit, the Anishnabeg, the Chippewa, the Haudenosaunee and the Wendat peoples and is now home to many diverse First Nations, Inuit and Métis peoples. We are grateful to have the opportunity to work on this land. We also acknowledge we are all responsible for reconciliation. CIFAR's AI & Society program seeks to advance our understanding of the societal implications of AI to design a future of responsible Al. A future of responsible Al includes one that centres the concerns of Indigenous communities. CIFAR is committed to prioritizing Indigenous perspectives in the development and design of responsible AI.

TABLE OF CONTENTS

- 2 EXECUTIVE SUMMARY
- **3** INTRODUCTION
- 4 THE HARMS PARADIGM
- 6 WIDER IMPACTS
- 8 **REGULATORY TRANSFORMATION**

CASE STUDY 1: HEALTHCARE

CASE STUDY 2: FINANCIAL SERVICES

13 REGULATORY IMPACTS ANALYSIS (RIA)

POLICY FRAMEWORK AND TOOL

CASE STUDY 3: NUCLEAR ENERGY

- **17** CONCLUSION
- **18** REFERENCES

199891

EXECUTIVE SUMMARY

Current efforts to regulate AI focus primarily on reducing harms and mitigating risks presented by the technology. These are important efforts, but their focus is incomplete. As a general-purpose technology, AI has the potential to fundamentally transform society. Framing AI regulation solely in terms of risk mitigation obscures the bigger picture: bracing for unprecedented economic and social change. The widespread diffusion of AI could, and is perhaps already beginning to, upend existing regulatory systems. Drawing on Canadian case studies in healthcare, financial services, and nuclear energy, this policy brief illustrates that AI could challenge the conventional targets and tools of regulation, which would have farreaching implications. We propose a practical tool - "regulatory impacts analysis" (RIA) - to assist policymakers in navigating these challenges and adapting governance infrastructure to an economy transformed by AI. We expect this framework to be useful for policymakers both in Canada and beyond.



1.0 INTRODUCTION

Regulatory responses to AI have to date largely focused on the potential harms and risks associated with the technology, ranging from algorithmic bias and misinformation to accidents caused by autonomous vehicles (Kaminski, 2023). This prevailing approach, which we describe as a "harms paradigm", is necessary but also incomplete (Hadfield and Clark, 2023). While AI does indeed pose significant risks in various domains, its potential implications are far broader. In particular, because AI is a general-purpose technology, it could transform many sectors of the economy and, in doing so, call into question the traditional targets and tools of regulation. Policymakers, we suggest, should prepare for new and changing conditions by analyzing the anticipated impact of AI on governance infrastructure. To facilitate this, we propose a concrete framework for conducting "regulatory impacts analysis" that Canadian and international policymakers could implement at various levels of government. While the framework is by no means comprehensive, it affords policymakers clear directional signals that could be invaluable in the face of disruptive socioeconomic change (Krehm, 1980).

The remainder of the policy brief is organized as follows. Section 2 outlines the prevailing harms paradigm in AI regulation. Section 3 explores the wider impacts of AI as a generalpurpose technology. Section 4 investigates how Al could upend existing regulatory frameworks, including the targets and tools of existing regulation. Section 5 describes the RIA process and illustrates how it could be applied in practice.

> ² This approach seems to align with the concept of "everything-bagel liberalism," according to which regulators seek "to accomplish so much with a single project or policy that ends up failing to accomplish anything at all" (Klein, 2023).

2.0

THE HARMS PARADIGM

SAME, SAME, BUT DIFFERENT.

As Al begins to play a growing role in the economy, regulators appear to have taken a leaf out of their traditional playbook. They focus on the harms potentially posed by Al and adopt regulatory approaches originally developed in conventional product safety and risk mitigation. Prominent examples include the EU Al Act, which is a quintessential risk categorization and mitigation regime (Kaminski, 2023). The Act categorizes Al systems based on their risk levels and imposes corresponding safety requirements. Canada's Artificial Intelligence and Data Act (AIDA), part of Bill C-27, similarly proposes regulating Al by classifying systems according to their level of impact, imposing the most stringent conditions on "high-impact" Al systems. The harms paradigm is also dominant in the United States, typified by the National Institute of Standards and Technology (NIST) Al Risk Management Framework, which proposes principles and procedures for Al enterprise risk management. Although these regulatory frameworks purport to address a wide range of risks from AI, including physical and psychological harm, property damage, economic losses, and discriminatory outcomes, they miss the risk that AI disrupts our ability to achieve regulatory goals across other products and sectors (Hadfield and Clark, 2023; Maas, 2019; 2022). [Moreover, they rely on high-level abstractions and principles that are difficult to implement in practice (Scassa, 2023; Hohma et al., 2023).] The focus on harms and risks of AI as a technology arguably obscures the impact of AI on regulation itself, and the growing need to adapt regulatory strategies and techniques to an economy shaped by AI. It is increasingly apparent that the technology presents complex questions that go beyond the prevailing harms paradigm. Consider, for example, regulation in the healthcare sector. How should healthcare regulators assess the safety and efficacy of AI systems that rely on simulated data instead of traditional human clinical trials? How should regulators divide responsibility and attribute liability between the developers of AI systems and companies deploying the technology in healthcare settings? Answering these questions solely through the narrow lens of risk mitigation fails to come to grips with broader, systemic issues.

