UNIVERSITY OF TORONTO
FACULTY OF LAW

November 6, 2023

Mr. Joël Lightbound, M.P.
Chair, Standing Committee on Industry and Technology
House of Commons
Ottawa, ON K1A 0A6

**Re: Submission to the Standing Committee on Industry and Technology Study of Bill C-27, the Digital Charter Implementation Act, 2022.**

Dear Mr. Lightbound,

As a scholar of artificial intelligence and emerging technology regulation, and a member of the Canadian AI Advisory Council, I am writing to provide comments on your committee's study of Bill C 27. The *Artificial Intelligence and Data Act* (AIDA) is a crucial piece of legislation. Artificial intelligence (AI) and other advanced technologies are incredibly powerful; they can both provide immense benefits and cause significant harm. Parliament should enact AIDA as soon as possible. However, it suffers from critical flaws. This brief covers some areas of concern and proposes six mitigating recommendations.

***Amend or supplement* AIDA *to address the systemic harms of AI and create a registry for large AI models.*** As the recent dramatic regulatory moves made by the U.S. and the U.K. in relation to frontier AI systems demonstrate, the individual harms that AIDA (like the EU AI Act) was drafted to address now sit alongside widespread recognition of fundamentally systemic and potentially catastrophic risks that large models might pose. One example of the difference between individual harms and systemic harms is the difference between personal financial losses and economic instability across entire markets. The risk of these systemic harms, paired with the speed of AI development, means that governments should quickly implement basic regulatory infrastructure. Whether this happens as part of AIDA or in separate legislation Canada should create a registry for large AI models that is mandatory to join for all developers of systems over a certain size. Limited disclosures would grant the government and the public basic insights into the actors at play and the risks or legal violations their models might engender.

***Operationalize the desired flexibility, longevity, and balance of* AIDA *through safe harbours and regulatory markets.*** AIDA's goals of flexibility, longevity, and balance between safety and innovation cannot be realized with the government's plan to rely on regulations that will take (at least) two years to develop, which creates substantial uncertainty for stakeholders in a dynamic area. To address this deficit, Canada should implement two low barrier regulatory schemes. First, safe harbours outline time limited guidelines for acceptable AI use such that any entity acting within these bounds is exempt from liability for negative impacts. Second, regulatory markets are private markets for regulatory services in which private third parties licensed by the government compete to provide regulatory services that ensure AI entities achieve government set regulatory metrics. Regulatory markets preserve public control over the direction of AI regulation while bringing private sector efficiency to implementation.

*Focus not on domains but degrees of impact in the definition of "high-impact systems".* The government should abandon its proposed amendment to define high impact systems by sector, which would unduly narrow AIDA's scope, impose stringent requirements where unnecessary, and ignore risky uses of technology. Rather, systems should be defined by their actual level of impact, regardless of sector.

*Promote mutual recognition with trade partners, rather than specific interoperability with the EU.* Although interoperability with the laws of other jurisdictions is important, especially for access to foreign markets, full alignment with EU regulation could impact AIDA's flexibility and abdicate Canadian sovereignty. The government should instead opt for mutual recognition, whereby anything in compliance with one jurisdiction's regulations is deemed compliant with the other's, much like driver's licences.

*Keep the proposed Artificial Intelligence and Data Commissioner as a ministerially-appointed office.* Commissioner decisions will necessarily involve both making political trade offs and balancing between innovation and safety. The pervasiveness of advanced technologies means they will necessarily impact topics of political significance and interest. Independent commissioners tasked with a protective mandate and statute are likely to focus more on harms than innovation, dampening potential positive societal outcomes. Both of these issues can be solved with a commissioner that maintains alignment with the elected government, leading to a more democratic approach.

*Retain* AIDA *'s regulation-focused structure.* Leaving most of the detail in AIDA to regulations was a sound decision in the original bill. Regulations are more responsive and agile than legislation. They can be more technical and specific with help from experts, and they can be amended through a more efficient process. The speed, scale, and complexity of AI demands a flexible approach to its regulation.

I would like to thank you and the committee for your important work on this bill.

