March 1, 2024

Standing Committee on Industry and Technology,
House of Commons, Parliament of Canada

cc: The Honourable François-Philippe Champagne, P.C., M.P. Minister of Innovation, Science and Industry

Dear Members of the Standing Committee on Industry and Technology,

Since the introduction of Bill C-27, the Canadian Chamber of Commerce, on behalf of and in consultation with our members from across sectors, including digital technologies, telecommunications, finance, insurance, retail, and manufacturing, among others, has welcomed the opportunity to provide continuous feedback, and engage with the members of the INDU Committee. Given the recent significant changes to Part 3 of the bill, the Artificial Intelligence and Data Act (AIDA) in particular, we are reaching out to share additional feedback on the amendments proposed by Minister Champagne and reiterate the importance of consistent dialogue between INDU, the Government of Canada and industry members at this decisive stage in the legislative process.

As Canada embarks on more digital transformation initiatives to foster a digital economy, it is critical for Bill C-27 not to narrow the window of opportunities that could be harnessed with AI technologies to grow and innovate in Canada. Our members are concerned that the new version of AIDA may unintentionally limit these opportunities. The proposed overly broad regulatory framework without sufficient regard for risk that AI can pose, combined with criminal penalties, unlike any other OECD peer, will disincentivize adoption of AI by Canadian businesses, especially SMEs. This will not only stifle innovation but also contribute to Canada's record of declining productivity. In turn, landing a clear, reasonable, and opportunity-driven legislative model will empower Canadian businesses to invest in the development and application of new technology, eventually bolstering Canada's lagging productivity.

Like Canada's approach on privacy reform, we urge parliamentarians to consider a legislative approach to AI that strikes an appropriate balance between safeguarding the public and fostering innovation and growth for the economic and social benefit of Canadians. It will also be important to provide adequate certainty for business in the legislation itself, and not leave scoping and important details to future regulation. Businesses thrive on certainty for which regulatory clarity and stability are a must to achieve.

The Canadian Chamber of Commerce's Future of AI Council continues to advocate for technologically neutral and risk- and principles-based legislation grounded in four primary tenants:

- Interoperability
- Standardization
- Legal and regulatory harmonization
- Open dialogue

As a follow up to our letters to the Committee dated November 23, 2023 and February 13, 2024, members of the Canadian Chamber of Commerce respectfully submit the attached comments on proposed amendments to AIDA in Annex A, to the Government of Canada and INDU Committee members to help strengthen AIDA based on industry best practices.

Kind regards,

Ulrike Bahr-Gedalia
Senior Director, Digital Economy, Technology & Innovation
Canadian Chamber of Commerce
C: 613.410.6629 E: ubahr-gedalia@chamber.ca

Annex: Artificial Intelligence and Data Act (AIDA) Amendment Feedback

Scope

It is critical that the bill align with other international definitions and focuses on the outcomes, as in the specific risks and impact on individuals, rather than the data processed by an AI system. From that perspective we would encourage ISED to remove references to "another technique" in the amended definition and also to more closely align with the OECD definition. Potential language to that effect could be: "***artificial intelligence system*** means a technological system that, using a model, makes inferences in order to generate output, including predictions, recommendations or decisions."

"High-Impact" Systems

It is important that the bill itself include a definition of "high-impact" systems to provide businesses with certainty regarding their obligations and entrench the core of the regulatory scheme in the law. While we appreciate that the Government has acknowledged this and provided a set of classes with criteria for defining "high-impact systems," our members have several concerns with the proposed framework. Instead of establishing a structured framework akin to the EU's AI Act for determining what constitutes high impact, the current approach relies on examples of intended use cases. The proposed Schedule 2 creates an overly broad classification system that encompasses scenarios not truly deemed high impact, especially since the list of high-impact uses is not linked to any associated harms or materiality threshold.

At a high level, we believe AI regulation should be risk-based and proportionate, focusing on the most sensitive types of AI applications and sectors. Several of the criteria proposed miss this mark and instead rely on overly broad language and concepts that would capture many low-risk AI systems. An example of an overly broad definition can be found in employment-related determinations; a determination of this nature should be limited to AI systems that make decisions that are not subject to human review (versus making a recommendation or prediction that is accepted or rejected by a human) and to decisions that have a "material legal or other similarly significant effect". Leaving it to future regulations to further specify these categories of high impact systems will likely have the unintended effect of dissuading many Canadian enterprises from adopting this transformative technology. While the Annex B to Minister Champagne's Letter to INDU dated November 28, 2023, provides helpful guidance, this guidance has no legal effect and does not constrain what can be regulated. We would therefore strongly encourage the Government to draft a much more precise set of high impact scenarios. Rather than further specify these in regulations, the government should rely on regulations to add additional precise high impact scenarios where necessary.