Regulators are not alone in subscribing to the harms paradigm. The discourse among technologists building AI systems and social scientists studying the impacts of AI also reinforces the harms paradigm. For example, the seminal academic paper on "foundation models" dedicates significant attention to the risks posed by these systems (Bommasani et al., 2021), as do prominent critics of large language models (Bender et al., 2021). Leading Al companies have adopted a similar posture. For example, OpenAI and Google DeepMind have adopted the harms paradigm when discussing the policy implications of cutting-edge AI systems (Weidinger et al., 2021, 2022; GPT-4 technical report; Anderljung et al., 2023). Legal scholars have followed suit. For instance, Kaminski analogizes Al regulation to risk mitigation encountered in environmental law, privacy law, and cybersecurity (Kaminski, 2023a; 2023b). Kolt similarly focuses on largescale risks from AI, not its broader impact on regulatory infrastructure (Kolt, 2023). This prevailing approach, however, is not a foregone conclusion. Policymakers, we suggest, can and should adopt a broader lens when confronting the societal impact of Al.

3.0 WIDER IMPACTS

Al is a general-purpose means of invention (Cockburn et al., 2019) and problem-solving that has the potential to impact every sector of society. In this respect, AI differs from many disruptive technologies of the past. It prompts a complex interplay of impacts rather than discrete, isolated opportunities and risks. This perspective is encapsulated in the characterization of AI as a general-purpose technology (Trajtenberg, 2018; Brynjolfsson et al., 2017; 2019; Crafts, 2021; Garfinkel, 2022; Goldfarb et al., 2023; Lipsey et al., 2005). The generality of Al can be seen in the many applications of the technology, including personal assistants like ChatGPT and Claude, which provide a broad array of services, ranging from legal and medical advice to computer programming and managing email correspondence. Al search tools such as Bing Chat and Google Bard purport to serve as information intermediaries and repositories of knowledge that could dramatically affect the way people access and consume information. Meanwhile, specialized scientific AI systems like AlphaFold can perform complex cognitive tasks that accelerate scientific research (Korinek, 2023). More recently, AI systems such as AutoGPT have been shown to act as independent agents that can spawn additional agents and work together to perform complex multi-step tasks (Chan et al., 2023).



Seen in this light, the applications and impact of AI extend well beyond the boundaries of any particular sector or regulatory context. It is also apparent that many current regulations were not designed with Al in mind. For instance, car safety regulation was not developed to address the advent of autonomous vehicles. Policies in educational institutions were not equipped to deal with the onslaught of Al-generated content (Colonna, 2022). Legal services were similarly unprepared for the automation of activities traditionally carried out by (human) lawyers (Kluttz and Mulligan, 2019). While these are seemingly isolated examples, each pertaining to a particular sector of the economy, they share in common an underlying problem: regulators do not, at least by default, anticipate and address the cascading effects of transformative technologies such as AI (Frank et al., 2019). In addition to driving change in particular sectors of the economy (Brynjolfsson et al., 2019, 2023), widely adopted AI tools could dramatically impact markets (Gal and Elkin-Koren 2017; Van Loo 2019) and bring about large-scale, unpredictable structural changes (Shevlane and Dafoe, 2021). We now turn to consider the impact of AI on regulation itself.

4.0 REGULATORY TRANSFORMATION

Regulation is commonly understood as "the enterprise of subjecting human conduct to the governance of rules" (Fuller, 1964), which is supported by complex and iteratively evolving social, economic, and legal infrastructure (Hadfield, 2017). Regulation, in other words, is not confined to formal statutes or judicial decisions. It is built up from intricate and highly context-dependent norms. As we will demonstrate, Al poses multiple challenges for regulation (Hadfield and Clark, 2023), a significant number of which are distinct from the challenges posed by other disruptive technologies (Brownsword, 2018; 2019).

While prior studies unpack some of these challenges, including the issue of regulation keeping pace with innovation (Liu et al., 2020, Maas, 2019; 2022, Hopster and Maas, 2023), we aim to provide an overview that will be of practical use to policymakers in Canada. The discussion covers two case studies: healthcare and financial services. Before proceeding, it is helpful to point out that each case study focuses on two distinct, but closely related, impacts of Al. The first concerns how Al could impact the targets of regulation, i.e. the entities to which regulation applies (or aims to apply). The second concerns how Al could impact the tools of regulation, i.e. the mechanisms used to govern the targets of regulation.

CASE STUDY 1: HEALTHCARE

Prior to the advent of Al, regulation has targeted human actors and entities comprised of humans, often within a particular domain. Healthcare regulation, for example, imposes a range of educational and licensing requirements on doctors, nurses, and other healthcare providers. Healthcare professionals must undergo training and, to varying degrees, demonstrate competence on an ongoing basis. Al, however, is beginning to complicate the regulatory landscape. While human professionals are likely to continue to be heavily involved in providing healthcare services, Al tools such as large language models are predicted to play a significant role in the sector (Moor et al., 2023, Lee et al., 2023), as they cut across various professions, particularly high-income occupations (Eloundou et al., 2023; Kreitmeir and Raschky, 2023; Noy and Zhang, 2023).³ As Al changes how healthcare services are delivered, regulation of healthcare professionals may no longer be sufficient and current approaches may no longer be appropriate.