Sincerely,

Gillian K. Hadfield

Schwartz Reisman Chair in Technology and Society
Professor of Law and Strategic Management, University of Toronto
AI Chair, Canadian Institute for Advanced Research
Director, Schwartz Reisman Institute for Technology and Society
Senior Fellow, Schmidt Futures AI 2050

### *Recommendation 1: Amend* AIDA *to address the systemic harms of AI and create a registry of large AI models.*

AIDA's focus on individual harms ignores the potential for AI systems to cause **systemic harms**, such as financial instability, uncertainty for investors, and undermining public confidence in elections. The use of increasingly sophisticated AI systems by businesses and the public sector creates a substantial risk of AI causing systemic harms to our economy and society. For example, if left unchecked, the interactions between financial trading algorithms could cause instability in Canadian financial markets. In fact, the Ontario Securities Commission recently released a report on the use of AI in financial markets expressing this very concern. AI systems could also be used to carry out elaborate fraud, undermining investor confidence in the Canadian economy. The well documented phenomenon of generative text applications promoting falsehoods could seriously harm political discourse and undermine public confidence in democratic institutions. Other risks implicate national security, such as elaborate cybersecurity incursions or interference with critical infrastructure.

These systemic harms have recently prompted the White House to issue new requirements for especially large AI systems that could be misused by malicious actors.

As presently drafted, AIDA fails to address these systemic AI harms. The specific harms related offences in the bill are limited to individual harms caused by AI. Although future regulations may address this issue, the magnitude and severity of potential risks calls for addressing the systemic harms of AI in the text of the legislation.

Thus, AIDA should be amended in two ways to address the potential for AI to cause systemic harms. First, the definition of "high impact systems" should be amended to include coverage of AI likely to cause systemic harms regardless of domain.

Second, AIDA should be amended to facilitate the creation of a registry for large AI models. The speed of AI development requires governments to put in place basic regulatory infrastructure as soon as possible. This will allow states to be proactive as they learn more about the risks posed by advanced technologies, as opposed to reacting to negative impacts that have already affected their citizens. Currently, what we know about existing models is entirely a function of what the relevant companies have chosen to disclose. Only the companies building this technology know what they are building or what safety tests they are performing. The public, through its governments, lacks a way to know who is building even more powerful models and to whom those models may be made available.

That's why Canada should move quickly to gather basic information about who is training the most sophisticated models for deployment in Canadian markets a goal that can be accomplished through a simple process of registration. This would furnish the government with basic insights into who is developing models and whether there is substantial risk that their use might violate existing laws and cause systemic harms. The Executive Order issued by the White House on 30 October 2023 takes a step in this direction, requiring developers to share safety tests and their results with government. Canada should follow suit via registration.

Registration is a familiar feature of modern legal systems. Corporations are subject to registration so people and other businesses can have confidence that a company with which they are transacting is not a sham or a front for illegal activity. Broker dealers of securities, companies

handling nuclear materials for civilian purposes, and labs handling dangerous pathogens or toxins all register with the government.

The proposed registration system should be straightforward and make it easy for responsible companies to achieve compliance. It would be tailored to protect intellectual property while enabling countries to forge a better understanding of how the technological frontier is moving. Registration allows governments to ensure that markets and innovation are driven by people and businesses who are following the rules. It draws a line between the more scrupulous actors and those who might be less inclined to comply with a targeted rule and thus merit more careful scrutiny.

Here's how this might work in practice. First, Canada should establish a national registry for large, sophisticated models over a threshold defined by size (number of parameters or amount of compute used for training, for example) and capabilities. Given the dramatic shift in capabilities demonstrated by OpenAI's GPT 4, the threshold should be set near and slightly above the capabilities of this model. Existing federal laws governing export controls, sensitive information, and related matters may allow initial progress toward a prudent registration scheme even in the absence of new legislation.

Second, developers should be legally required to participate in this registry and to disclose confidentially to the registry descriptions of the size, training methods, training data, and known capabilities of these models. The models and the data files used for training the models would not be transferred to the registry. Inadequate or deceptive disclosures should bring substantial penalties including, at a minimum, de registration. The registry should be highly secure to protect against adversarial efforts to hack into the information shared by developers.