AIDA's high-impact system amendments also raise issues related to interoperability, especially with our closest trading partners. For example, as proposed, the scope of AI systems regulated under AIDA as high-impact is materially broader than "high-risk" AI systems classified under the EU's AI Act. This means that far more low-risk uses of AI in Canada will be subject to far more requirements than the EU. Ultimately, we are concerned that the framework envisioned will create

a disproportionate regulatory approach that severely burdens innovation in Canada. A principles-based framework should limit the scope to high risk arising in the B2C context impacting consumers directly.

Recommendation: Reframe definition of high-impact to focus on consequential uses like those that have a material or legal impact (or similarly significant harm) to housing, employment, credit, education, healthcare, criminal justice or insurance.

Recommendation: Add a consideration to the new section 36.1 requiring the regulator to consider economic impact, trade impact and innovation impact when modifying or adding new high-impact categories.

In addition to broad concerns with the definition of high-impact provided in AIDA amendments, we would like to note the following issues with the proposed classes:

Use of AI in the Provision of Services: Use of the term "services" to describe this proposed class of AI systems is incredibly broad and in need of clarification. This could potentially cover everyday consumer interactions, such as home maintenance or entertainment, or it could concern more sensitive areas, such as legal or employment services. Such a range of scenarios presents varying levels of risk, and the bill's "high-impact" framework should be tailored accordingly to ensure proportionate regulation that does not burden the ability of the technology to deliver consumer benefits.

Recommendation: Clarify that the proposed class of "high-impact" AI systems related to the provision of services to an individual, including whether to provide services, determining the type and cost of services, and the prioritization of services to be provided, is in reference to government services.

Use of AI Systems to Process Biometric Information: Expanding this obligation to the processing of biometric information in matters related to an "individual's behaviour or state of mind" is an incredibly vague standard. Nearly any information about an individual could be thought of as relating to their behaviour or state of mind. For example, there are popular extended reality, gaming, and wearable products that rely on data about the body to facilitate basic product functionality. Hand tracking technology could be implicated even if data about the position of an individual's hands is used only to enable them to interact with an AR/VR interface or operate a smartwatch, as how they move their hands could be considered a behaviour.

Similarly, an in-device camera may analyze an individual's mouth movements to render and animate their avatar—not to learn any information about their emotional state – yet such movements might be considered to relate to their state of mind. If 3(b) were limited to AI systems actually used to identify emotional states, it would be appropriately to the potentially risky and sensitive uses of data about the body, without also capturing low-risk uses that often provide essential device functionalities expected by consumers.

Recommendation: Limit language referring to "matters relating to . . . an individual's behaviour or state of mind" to the use of an AI system that processes data about the body for the purpose of identifying an individual's emotional state.

Use of AI Systems for Content Moderation or Prioritization: The proposed inclusion of the use of AI systems to moderate or prioritize content that is found on an online communications platform, including a search engine or social media service, is very broad and not coherent with the potential harmful impacts related to other types of AI systems listed. Further, automated systems used for content moderation, or the prioritization of content have not been categorized as "high-impact" AI in any other jurisdiction. The other proposed high-impact AI systems are used in areas such as employment, identification of an individual, health care services, administrative decisions about an individual and assistance of peace officers. Automated systems used for "content moderation" or "prioritization of content" are not comparable to those other high-impact automated decisions. High-impact decisions are made with much less frequency and create a direct impact on an individual's life. We would also note that general data protection laws already regulate organic content prioritization or personalization. We therefore would encourage a careful assessment of the needs to regulate specific aspects of automated systems used for the delivery of personalized content to avoid duplicative and burdensome regulations that lead to confusion and often a poor experience for consumers.

We are also concerned about the unintended consequences of such a broad definition, particularly around the definition of "online communications platform". This could be interpreted as an internal online tool for the organization. There are many different types of communication platforms available, each with its own set of features and capabilities. That being the case, a line of business application could easily be considered an online communications platform and this would put any generative AI that is geared at helping businesses better prioritize data and insights for customers into the high impact category.

Recommendation*:* Remove the use of AI systems for content moderation or prioritization from the list of classes included in the proposed amendment and instead maintain criteria related to systems that make decisions regarding matters of consequence to an individual's life or access to basic necessities.