Consider, for example, Google's Med-PaLM, which can perform a variety of biomedical tasks, including mammography and dermatology image interpretation, radiology report generation and summarization, and genomic variant calling (Tu et al., 2023). Tools like this arguably shift the regulatory focus from doctors and conventional healthcare processes to software engineers and the development of AI products and services, raising a host of new questions. For instance, which actors should be required to undergo educational and professional training – human healthcare specialists or AI developers building healthcare applications, or both? Which entity is liable for defective AI-generated medical advice? How can responsibility be shared between these different actors?

³ Compare Lamb 2016, which suggested that bluecollar jobs will be the most impacted by Al in Canada. In addition to questioning who should be the targets of healthcare regulation, Al also prompts us to reflect on which regulatory tools might be appropriate for governing a sector impacted by AI technologies. For example, should standards related to hygiene, patient care, and safe medical practices be changed in light of the use of AI in healthcare settings? What mechanisms can be developed to monitor and enforce compliance? What resources do regulators need in order to make these assessments? While some components of healthcare regulation, such as medical device regulation, could be adapted to the use of AI tools, other components will need to be entirely revisited.

How have Canadian healthcare regulators responded to these challenges? In short, significant work lies ahead, both with respect to addressing new regulatory targets and developing new regulatory tools. Just as Canadian regulators have not yet adapted to the role of new AI-focused entities in providing or supporting healthcare services, they have not yet adapted to the provision of healthcare services through new AI-driven or AI-mediated means (Da Silva et al., 2022). Health Canada, for instance, faces significant challenges in implementing new licensing regimes for AI medical devices (CIFAR, 2020; Régis and Flood, 2021; Da Silva et al., 2022).

The case of medical device regulation is particularly illuminating. Health Canada currently aims to address medical AI using a product safety approach supported by provisions in the Food and Drugs Act and Medical Devices Regulation and a supporting set of AI-specific principles (Health Canada, 2021; 2022). In large part, however, medical devices are categorized into four classes, ranging from Class I (low-risk devices like wheelchairs) to Class IV (high-risk devices like defibrillators) (Medical Devices Regulations to Government of Canada, 1998 as amended in 2023; Health Canada, 2019), which determines the stringency of the applicable requirements. The classification of AI largely remains an open question. At present, AI developers would likely be considered medical device manufacturers, resulting in them being responsible for determining the risk classification of medical Als (Da Silva et al., 2022). The problem, however, is that Al systems are not necessarily standalone devices or products. They are dynamic tools that are highly sensitive to the contexts in which they are deployed (Gulshan et al., 2016). Seen in this light, how can a regulatory regime designed for evaluating traditional, narrow-purpose medical devices be applied to general-purpose medical Als (Gerke et al., 2020)? Tracing the causal chain between an adverse outcome and the AI, alongside other human and organizational factors that contribute to that outcome is not trivial (Hadfield and Clark, 2023). Allocating liability among different actors is equally challenging.

A failure to address these issues could have significant implications for patients, healthcare providers, and the entire healthcare system. Critically, regulators are often at a notable information disadvantage compared with Al developers in the medical domain. For example, there is little historical data available for comparing different Al-driven personalized treatments (Dankwa-Mullan and Weeraratne, 2022). As a result, the public may lose trust in healthcare providers that use AI tools or, worse still, repose trust in healthcare providers and technologies where trust is unwarranted (Kelly et al., 2019). Without thoughtful intervention, these dynamics could undermine the broader integrity of the healthcare system.

CASE STUDY 2: FINANCIAL SERVICES

Financial regulation aims to protect the financial system by tackling fraud, market manipulation, and unfair practices. To achieve this goal, regulators have traditionally targeted human actors, specifically individuals involved in financial activities, and traditional institutions such as banks and investment firms. Regulatory tools have typically included a combination of mandatory disclosure, monitoring, auditing, and professional training and licensing. To be clear, these tools are designed to govern human actors and institutions, ensuring that all parties understand their rights, obligations, and expectations when it comes to providing or receiving financial services. For example, financial analysts must obtain a Chartered Financial Analyst certification to carry out various financial services and adhere to legally binding financial services agreements, which outline the terms and conditions of their relationship with clients and financial institutions. These requirements are complemented by the threat of personal sanctions in the event of non-compliance (Azzutti et al., 2021).

Until recently, this approach to financial regulation appeared satisfactory, at least in principle.⁴ But this is likely to change as AI tools are integrated more broadly and deeply into the financial sector. For instance, algorithmic trading now accounts for a growing fraction of market activity (Fortune Business Insights, 2023). AI-driven financial and investment services are also available to retail clients through trading platforms like Aiden Arrival used by the Royal Bank of Canada (Borealis AI, 2022). In addition to exposing investors to greater risk, due to their vulnerability to minor market fluctuations,

Al tools could also pose systemic risks to financial markets.⁵ For example, using popular large language models to assist in predicting stock performance (Lopez-Lira and Tang, 2023) could result in highly correlated or homogenous investment decisions that have cascading effects on markets. In addition, efforts to build autonomous agents that can independently manage an online trading business (Suleyman, 2023) could, if successful, spawn innumerable Al agents that interact with each other in unpredictable and potentially destabilizing ways.⁶

⁴ Of course, financial regulations can and do fail, e.g. 2007–2008 financial crisis.