Third, Canada should make it unlawful to deploy or use the services of an unregistered model. While only developers of large models need to register, this obligation not to use unregistered models is aimed at both developers and the entities that purchase or use the services of models. If the registration requirement applies, for example, to the next iteration of GPT 4, then it would be unlawful to use that model unless it is registered. Developers could be asked for evidence of registration by users, customers, or service providers.

A registry would allow the government to understand who is working on and responsible for the use of these models. Civil servants would be better able to enforce prudent export control limits, restrictions on the development of biological weapons, or other existing laws. Public officials would also build a nucleus of capacity and expertise to inform further policy on advanced technologies, particularly as systems get closer to acquiring the capacity for self improvement.

The proposed registry has some key features and benefits. First, required disclosures would only be made to governments and under duties of confidentiality that would protect trade secrets appropriately while sharing other information publicly. Second, registration would not necessarily entail any other substantive requirements for the time being; existing laws already prohibit many of the most troubling activities that these models could facilitate. Third, registration is a key step in developing a licensing regime, something policymakers may desire in the future.

Today, the only people who have full information about the scale, training methods, training data, and capabilities of advanced models are those inside the technology companies building them. Even though these companies are by and large mindful of and careful about the risks, it is not

democratically legitimate for this visibility to be exclusively within their purview. Decisions about hugely consequential technologies how fast they roll out, how much they disrupt economies and societies, what is considered a good trade off between benefit and harm, what kinds of tests should be required prior to deployment should not be solely under the exclusive control of even well intentioned business executives who are legally obligated to act only in the interests of their shareholders. Precisely because society will benefit from further innovation and development of advanced models, regulation should start with a basic registration scheme to enable visibility into the development of these technologies and to ensure that prudently designed policies can be carefully targeted.

*Recommendation 2: Operationalize the desired flexibility, longevity, and balance of AIDA through safe harbours and regulatory markets.*

The goals underlying AIDA are laudable. It is supposed to be flexible and agile, to stand the test of time and be technology neutral, and to balance AI safety and innovation. However, AIDA is not receiving the support it needs to achieve those goals.

Regulatory mechanisms do not move quickly. Advanced technologies do. AIDA must be dynamic and capable of responding to rapid advancements in technology. For the most part, AIDA does just that. It is a general piece of legislation that leaves specifics to regulation, allowing the government to respond quickly to a dynamic field in the long term.

The transformative power of advanced technologies cuts both ways. They hold great promise for significant developments in nearly all domains while also harbouring potential harms. The peaceful cohabitation of innovation and safety is a balancing act, but one that must be undertaken to realize the benefits of advanced technologies while ensuring they are not doing harm. AIDA can do both, but only if the government moves quickly and embraces innovative regulatory mechanisms to meet this moment.

The government currently anticipates developing regulations on a two year timeframe once AIDA becomes law. Any benefits from non specific and balanced legislation will be largely undone by this lag in regulation development. Not only does this leave organizations using AI uncertain about the legal status of their actions, it suggests that future updates to these regulations will also move slowly.

Part of the anticipated two year timeline allows for consultations, feedback, and revisions. While this is important, Canada can make itself a leader in AI regulation by implementing two regulatory schemes that will provide AIDA with the flexibility and dynamism it needs and will offer increased certainty and stability to industry and investors. One is safe harbours. The other is regulatory markets.

Safe harbours allow organizations to continue AI use and innovation without facing uncertainty about legal repercussions. They set out specific, time limited guidelines for acceptable AI use and make it clear that, for the time being, organizations following these guidelines or performing to specified standards will not be held responsible. For example, an organization could deploy an algorithmic decision making system and be protected from liability if it is able to show that its system performs with equal accuracy across a defined set of demographic groups. Early regulations could implement this recommendation by identifying cases in which initial safe harbours could be deployed with a high degree of confidence that harms would be largely

prevented. Over time, the boundaries of safe harbours would evolve and organizations that believe the safe harbour is overly conservative could follow a conventional approach to managing the risk of liability for harms.

