

Policy Brief

Canadians' Perspectives on Governing AI Chatbots

February 2026

Executive Summary

This report adds to an evidence-based mandate for AI chatbot regulation in Canada, drawing on nationally representative survey data from 1,454 Canadians and mapping public preferences onto established international regulatory frameworks. The findings demonstrate overwhelming public concern across all surveyed risk categories and clear attribution of accountability to AI companies, providing democratic legitimacy for mandatory intervention.

Survey data reveal that over 70% of Canadians express concern about every major surveyed chatbot risk, including: exposure to misinformation (81%), privacy violations (79%), impacts on critical thinking (79%), mental health deterioration (78%), learning deficits (75%), emotional dependency (75%), and harmful behaviors (73%). Most significantly, 56% of respondents hold AI companies primarily responsible when youth are harmed—far exceeding attribution to parents (12%), youth themselves (9%), or government (8%).

This assignment of corporate accountability establishes clear public expectations for governance. 69% want stricter regulations, while 74-80% support specific safeguards including age verification, parental oversight, data limits, and content restrictions. These policy preferences could be accounted for through a series of coordinated interventions: safety-by-design obligations in AI legislation, strengthening privacy protections through rights-based reform, implementing proportionate age assurance mechanisms, and establishing independent oversight with enforceable transparency requirements.

Sommaire exécutif

Ce rapport met en évidence, à partir de données probantes, la nécessité de réglementer les chatbots d'intelligence artificielle (IA) au Canada. Il s'appuie sur les résultats d'un sondage représentatif pancanadien mené auprès de 1 454 Canadiens et compare les préférences du public aux cadres réglementaires internationaux existants. Les résultats montrent une forte inquiétude du public face à l'ensemble des risques évalués. Ils indiquent aussi que les Canadiens en attribuent clairement la responsabilité aux entreprises d'IA, ce qui renforce la légitimité démocratique d'une intervention réglementaire obligatoire.

Les données du sondage indiquent que plus de 70 % des Canadiens se disent préoccupés par chacun des principaux risques associés aux chatbots qui ont été sondés, notamment : l'exposition à la désinformation (81 %), les atteintes à la vie privée (79 %), les impacts sur la pensée critique (79 %), la détérioration de la santé mentale (78 %), les difficultés d'apprentissage (75 %), la dépendance émotionnelle (75 %) et les comportements nuisibles (73 %). Plus significatif encore, 56 % des répondants estiment que les entreprises d'IA sont les premières responsables lorsque des jeunes subissent des préjudices — un chiffre largement supérieur à l'attribution de responsabilité aux parents (12 %), aux jeunes eux-mêmes (9 %) ou au gouvernement (8 %).

Cette attribution de la responsabilité aux entreprises reflète des attentes publiques claires en matière de gouvernance. Ainsi, 69 % des répondants souhaitent des règles plus strictes, tandis que 74 à 80 % appuient des mesures de protection précises, notamment la vérification de l'âge, la supervision parentale, des limites à la collecte et à l'utilisation des données, ainsi que des restrictions sur certains contenus. Ces préférences pourraient être prises en compte au moyen d'une série d'interventions coordonnées : l'intégration d'obligations de sécurité dès la conception dans la législation sur l'IA, le renforcement des protections de la vie privée par une réforme fondée sur les droits, la mise en place de mécanismes d'assurance de l'âge proportionnés, ainsi que l'établissement de mécanismes indépendants de supervision appuyés par des exigences de transparence contraignantes.

Introduction

AI chatbots, including AI companions,¹ have become an increasingly important part of young people's daily lives. Unlike traditional digital platforms, conversational AI systems provide personalized, real-time responses that can feel authoritative or emotionally attuned and often blur boundaries between information provision, advice, and social interaction. These affordances are what make them so popular: a recent survey indicates that over 70% of American teenagers have used AI companions, with more than 50% engaging with them on a regular basis. Meanwhile, experts, civil society, and journalists increasingly raise flags about the risks associated with their use, including the provision of inaccurate or harmful information, privacy challenges, and developmental and mental health effects. These risks are particularly acute for minors. In fact, recent reported incidents show that, in extreme cases, AI chatbots have steered youth users toward violence and suicide, leading to legal challenges and public pressure for stronger protections.

It is within this context that the Centre for Media, Technology, and Democracy's [Media Ecosystem Observatory](#) launched a national survey to better understand Canadians' views on AI chatbots and governance mechanisms to regulate them. Fielded between November 28 and December 7, 2025, the nationally representative survey of 1,454² Canadians sought to answer three questions:

1. To what extent do Canadians perceive risks in young people's use of AI chatbots?
2. Who do Canadians believe should be held responsible for harm caused by AI chatbots?
3. What measures do Canadians believe should be put in place to protect youth using AI chatbots?

Drawing on findings from this survey, and the Canadian and international policy context for AI chatbots, this report offers best practices for AI chatbot governance in Canada.

Findings

Canadian digital policy has historically approached online harms through content-based regulation: harmful material is produced, distributed, and consumed, and regulation intervenes at one or more of these points. This paradigm underpins existing approaches to broadcasting, platform moderation, and online safety. AI chatbots complicate this model because they: 1) generate novel outputs in real time; 2) adapt to user inputs over repeated interaction; and 3) simulate responsiveness, empathy, and social presence.