⁵ A canonical example is the 2010 Flash Crash (Kirilenko et al., 2017).

⁶ See also: Ezrachi and Stucke (2017), on how Al collusion undermines competition law.

These developments pose difficult questions for financial regulators. Facing the overarching challenge that "AI systems are a somewhat unusual animal ... to regulate" (Azzutti et al., 2022), regulators must contend with new financial actors for which traditional regulatory tools are ill-suited. For example, are robo-advisors subject to fiduciary obligations comparable to those imposed on their human counterparts (Lee, 2020)? What safeguards should be established to prevent over-reliance on these AI systems? How might antitrust regulators ensure that smaller market participants have equal opportunities to compete amidst the growing prevalence of algorithmic trading and robo-advisors (Agrawal et al., 2019)? Who bears liability for Al-based investment advice that results in substantial losses? How can regulators prevent malicious actors from using open source AI tools (such as BloombergGPT) from orchestrating market manipulation and engaging in financial crime?

Most of these questions remain unanswered. While this may result from the organizational culture of regulators, it also stems from a lack of AI expertise and resources - which might also explain why the regulatory responses to the adoption of AI in the Canadian financial services industry are limited to a few white papers and reports. To date, there is no binding Al-specific regulation for the financial sector (Savoie, 2023; Aziz et al., 2021). While AIDA could change the regulatory landscape by imposing stringent obligations on "high impact" Al systems that result in "economic loss to an individual", its consequences for the financial services sector remain uncertain. Moreover, the focus on individual harm, rather than systemic risk, arguably obscures the central issue presented by AI (Scassa, 2023).7 To better contend with these issues, policymakers will need to more carefully consider the impact of AI on financial markets and, just as importantly, examine how AI could challenge the traditional targets and tools of financial regulation.

⁷ But see discussion of "collective harms" in the AIDA companion document.

5.0 REGULATORY IMPACTS ANALYSIS (RIA)

Adapting to regulatory change is difficult for policymakers at the best of times. Adapting to regulatory transformation prompted by AI is an even more formidable challenge. As discussed, AI technologies are likely to significantly shift the targets of regulation and, to varying degrees, render many existing regulatory tools ineffective or obsolete. While policymakers cannot predict the precise scope or magnitude of this regulatory transformation, they can take concrete steps to better prepare for at least some of the changes that are in store. This section proposes that policymakers conduct regulatory impacts analysis ("RIA") and illustrates how this might be done. Our aim is to assist policymakers in understanding and anticipating the impact of AI on regulation in their respective domain, and provide them with information to recalibrate the targets and tools of existing regulatory regimes. The discussion begins with a brief description of the RIA process, which is then followed by a case study demonstrating how policymakers could use RIA in practice.

POLICY FRAMEWORK AND TOOL

Given the pervasiveness of the harms paradigm in Al policy and governance, as discussed above, it is no surprise that there exist many frameworks for evaluating the risks posed by Al systems (Hendrycks and Mazeika 2022; Khlaaf, 2023), along with methods for documenting and communicating the resulting findings (Mitchell et al., 2018; Gebru et al., 2018; Gilbert et al., 2022). To date, however, there is no accepted framework for evaluating the impact of Al on regulatory regimes or systems (Raji et al., 2022; Costanza-Chock et al., 2022).

To fill this gap, we propose regulatory impacts analysis ("RIA"): a novel framework and procedure for analyzing the impact of AI on regulatory systems. RIA is designed to (1) assess the likely impact of AI on the targets and tools of regulation, and (2) assist policymakers in adapting governance institutions to the new and changing conditions arising from AI. This framework will be particularly useful to policymakers involved in designing regulatory tools and institutions.⁸

Concretely, RIA is facilitated by presenting policymakers with a questionnaire. The purpose of the questionnaire is to prompt policymakers in a given domain to engage in structured analysis of the likely impact of AI on the regulatory institutions and methods that they are tasked with administering. To be clear, the RIA questionnaire is not supposed to serve as an exhaustive checklist describing all regulatory implications of Al. Rather, the questionnaire seeks to encourage policymakers to engage in frank conversations about current and anticipated regulatory challenges arising from AI, to engage necessary technology and other experts to understand where change is likely to come and when, and to begin to consider what steps can be taken to better prepare for these challenges. We envision the questionnaire guiding discussions or workshops with a wide range of experts, from within and without government. Sample questions are displayed in Box 1.

> ⁸ RIA could also inform the work of technologists building AI and academic researchers studying the broader social and economic impacts of AI.