Some worry the bare bones formulation of AIDA cannot solve the uncertainty currently being faced by industry and investors. However, the combination of AIDA and safe harbours would help mitigate this uncertainty. Again, consider algorithmic discrimination. Without a safe harbour, a company developing an AI model that decides which applicants receive loans currently faces legal uncertainty as to what is expected to mitigate potential algorithmic bias, as it might be judged after the fact by regulators and courts. If Canada were to hold off on AIDA or even follow its current two year regulatory trajectory, the company's options are to refrain from using the model or to move forward and risk repercussions. However, with safe harbours and the promise of regulation to follow, the company could, for example, place its product on the market after undergoing an impact assessment and performance test to meet a government defined metric or after using a specified safety procedure in AI development. Safe harbours give the government time to develop appropriate AI regulations while imposing basic standards to mitigate serious AI risks, creating certainty for AI businesses and maintaining flexibility in the overall regulatory apparatus.

Regulatory markets incorporate the dynamism and efficiency of private markets into the regulatory process to facilitate the necessary flexibility and innovation to regulate AI. They consist of three principal actors: the targets of regulation, licensed private regulators, and governments. Targets are AI businesses and other organizations that governments seek to regulate. Private regulators are for profit and non profit organizations that are licensed to develop and supply regulatory services that they compete to sell to targets. Governments require targets to purchase regulatory services and directly regulate the market for regulatory services, ensuring it operates in the public interest.

Private regulators translate generally worded government regulatory goals into defined processes and technologies. Private regulators could employ, but would not be limited to, conventional means of regulation, like required training procedures enforced by fines and orders. A private regulator might also develop technologies that directly control or shape the business decisions of the targets it regulates. They would gain their authority to regulate via a contract with the target and authorization from governments to collect fines or impose specific requirements on the targets that submit to their regulatory system.

In order to participate in the market, private regulators must be first licensed by the government. To facilitate competition, multiple regulators must be licensed in any given domain. Targets are mandated by the government to choose a regulator, but they have the right to choose which regulator to work with and switch if they are unsatisfied. Presumably, they will do so by comparing across regulators in terms of the cost and efficiency of the services provided by the private regulators.

Government ensures regulations align with the public interest through the licensing system. To obtain and maintain a licence, regulators must demonstrate their regulatory approach achieves government mandated outcomes. The government will develop and implement these outcomes through its typical policymaking and consultation processes. Through this mechanism, the delegation of regulatory oversight of the target to private actors is made legitimate and in alignment with the public interest and political priorities.

The key element of regulatory markets is a division between the policy setting and implementation aspects of regulatory activity. Implementation methods are developed by the private regulators and tested for robustness by governments. Government testing would occur through a combination of upfront evaluation of the capacity for a regulator's system to satisfy government goals and ongoing auditing and oversight. Regulators that fail to pass the tests set by government would risk having their licences suspended, conditioned, or revoked. Private regulators that fail to provide efficient and price sensitive regulatory services to targets will face the competitive pressures and dynamism of market systems, ideally leading to further innovation and technological development of means to regulate AI systems.

We already see startups developing these technologies in the market. To date, they are focused on addressing the demand for risk reduction in areas such as algorithmic discrimination, where existing legal standards and consumer awareness create liability and business risks. Canada could stimulate this domestic market by increasing the demand for such services through licensing and a mandate to purchase these private regulatory services to meet some of the goals of AIDA.

AI is a transformative technology, but we do not have the regulatory levers in place to both unlock its potential and protect society from its harms. To achieve this necessary balance, Canada needs adaptable regulators and smart, pragmatic, and balanced regulation. The capacity of rapid adaptive technology requires us to think in creative ways. Luckily, two novel solutions safe harbours and regulatory markets are actions the government can take now to ensure AIDA realizes its commendable goals. Both lead to lower cost and effective means of demonstrating compliance and enabling the development of technology in an agile way. Rather than stifling innovation, effective technology regulation can promote discovery while ensuring citizens remain protected from risks.

### *Recommendation 3: Focus not on domains but degrees of impact in the definition of "high-impact systems".*

One amendment to AIDA proposed in Minister Champagne's letter of 3 October 2023 defines high impact systems as those with intended uses in particular areas, namely, employment, service provision decisions, biometric information, content moderation, health care, judicial and administrative decisions, and law enforcement. This domain based approach to defining AIDA's scope is both under and over inclusive, limiting the law's effectiveness and impairing its flexibility and balance between safety and innovation. A domain based approach runs the risk of focusing on domains that just happen to have garnered public attention: The EU's *AI Act* draft list of "high risk" applications, for example, includes credit scoring, which has been high profile in the news, but omits housing, which has not drawn much attention.