Obligations for General Purpose AI Systems

The risks associated with General Purpose AI systems are dependent on the context in which they are deployed. This means, the necessary risk identification and mitigation strategies that need to be developed would also depend on the context, and model developers are unlikely to be able to reasonably foresee (and address) all possible risks in the multitude of downstream applications of General Purpose AI systems. It is unnecessary, therefore, to introduce distinct requirements for providers of General Purpose AI systems independent of any context specific use considerations. Deployers of such systems (i.e., downstream users) are best positioned to comply with the obligations that arise from their use in high-impact scenarios. The government should ensure AIDA does not apply impractical and unworkable obligations on upstream providers

but provide distinct obligations based on their role in the supply chain. For example, a business might integrate use of a large language model (LLM)-based chatbot in its customer service interactions or leverage it for marketing to prospective customers. In this case, absent any contractual arrangement, the developer of the LLM would not have insight into how it was being used in these specific business contexts and any risks related to that use, and therefore, would not be able to meet its regulatory obligations under the Act. AI systems which can complete multiple, distinct tasks are relatively nascent, even in the field of AI, which is itself an emerging technology. We have seen such systems described variably as general-purpose AI and foundation models. There remains a lack of consensus about how to name and define these kinds of systems. In other words, it is still early, so it is important to approach regulation in a way that recognizes that we don't have the full picture of AI and its applications yet.

Instead of focusing on an approach to regulate General Purpose AI systems – based on a current understanding of model capabilities – AIDA should adopt a technologically neutral approach that focuses on risk-of harm and functional effects of this technology. Like all other areas of regulation, AI laws and regulations, including those applying to different types of AI systems/model development and deployment, should be designed such that they develop independently of any specific technology and should continue to apply equally across technologies as they emerge, without favouring or discriminating between new and old. Such an approach presents the promise of sustainable laws in a time of rapid technological change, thereby, "future-proofing" the AI regulation regime to some degree such that the same laws continue to apply to new and emerging AI systems and models. Such an approach produces more comprehensible legislation for business and the public who are expected to abide by it.

The government's amendments also fail to make important distinctions between "systems" and the "models" on which they are built. In particular, the GPAI regime proposed in the amendments relies on the term "general-purpose system," but the definition describes a GPAI model (i.e., an AI system that is "adapted for use"). The amendments should be revised to address this confusion.

The government's current approach to General Purpose AI systems is significantly out of step with some of its largest trading partners like the US and the EU, and raise concerns about over-regulation of GPAI. For example, even in the EU which currently has the most stringent and prescriptive AI regulatory framework globally, the AI Act distinguishes requirements for General Purpose AI systems according to risk. This allows a focusing of resources and requirements where the risk is highest, rather than having the same requirements for all General Purpose AI systems, regardless of risk.

Rather than a rush to create comprehensive regulation based on ambiguous concepts, it would be preferable to allow time for more concrete definitions to evolve that reflect consensus among key stakeholders. If regulatory obligations are still pursued despite the nascency of these issues, they should be consistent with a risk- based approach to ensure they provide certainty and are tailored to address discrete harms.

Recommendation: Narrow any distinct obligations imposed on General Purpose AI Systems to the most high-risk types of uses, such as use of the technology to produce a reasonably foreseeable legal or similarly significant effect on an individual. Address ambiguity in key terms by creating distinct definitions for 'general-purpose systems' and 'general-purpose models.'

Alignment with International Norms and Global Leadership in Areas of Canadian Comparative Advantage

We support efforts to ensure any federal AI framework is consistent with international norms to avoid a patchwork approach that adopts some aspects of certain frameworks. Definitions and scope of responsibilities should be consistent to prioritize interoperability.

There are terms and implementation frameworks proposed in the amendments that depart from comparable international frameworks and raise interoperability concerns. While we recognize the improvements proposed to the definition of AI, further changes are needed to align the bill with the OECD definition. In particular, the definition should build from OECD and center its definition around software that is machine-based and able to learn over time. Such an approach strikes the right balance between focusing-in on the novel risks certain AI systems may present with an overly broad approach that would capture *any* software. Moreover, the term 'machine-learning model' is not referenced within the definition of AI, general-purpose system, or high-impact system, and functions as superfluous to the high-impact and GPAI frameworks set forth in the bill. Inclusion of this term and related obligations is confusing and inconsistent with the creation of a clear taxonomy of actors on which the bill's obligations should be based. The term should therefore be excluded from the bill. The amendments also apply high-impact system operation obligations retroactively once the regulations are in force. This is an overly broad provision for implementing high-impact system obligations that is onerous and beyond what comparable laws internationally (i.e., E.U. AI Act) have contemplated. This amendment should be excluded from the bill.