BOX 1: SAMPLE RIA QUESTIONNAIRE	
REGULATORY TARGETS	 Currently, who are the primary targets of regulation in your domain? Who else are you responsible for regulating? Do most regulatory requirements currently apply to these people and/or organizations? Which actors are not currently regulated but should be regulated? How will AI technologies and applications change your answers to questions 1–3? Are these changes already taking place? If so, what are the most significant changes to date? In what timeframe do you anticipate further changes to take place? If AI were deployed more widely in your domain or were relied upon to a greater extent, in which additional ways would your answers to questions 1–3 change?
REGULATORY TOOLS	 6. Currently, what are the primary tools, mechanisms, and methods of regulation in your domain? 7. How do you currently administer and apply these regulatory tools, mechanisms, and methods? 8. Which regulatory tools do you currently refrain from using, and why? 9. How will AI technologies and applications change your answers to questions 7–9? Are these changes already taking place? If so, what are the most significant changes to date? In what timeframe do you anticipate further changes to take place? 10. If AI were deployed more widely in your domain or were relied upon to a greater extent, in which additional ways would your answers to questions 7–9 change?
NEXT STEPS	 Based on your answers above, list (a) regulatory targets in your domain that Al will render less important, and (b) regulatory targets in your domain that Al will render more important. How do you plan to regulate these increasingly important regulatory targets? Provide 3–5 concrete recommendations that your organization should implement, and the timeline for implementation. Based on your answers above, list (a) regulatory tools that Al will render less effective in your domain, and (b) new regulatory tools that should be developed in light of the use of Al in your domain. How do you plan to develop these new regulatory tools? Provide 3–5 concrete recommendations that your organization should implement, and the timeline for implementation. Do you have the resources to implement the recommendations listed in questions 12 and 14? If not, what additional resources do you require? How could you acquire those additional resources?

CASE STUDY 3: NUCLEAR ENERGY

To illustrate how RIA could operate in practice, it is helpful to consider a concrete example: the regulation of nuclear energy and materials. In Canada, regulation in this domain is led by the Canadian Nuclear Safety Commission (CNSC). The CNSC currently focuses its attention on organizations that operate nuclear facilities and activities, including engineers and other personnel in these organizations. The CNSC employs various regulatory tools to oversee and enforce safety requirements and standards that apply to the use and transportation of nuclear materials and equipment. Key regulatory tools include licensing regimes and certification requirements (CNSC, 2020).

As AI systems are increasingly used in nuclear energy (Degrave et al., 2022; Lu et al., 2020), how can the CNSC use the RIA framework to adapt its regulatory targets and regulatory tools? First, by responding to questions 1–5 (see Box 1), the CNSC could gain further clarity around new regulatory targets it needs to consider. For example, AI technologies for conducting offensive cyber operations could potentially expand the range of malicious actors who seek to gain access to, or exploit, Canadian nuclear materials and technologies. Second, by responding to questions 6-10 (see Box 1), the CNSC could better understand which kinds of regulatory tools are increasingly important as AI is integrated into the infrastructure used for nuclear power in Canada. For instance, tools that monitor cybersecurity threats and enforce cybersecurity standards will be critical if nuclear materials are controlled (whether completely or partially) by AI systems. Finally, questions 11-15 (see Box 1) will assist the CNSC in preparing for these changes and building a concrete plan for recalibrating their regulatory targets and developing new regulatory tools, including by encouraging the CNSC to examine whether it has the resources to implement the necessary changes.

Both in nuclear energy and in other regulatory contexts, a few important clarifications are warranted. While RIA, if conducted appropriately, could elicit useful information about the regulatory impact of AI, a one-time procedure is unlikely to be sufficient. Given the impacts of AI are evolving and often unpredictable, RIA will need to be conducted periodically. Regular RIA audits, perhaps on an annual basis, could assist policymakers in revisiting the implications of AI for regulatory targets and tools in a sufficiently timely manner.

Another issue concerns implementation. Like any policy tool, the success of RIA will turn on a policymaking body's resources and organizational culture. Without the resources to implement the recommendations emerging from the RIA procedure, progress on regulatory priorities cannot be made. Similarly, rigid organizational cultures could frustrate efforts to address new regulatory targets and administer new regulatory tools. While exploring strategies for overcoming these obstacles is outside the scope of this policy brief, suffice to say that RIA will be most effective if undertaken by actors with significant institutional support and internal organizational clout.

6.0 CONCLUSION

This policy brief has aimed to demonstrate that regulatory responses to AI focused only on the technology's harms are incomplete. As illustrated in a growing number of applications, Al is a general-purpose technology that stands to transform regulation itself - both the targets and tools of regulation. Policymakers need to proactively prepare for these developments by analyzing the anticipated impact of AI on governance frameworks and, where appropriate, re-designing those frameworks for an economy shaped by AI. To facilitate this process, we propose a practical tool for conducting regulatory impacts analysis. We expect the tool to assist policymakers at various levels of government both in Canada and beyond.

7.0 REFERENCES

Kaminski, Margot. E. (2023). "Regulating the Risks of Al." Forthcoming, Boston University Law Review, 103. <u>http://dx.doi.org/10.2139/ssrn.4195066</u>

Klein, Ezra. (April 2, 2023). The problem with everything-bagel liberalism. The New York Times. <u>https://www.nytimes.com/2023/04/02/opinion/democrats-liberalism.html</u>

Hadfield, Gillian. K., & Clark, Jack. (2023). "Regulatory Markets: The Future of Al Governance." arXiv preprint arXiv:2304.04914.