Under inclusiveness can be addressed by amending the list of domains over time, although this promises to be contentious (especially because of the over inclusion problem) and may bog regulators down. But the risk of over inclusion, which can curb innovation unnecessarily, is harder to fix and a core reason not to use a domain based approach. While advanced technologies are powerful, the response should not be to over regulate. Rather, we should aim for the appropriate level of regulation, a level that promotes the positive outcomes of these powerful technologies while preventing their negative effects. AIDA should not dampen innovation in entire sectors because of the potential for significant negative outcomes in some

contexts. Instead, it should focus on the degree of a system's impact, regardless of its sector of use, attaching requirements only to systems that could, in fact, have a high impact.

For example, consider two uses of AI in education, a domain that was not included in the Minister's proposed list. Incorporating AI into grammar checking and predictive text software could be convenient for students and teachers alike while not drastically changing how schools operate. These uses in education are likely not "high impact". However, replacing teachers or tutors with chatbots could negatively impact a student's development and the adaptability of a classroom, resulting in significant harms, especially if this implementation is unequally applied across groups of students.

Consider next two uses of AI in healthcare, a domain that was listed by the Minister. Using AI systems to manage appointment systems could have an impact on patient care but is not obviously high impact. Using AI to make important decisions about patient care with little oversight from healthcare professionals, however, could well be a high impact application worth significant risk management and oversight.

There is a spectrum of uses for advanced technologies in any sector. AIDA should thus not unduly stifle innovation by placing strict requirements on entire sectors, as this will prevent or slow the development of not only risky uses of AI, but relatively innocuous and helpful uses of these technologies. By focussing on the actual impact of a given system, AIDA can target the uses of AI most likely to have a high impact without restricting innovation in entire sectors and missing risky applications in others.

### *Recommendation 4: Promote mutual recognition with trade partners, rather than specific interoperability with the EU.*

Minister Champagne has proposed in his letter of 3 October 2023 amending AIDA to promote alignment with the *EU AI Act* as well as other advanced economies. Although the text of these amendments remains unavailable at the time of writing, adopting interoperability through legislation raises significant concerns about AIDA's flexibility and the retention of Canadian control over AI regulation. The goal of developing regulatory interoperability across jurisdictions is a laudable one, but a framework of mutual recognition is a solution that achieves the same goals without compromising AIDA's flexibility to the same extent.

Legislative interoperability is the process of modifying the provisions of domestic legislation to match other jurisdictions, thereby allowing domestic companies easier access to foreign markets. Implementing interoperability at the legislative level for AI, however, unnecessarily abdicates Canadian sovereignty to more powerful global actors like the EU. If Canada commits to achieving interoperability with the EU through legislation, it will have to incorporate some elements of EU law into AIDA. Yet, no Canadian policymaker had any substantive input into EU law, which takes a different approach to regulating AI and envisions a vastly different role for AI in its society and economy. The EU's approach is more aggressive, motivated predominantly by mitigating harms from AI, while Canada has taken a more flexible, pro innovation approach, attempting to balance the potential harms from AI with support for the economic and social benefits that AI development can bring. Committing to interoperability with the EU at the legislative level will abdicate Canadian sovereignty over this issue to a foreign jurisdiction with a demonstrably different approach.

Moving beyond sovereignty, legislative interoperability will also severely impair the government's ability to regulate AI flexibly. Legislation is the most complex route by which the government makes law. AIDA itself speaks to this it has been before Parliament for nearly a year and a half at the time of writing. Interoperability at the legislative level would commit Canadian AI regulation to that of any applicable jurisdiction. Given how quickly AI has developed and will continue to develop, committing to legislative interoperability will cause serious issues for the effectiveness of AIDA as a means to regulate AI flexibly.