Given work happening on AI at a global level, we encourage Canada to avoid rushing ahead with any legislation or regulation which could put us out of step with major trading partners, discourage future investment and innovation, or unintendedly harm the early advances of our globally renowned AI ecosystem.

Canada can play a leadership role internationally where it has a comparative advantage. For example, Canada can lead in developing ways to effectively assess and address potential risks related to the safe and secure use of AI in critical infrastructure sectors, including preventing ways in which deploying AI could make critical infrastructure systems more vulnerable to cyber attacks, and consider ways to mitigate these vulnerabilities. To that end, we encourage the government to support research and development efforts and investment in Canada specifically to help prevent the malicious use of AI systems to enable offensive cyber operations through, for instance, automated vulnerability discovery and exploitation. Canada's goal should be to position itself as a global leader in preventing the malicious use of powerful AI systems to perpetrate cyber attacks or other malicious activity that could cause widespread disruption and destruction.

Recommendation: Insert regulatory guidance or principles directly within AIDA that specify forthcoming regulations must align with international AI-governance frameworks and standards and promote interoperability for Canadian companies abroad. Remove ambiguous terms and unworkable obligations proposed in the amendments.

Criminal Liability and Intentional Harm

Absent among the proposed amendments is much needed clarification on the scope of criminal liability in Bill C-27. Canada is the only country that has sought to include criminal liability in its legislation and is a significant departure from international AI governance norms. To clarify its use for intentional and egregious use of AI systems for physical harm or serious fraud will minimize perceived risk and allow Canada to lead on addressing the use of AI systems for intentional harm. Without such clarifications, the increased risk to operate in Canada may discourage investment in Canada's global-leading AI.

Recommendation: Clarify use of Criminal Liability for intentional and egregious use of AI systems for physical harm or serious fraud will minimize perceived risk and allow Canada to lead on addressing the use of AI systems for intentional harm.

AI and Data Commissioner

Audit rights of Commissioner are overly broad and should focus on necessary information for compliance. For technology companies heavily reliant on AI, the information requests alone within the current legislation could be crippling. Orders should be technically feasible and in line with industry standards, the least restrictive means of demonstrating compliance, flexible to allow for alignment with business practices. The Commissioner should also be independent of ISED.

The amendments also propose a remote access provision, which would grant the government broad authority to remotely access company's IT systems, including companies based outside of Canada. This raises significant corporate espionage and cybersecurity concerns. Regardless of the Canadian government's intent for this legal power, it would set a worrying precedent internationally for governments overseas if passed into law by Canada, particularly with respect to governments that have troubling human rights records. These countries may introduce copycat legal powers with a different intent behind their use, such as to increase law enforcement access to user data. This regulation raises the prospect for the Canadian regulator to become a means of a backdoor access to overseas data by Canadian Intelligence agencies – circumventing the standards and processes upheld by international agreements on cross border data access to which it is a party.

This remote access requirement would be even more alarming if it enables access to user data, and if that access is in real time since this would severely exacerbate risks to user privacy and system integrity. If this remote access provision is not removed, it should be limited to exclude access to user data or source code and prohibit real time access.

This is a departure from normal regulatory powers. Traditionally regulators use information gathering powers such as requests for information, etc., then progress to investigation and/or audit of a service provider to draw conclusions as to compliance. This bypasses that long-standing model without justification. Furthermore, it has been proposed without any justification of its necessity, or safeguards to ensure proportionality and protect against abuse.

Recommendation: Limit the scope of the Commissioner's audit powers on necessary information for compliance. Remove remote access authority altogether. The Commissioner's powers should be organized similar to the Competition Bureau.

Delegation: AIDA should allow for sectoral regulators to regulate the use of AI in their sector. Sectoral regulators, whether in financial services, transportation, health or others, are closest to their industries and are familiar with the ways in which they are already using AI and will be using AI in future. These regulators are best equipped to understand sector-specific issues. Leaving regulators with strong knowledge of their industries to regulate AI in their industries would ensure greater agility and innovation in these sectors.

We continue to believe that AIDA is fundamentally problematic in that it relies on a single Minister to identify high impact use cases and to develop corresponding regulations. AIDA should instead ensure that all relevant Ministers are delegated to review the use cases under their normal purview and to determine whether they are high impact – based upon principles developed within AIDA and modified by the Commissioner as needed. This would ensure that regulators are much closer to, and much more familiar with, the use cases they intend to regulate. It also spreads responsibility amongst a greater number of bureaucratic staff and builds necessary expertise within departments. The Commissioner and accompanying staff could then support the relevant Ministers and departments with deep AI policy and standards capacity, and ensure harmonization where desirable.