Krehm, William. (1980). How to Make Money in a Mismanaged Economy, and Other Essays. COMER Publications.

Maas, Matthijs. M. (2019). "International Law Does Not Compute: Artificial Intelligence and the Development, Displacement or Destruction of the Global Legal Order." Melbourne Journal of International Law, 20(1), 29-57.

Maas, Matthijs. M. (2022). "Aligning AI Regulation to Sociotechnical Change." Oxford Handbook on AI Governance (Oxford University Press, 2022 forthcoming).

Scassa, Teresa. (2023). "Regulating AI in Canada: A Critical Look at the Proposed Artificial Intelligence and Data Act." The Canadian Bar Review. 101(1).

Hohma, Ellen. Boch, A. Trauth, R. Lütge, C. (2023). "Investigating Accountability for Artificial Intelligence Through Risk Governance: A Workshop-Based Exploratory Study." Frontiers in Psychology. 14.

Bommasani, Rishi., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). "On the Opportunities and Risks of Foundation Models." arXiv preprint arXiv:2108.07258.

Bender, Emily. M., Gebru, T., McMillan-Major, A., & Mitchell, M. (2021, March 1). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? [FAcct '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency]. Association for Computing Machinery.

Weidinger, Laura., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P. S., ... & Gabriel, I. (2021). "Ethical and Social Risks of Harm From Language Models." arXiv preprint arXiv:2112.04359.

Weidinger, Laura., Uesato, J., Rauh, M., Griffin, C., Huang, P.-S., Mellor, J., Glaese, A., Cheng, M., Balle, B., Kasirzadeh, A., Biles, C., Brown, S., Kenton, Z., Hawkins, W., Stepleton, T., Birhane, A., Hendricks, L. A., Rimell, L., Isaac, W., ... Gabriel, I. (2022). "Taxonomy of Risks Posed by Language Models." 2022 ACM Conference on Fairness, Accountability, and Transparency. <u>https://doi. org/10.1145/3531146.3533088</u>

OpenAI (2023). GPT-4 Technical Report. ArXiv. https://doi.org/arXiv:2303.08774v3

Anderljung, Markus., et al., (2023). "Frontier AI Regulation: Managing Emerging Risks to Public Safety." arXiv preprint arXiv:2307.03718

Kaminski, Margot. E. (2023, April 20). The Developing Law of Al Regulation: A Turn to Risk Regulation.Lawfare. <u>https://www.lawfareblog.com/developing-law-ai-regulation-turn-risk-regulation</u>

Kolt, Noam. (2023). "Algorithmic Black Swans." Washington University Law Review, 101. https://ssrn.com/abstract=4370566

Cockburn, Iain. M., Henderson, R., & Stern, S. (2019). The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis. In The Economics of Artificial Intelligence. University of Chicago Press.

Trajtenberg, Manuel. (2018). "Al as the next GPT: a Political-Economy Perspective" (No. w24245). National Bureau of Economic Research.

Brynjolfsson, Erik., & Mitchell, T. (2017). "What Can Machine Learning Do? Workforce Implications." Science, 358(6370), 1530-1534.

Brynjolfsson, Erik., Rock, D., & Syverson, C. (2019). "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." (No. c14007). National Bureau of Economic Research.

Crafts, Nicholas. (2021). "Artificial Intelligence as a General-Purpose Technology: An Historical Perspective." Oxford Review of Economic Policy, 37(3).

Garfinkel, Ben. (2022). The Impact of Artificial Intelligence: A Historical Perspective, in Justin B. Bullock and others (eds), The Oxford Handbook of Al Governance. <u>https://doi.org/10.1093/oxfordhb/9780197579329.013.5</u>

Goldfarb, Avi., Taska, B., & Teodoridis, F. (2023). "Could Machine Learning be a General Purpose Technology? A Comparison of Emerging Technologies Using Data From Online Job Postings." Research Policy, 52(1), 104653.

Lipsey, Richard. G., Carlaw, K. I., & Bekar, C. T. (2005). Economic Transformations: General Purpose Technologies and Long-term Economic Growth. OUP Oxford.

Korinek, Anton. (2023). "Language Models and Cognitive Automation for Economic Research." (No. w30957). National Bureau of Economic Research.

Chan, Alan., Salganik, R., Markelius, A., Pang, C., Rajkumar, N., Krasheninnikov, D., Langosco, L., He, Z...et al.,. (2023). "Harms From Increasingly Agentic Algorithmic Systems." In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (pp. 651–666). June 2023.

Colonna, Liane. (2022). "The AI Regulation and Higher Education: Preliminary Observations and Critical Perspectives." Faculty of Law, Stockholm University Research Paper No. 102.

Kluttz, Daniel., & Mulligan, Deirdre. K. (2019). "Automated Decision Support Technologies and the Legal Profession." Berkeley Technology Law Journal, 34, 853-890.

Frank, Morgan. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., ... & Rahwan, I. (2019). "Toward Understanding the Impact of Artificial Intelligence on Labor. Proceedings of the National Academy of Sciences, 116(14), 6531-6539.

Brynjolfsson, Erik., Li, D., & Raymond, L. R. (2023). "Generative AI at Work." (No. w31161). National Bureau of Economic Research.