There is a viable alternative that will achieve regulatory interoperability without infringing Canadian sovereignty or limiting flexibility to the same degree. Mutual recognition is a system whereby anything deemed in compliance with one jurisdiction's regulations is also deemed to comply with the other's. A simple example is driver licensing. Ontario can issue an individual a driver's licence that they can then use to operate a vehicle in any Canadian province or territory, American state, or foreign country like Germany, South Africa, or Brazil. Importantly, they do not need to take a new driving test or secure any paperwork to use their driver's licence outside of Ontario the other jurisdictions have simply recognized that Ontario's licensing system is sufficiently rigorous that anyone who has been licensed under it ought to be deemed a competent driver under their own laws.

Such a system could be put in place for AI regulation. Under it, states could recognize that any AI system deemed compliant with a foreign framework that achieves a sufficient level of rigour and protection is also deemed compliant with their own laws. This solution preserves Canadian sovereignty over AI policy; Canada will not be bound to keep its legislation in strict alignment with the EU. Instead, the government can balance the benefits of interoperability with particular jurisdictions with changes to its own regulatory framework as part of a negotiated process of mutual recognition. This in turn allows Canada to retain control over amending its AI regulations as it sees fit.

A simple new provision could facilitate this process. It would allow the Minister to promulgate regulations designating which jurisdictions have mutual recognition under AIDA, similar to the mechanism by which provincial privacy and data protection legislation can be deemed substantially similar to the *Consumer Privacy Protection Act*.

### *Recommendation 5: Keep the proposed Artificial Intelligence and Data Commissioner as a ministerially-appointed office.*

The government's decision to delegate enforcement of AIDA to a ministerial appointee, the Artificial Intelligence and Data Commissioner (AIDC), has been the subject of criticism that suggests that the AIDC be made an independent officer of Parliament. But preserving the current nature of the AIDC's office will strengthen AI governance in Canada, as the decisions it makes will require inherently political trade offs and a sensitive balance between mitigating AI harms and support for AI development and entrepreneurship.

Although AI regulation may seem highly technical and remote from politics, the trade offs the AIDC will have to make in its decisions have an inherently political nature that would benefit from alignment with the elected government. Many decisions about enforcement and regulation will amount to balancing the economic, social, and political interests of citizens, the kinds of decisions for which input and oversight from the elected government is integral for democratic accountability. Furthermore, the current structure by which the AIDC is appointed will allow for

closer coordination between the government and the AIDC. Given the quickly evolving nature of AI, such close coordination will be an integral aspect of effective AI regulation.

There is also a risk that an independent commissioner will focus mostly on AI harm minimization rather than a balanced, evidence based approach that does not unduly and negatively impact responsible and effective AI users. By maintaining the present structure of the bill, this committee can promote a balanced and democratically accountable regulatory approach that ensures the legitimate interests of AI users are not lost to a hyperfocus on hypothetical AI harms. The risks of AI are real, but where they are not severe, they can and should be managed in a way that helps preserve our ability to reap the benefits of AI.

Thus, the AIDC should remain a ministerial appointee as in the current draft of the bill. Alternatively, should Parliament or the committee decide to make the AIDC an independent officer of Parliament, the commissioner's mandate should emphasize the need for a regulatory approach that balances the risks and benefits of AI use; the Commissioner should be tasked with promoting AI innovation and use as well as harm elimination or reduction as appropriate.

### *Recommendation 6: Retain* AIDA *'s regulation-focused structure.*

The initial draft of AIDA was very general and left much regulatory detail to future regulations. However, in his letter of 3 October 2023, Minister Champagne suggested that this basic element of AIDA's structure may be changed, with much more regulatory detail placed in the text of the legislation itself. This change is not advisable and will likely impair the effectiveness of AIDA as a tool to regulate AI.

The decision to have the original draft of AIDA leave most of the detail to regulation was good for innovation and regulatory flexibility. Regulations can be much more responsive and agile than legislation. Governing AI through regulation has the added benefit of allowing the government to implement more technical and detailed rules developed by expert agencies to address the problems raised by AI in a more technically informed manner. Finally, regulations can be amended through Ministerial directives and orders in council, avoiding the lengthy and inflexible legislative process. If AIDA's structure is changed to place more details in the text of the legislation, it will impair the government's ability to modify the rules to adapt to evolving AI systems.

For these reasons, AIDA should retain its original structure of governing AI through regulation and not imposing requirements through the text of legislative provisions.