Recommendation: Insert powers for other Ministers who may be required to regulate AI on a sectoral basis (i.e. Minister of Finance, Minister of Transport). Further, AIDA should contain principles to guide the making of regulations. These principles should stress the importance of decentralized responsibility for regulation by the relevant Ministers.

Transparency Obligations

Transparency obligations will be a critical part of the new AI framework. However, there are important exceptions that will be required so that firms are able to use AI in important areas meant to also protect consumers and Canadians. We would suggest that the requirements to public disclosures by the relevant actors under AIDA be amended to include appropriate exceptions to ensure that firms are not compelled to disclose proprietary or sensitive information, which will ultimately protect the public and organizations from avoidable harm that can potentially be caused by broad disclosures of risks or mitigation measures. Certain information related to certain AI systems that is currently captured by the transparency obligations could, for example, be used by actors that pose a cyber risk to firms or fraud risk to customers.

This exception should include the obligation on persons responsible under the Act to disclose prescribed information; as well as the right to disclose information by i) the Minister or Commissioner to others and ii) the right to publish information by the Minister or the Commissioner.

From this perspective we recommend removing the requirement on organizations to publicly disclose mitigation measures and risks under s.11(1)(f)(iii) (s.11 of the original Act) and 7(1)(f)(ii) and s.8(1)(a) respectively, as this could result in the disclosure of sensitive or otherwise confidential information of an organization or of a third-party, and additionally introduce risk to the organization or the public (e.g. compromising the efficacy of AI systems, exposing exploitable information related to critical systems) that may not be balanced with the benefits such a disclosure is intended to provide.

More generally, if the proposed disclosure requirements under s.11(1)(f) (s.11(1) of the original Act) and s.8(1)(a) are adopted, an exception should be introduced to provide greater certainty that organizations will not be compelled to reveal confidential business information or information regarding sensitive systems as this would result in greater (and avoidable) risk for organizations and the public (e.g. systems related to security or used to prevent financial crime)[1].

To strike a better balance between the intent of AIDA to enhance the transparent use of AI systems with the need to protect organizations' confidential and sensitive information, we recommend that a general account of AI systems, including high-impact and general-purpose AI systems, similar to the provisions outlined in CPPA's s.62 (2), would provide an appropriate level of detail to the public without divulging system specific information. Incremental information, or detailed information on specific AI systems, could be provided to the Minister or Commissioner upon request or through audits.

For (s. 11(1)(f)) we believe the following language would meet the objectives of AIDA while the previously outlined risks: *A person who manages the operations of a high-impact system must in the time and manner prescribed by regulation, publish on a publicly available website a plain-language description of the system that includes the following information:*

(i) how the system is being used,
(ii) the types of output that it generates,
(iii) the mitigation measures established under paragraph (b) in respect of it, and
(iv) any other information that may be prescribed by regulation.

Regarding article (s. 7(1)(f)) we would suggest the following language: *Before a general-purpose system is made available in the course of international or interprovincial trade and commerce for the first time, the person who makes it available for that first time must ensure that*

(i) a plain-language description has been prepared of
(ii) the system's capabilities and limitations
(ii) the risks of harm or biased output referred to in paragraph (c), and
(iii) any other information prescribed by regulation.

[1] S.11(1) of the original text of AIDA refers to requirements that are now listed under S.11(1)(f) of the proposed amendments.

Lastly, we believe the following language should be used to amend: *(s. 8(1)(a)): A person who makes a general-purpose system available must:*

> *(a) make the plain-language description referred to in paragraph 7(1)(f) available to users of the system or, if the system is made available to the public, publish that description, in the time and manner that may be prescribed by regulation, on a publicly available website; and*
>
> *(b) take any measures prescribed by regulation.*

Publication Without Consent

S.28(1) currently permits publication of information by the Minister without consent or notification to the person to whom the information relates. An obligation should be added under this section to ensure that the Minister notifies and consults with the impacted organisation to ensure that the publication does not lead to the release of information that is competitively sensitive, highly confidential, or both.

Implementation Timeline

Due to the complexity of this issue and the importance of the consultation phases in developing the regulations we believe that the legislation should provide a 24-month timeline for implementation starting when a significant portion of the regulations have been developed.