Gal, Michal. S., & Elkin-Koren, Niva. (2017). "Algorithmic Consumers." Harvard Journal of Law and Technology, 30(2), 309.

Van Loo, Rory. (2018). "Digital Market Perfection." Mich. L. Rev., 117, 815.

Shevlane, Toby, Dafoe, Allan. (2021). "The Machinery of Power: Artificial Intelligence as a General-Purpose Power Technology." unpublished

Fuller, Lon. L. (1969). The Morality of Law: Revised Edition. Yale University Press.

Hadfield, Gillian. K. (2017). Rules for a Flat World: Why Humans Invented Law and How to Reinvent it for a Complex Global Economy. Oxford University Press.

Brownsword, Roger. (2018). "Law and Technology: Two Modes of Disruption, Three Legal Mind-sets, and the Big Picture of Regulatory Responsibilities." Indian JL & Tech., 14(1).

Brownsword, Roger. (2019). "Law Disrupted, Law Re-imagined, Law Re-invented." Technology and Regulation, 2019, 10-30.

Liu, Hin-Yan., Maas, M., Danaher, J., Scarcella, L., Lexer, M., & Van Rompaey, L. (2020). "Artificial Intelligence and Legal Disruption: A New Model for Analysis." Law, Innovation and Technology, 12(2), 205-258.

Hopster, Jeroen K. G., & Maas, Matthijs. M. (2023). "Triaging the Technology Triad: Disruptive AI, Regulatory Gaps and Value Change." AI and Ethics.

Moor, Michael., Banerjee, O., Abad, Z. S. H., et al.,. (2023). "Foundation Models for Generalist Medical Artificial Intelligence." Nature, 616, 259–265.

Lee, Peter. Bubeck, Sebastien., & Petro, Joseph. (2023). "Benefits, Limits, and Risks of GPT-4 as an AI Chatbot for Medicine." The New

England Journal of Medicine.

Eloundou, Tyna. Manning, S. Mishkin, P. Rock, D (2023). "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." arXiv preprint arXiv:2303.10130.

Kreitmeir, David. Raschky, Paul. (2023). "The Unintended Consequences of Censoring Digital Technology - Evidence from Italy's ChatGPT Ban." Available at SSRN: <u>https://ssrn.com/abstract=4422548</u>.

Noy, Shakked. Zhang, Whitney. (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." Science, 381(6654), 187-192.

Lamb, Creig. (2016). The Talented Mr. Robot: The Impacts of Automation on Canada's Workforce. Brookfield Institute. June 2016.

Tu, Tao. Azizi, S. Driess, D. Schaekermann, M. Amin, M., Chang, P. C., ... Natarajan, V. (2023). "Towards Generalist Biomedical Al." arXiv:2307.14334.

Da Silva, Michael., Flood, C. M., Goldenberg, A., & Singh, D. (2022). "Regulating the Safety of Health-Related Artificial Intelligence." Healthcare Policy, 17(4), 63.

Canadian Institute for Advanced Research (CIFAR). (2020), July. Building a Learning Health System for Canadians: Report of the Artificial Intelligence for Health Task Force.<u>https://cifar.ca/wp-content/uploads/2020/11/AI4Health-report-ENG-10-F.pdf.</u>

Régis Catherine. Flood Colleen. M. (2021). "Al & Health Law in Canada" In Bariteau-Martin F., Scassa T., eds., Artificial Intelligence and the Law in Canada (pp. 203–28). LexisNexis.

Health Canada (December 2022). Regulating Advanced Therapeutic Products. <u>https://www.canada.ca/en/health-canada/services/</u> <u>drug-health-product-review-approval/regulating-advanced-therapeutic-products.html.</u>

Health Canada. (2021). Good Machine Learning Practices for Medical Device Development: Guiding Principles. Ministry of Health.

https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/good-machine-learning-practice-medical-device-development.html.

Medical Devices Regulations (SOR/98-282), Amended by SOR/2023-19 (1998). Government of Canada.

Health Canada. (2019). Guidance Document, Software as a Medical Device (SaMD): Definition and Classification. Ministry of Health. December 18, 2019.

https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/application-information/guidance-documents/software-medical-device-guidance-document.html.

Gulshan, Varun. Peng, L. Coram, M. Stumpe, M. C., Wu, D., Narayanaswamy, A. Venugopalan, S. Widner, K. Madams, T. Cuadros, J., Kim, R. Raman, R. Nelson, P. C. Mega, J. L. & Webster, D. R. (2016). "Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs." JAMA, 316(22), 2402-2410.

Gerke, Sara. Minssen, T. & Cohen, I. G. (2020). "Ethical and Legal Challenges of Artificial Intelligence-Driven Healthcare" In A. Bohr & K. Memarzadeh (Eds.), Artificial Intelligence in Healthcare (1st ed). Elsevier. ISBN: 9780128184387.

Dankwa-Mullan, Irene., & Weeraratne, Dilhan. (2022). "Artificial intelligence and machine learning technologies in cancer care: Addressing disparities, bias, and data diversity." Cancer Discovery, 12(6), 1423-1427.

Kelly, Christopher. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). "Key Challenges for Delivering Clinical Impact with Artificial Intelligence. BMC medicine, 17, 1-9.

Azzutti, Alessio., Ringe, W., & Stiehl, H. (2021). "Machine Learning, Market Manipulation, and Collusion on Capital Markets: Why the "Black Box" Matters." University of Pennsylvania Journal of International Law, 43(1), 79-136.

Fortune Business Insights. (2023). Market Research Report (Report No. FBI107174). <u>https://www.fortunebusinessinsights.com/algorithmic-trading-market-107174</u>.

Borealis AI. (2022, November 15). RBC Capital Markets announces the launch of Aiden® Arrival, the second algorithm on the Aiden® platform. RBC Capital Markets, Borealis AI. <u>https://www.newswire.ca/news-releases/rbc-capital-markets-announces-the-launch-of-aiden-r-arrival-the-second-algorithm-on-the-aiden-r-platform-863109623.html.</u>

Kirilenko, Andrei. Kyle, A. S. Samadi, M. & Tuzun, T. (2017). "The Flash Crash: High-Frequency Trading in an Electronic Market." The Journal of Finance, 72(3), 967-998.

Lopez-Lira, Alejandro. & Tang, Yuehua. (2023). "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large

Language Models." SSRN. https://ssrn.com/abstract=4412788.

Suleyman, M. (2023, July 14). "Mustafa Suleyman: My new Turing test would see if Al can make \$1 million." MIT Technology Review.

https://www.technologyreview.com/2023/07/14/1076296/mustafa-suleyman-my-new-turing-test-would-see-if-ai-can-make-1-million.

Ezrachi, Ariel. & Stucke, Maurice. E. (2017). "Artificial Intelligence & Collusion: When Computers Inhibit Competition." University of Illinois Law Review, Vol.2017(2017).

Azzutti, Alessio. Ringe, W.-G. & Stiehl, H. S. (2022). "The Regulation of Al Trading from an Al Life Cycle Perspective." European Banking Institute Working Paper Series, 2022 - no. 130.

Lee, Joseph. (2020). "Access to Finance for Artificial Intelligence Regulation in the Financial Services Industry." European Business Organization Law Review, 21, 731–757.

Agrawal, Ajay. Gans, J. Goldfarb, A. (2019). "Economic Policy for Artificial Intelligence." Innovation Policy and the Economy (Vol. 19). National Bureau of Economic Research.

Savoie, Isabelle. (2023, April 18). "Mitigating Material Risk: Leveraging Consumer Protection Legislation as a Regulator of Client-Facing Al in the Canadian Financial Services Sector." Schwartz Reisman Institute for Technology and Society. University of Toronto.

Aziz, Peter. A., Daniels, M. A., & Goulard, B. (2021). "New Regulations Complete Overhaul of Bank Act Consumer Provisions. Torys LLP.

Innovation, Science and Economic Development Canada (ISED). (2023). "The Artificial Intelligence and Data Act (AIDA) – Companion document." Government of Canada.

https://ised-isde.canada.ca/site/innovation-better-canada/en/artificial-intelligence-and-data-act-aida-companion-document.

Hendrycks, Dan., & Mazeika, Mantas. (2022). "X-risk Analysis for Al Research." arXiv preprint arXiv:2206.05862.

Khlaaf, Heidy. (2023). "Toward Comprehensive Risk Assessments and Assurance of Al-Based Systems." Trail of Bits.

Mitchell, Margaret., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., ... & Gebru, T. (2018). "Model Cards for Model Reporting." In Proceedings of the conference on fairness, accountability, and transparency (pp. 220-229).

Gebru, Timnit., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Iii, H. D., & Crawford, K. (2018). "Datasheets for Datasets." Communications of the ACM, 64(12), 86-92.

Gilbert, Thomas. K. Dean, S. Lambert, N. Zick, T. Snoswell, A. (2022). "Reward Reports for Reinforcement Learning." arXiv preprint arXiv:2204.10817.

Raji, Inioluwa Deborah. Kumar, I. E. Horowitz, A. Selbst, A. (2022). The fallacy of AI functionality. In Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (pp. 959-972).

Costanza-Chock, Sasha. Raji, I. D. Buolamwini, J. (2022). "Who Audits the Auditors? Recommendations From a Field Scan of the Algorithmic Auditing Ecosystem." In Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (pp. 1571-1583).

Canada Nuclear Safety Commission (CNSC). (2020). CNSC Departmental Results Report 2019-20. Minister of Natural Resources. Government of Canada.

http://nuclearsafety.gc.ca/eng/resources/publications/reports/departmental/drr-2019-2020/index.cfm.

Degrave, Jonas. Felici, F., Buchli, J. et al. (2022). "Magnetic Control of Tokamak Plasmas Through Deep Reinforcement Learning." Nature, 602, 414–419.

Lu, Chao. Lyu, J. Zhang, L. Gong, A. Fan, Y. Yan, J. Li, X. (2020). "Nuclear Power Plants With Artificial Intelligence in Industry 4.0 Era: Top-Level Design and Current Applications—A Systematic Review." IEEE Access, 8, 194315-194332.



1]

MaRS Centre, West Tower 661 University Ave., Suite 505 Toronto, ON M5G 1M1 Canada

<u>www.cifar.ca/al</u>