Public concern suggests that Canadians recognize this shift. More than 70% of respondents report being somewhat or very concerned about

Figure 1.
Concern about the
effects of AI chatbots





Figure 2.
Responsibility for
harm caused by AI
chatbots

If a youth is harmed by inaccurate or harmful chatbot content, who should be held most responsible?



a wide range of potential harms associated with chatbot use by young people. Notably, concern exceeds 70% for risks related to exposure to misinformation (81%), impacts on critical thinking (79%), mental wellbeing and social isolation (78%), emotional dependence (75%), encouragement of risky behaviour including self-harm (73%), and privacy and data collection (79%).

These concerns fall into three analytically distinct categories:

- 1. Epistemic Risks:** Canadians are concerned about inaccurate or misleading information in domains with immediate personal consequences.
- 2. Social risks:** Canadians express strong concern that frequent reliance on AI chatbots may interfere with youth cognitive development, emotional regulation, and social skill formation. 67% agree that frequent chatbot interactions negatively interfere with emotional and social development, while only 6% disagree - an asymmetry that suggests directional consensus. Likewise, concern about emotional dependence, reported by nearly three quarters of respondents, indicates unease with chatbot features that encourage sustained engagement, anthropomorphism, or perceived intimacy. Public concern thus appears directed less at isolated harms than at design incentives misaligned with youth wellbeing.
- 3. Structural and data-related risks:** High concern regarding privacy and personal data collection (79%) indicates public awareness that chatbot risks are embedded in data practices and system architectures, not solely in outputs.

Responses and Responsibility

With clear awareness of the risks and the responsibility AI companies bear for their products, Canadians overwhelmingly support stronger safeguards governing their design, use, and oversight. 69% of Canadians advocate for stricter rules for how AI chatbots interact with youth, compared with only 6% who support loosening them.

To do so, Canadians indicate strongest favour for opt-out mechanisms and limits on data collection. By contrast, support is lower for downstream or complementary measures, including AI literacy education (62%) and the development of youth-specific chatbot platforms (46%).

These findings expose tensions within existing policy approaches for online harms:


- 1. Content versus system regulation:** public concern focuses on interaction dynamics, learning impacts, and emotional dependency—risks not adequately addressed by content-centric frameworks.
- 2. User responsibility versus design accountability:** responsibility attribution overwhelmingly favours AI companies, challenging regulatory and governance models that emphasize individual or parental responsibility.
- 3. Voluntary/self-regulatory commitments versus enforceable policy:** public support for safeguards raises questions about the legitimacy of relying on voluntary industry measures.

Policy Pathways

Taken together, this supports consideration of three distinct governance mechanisms:

Figure 3.
Support for measures
to protect youth using
AI chatbots

Do you support or oppose the following measures being put in place to protect youth using AI chatbots?



1. Safety-by-design obligations, which could be implemented in the federal government's likely forthcoming Online Harms Act;
2. Amended privacy protections including extended privileged protections for children's data, which could be implemented in the likely forthcoming amendments to the [Personal Information Protection and Electronic Documents Act](#); and,
3. Independent oversight with enforceable transparency mechanisms, which had been called for in the federal government's previous draft of the Online Harms Act.

1. Safety-by-Design Obligations

Responding requires moving intervention upstream by imposing statutory duty of care on developers, modeled on the UK's [Online Safety Act](#), EU's [Digital Services](#) and [AI Acts](#), and proposed U.S. legislation.³ This duty could require: mandatory and ongoing risk assessments identifying foreseeable harms before deployment and as systems evolve; default settings configured to the highest privacy and safety standards for users under 18; content filters preventing advice on self-harm or violence subject to regular circumvention and stress testing; and restrictions on features exploiting developmental vulnerabilities, including engagement-maximizing personalization, dependency formation, or simulated intimate relationships.

2. Privacy Protections

Privacy reform must move beyond consent to structural protections. Drawing on [California's Age-Appropriate Design Code](#) and the EU's [General Data Protection Regulation](#), federal legislation should: classify all data from users under 18 as categorically sensitive; mandate privacy-by-default with minimized collection and limited retention; prohibit using minors' data for profiling, advertising, or training AI; require transparency reports disclosing data practices; and strengthen Privacy Commissioner enforcement through proactive audits, binding orders, and revenue-scaled penalties. These reforms would fundamentally constrain data practices enabling manipulation, responding directly to the 80% of Canadians demanding data limits and operationalizing the 69% supporting stricter rule-setting by establishing enforceable baseline protections that cannot be negotiated away through terms of service.

3. Independent Oversight and Transparency

Effective oversight requires regulatory capacity and mandatory transparency. Drawing on the approach taken in the EU Digital Services Act, Canada should require developers to: publish risk assessments and mitigation strategies in accessible formats; provide regulators and accredited researchers with access to training data, model documentation, and anonymized interaction logs under confidentiality; submit

independent audits with public reports; maintain incident reporting documenting serious harms with mandatory regulatory notification; and establish accessible user reporting mechanisms with documented response protocols. These activities should be monitored and reviewed by a federally-established independent commission with dedicated expertise and statutory authority. The commission could include an ombudsman office empowered to investigate complaints, issue compliance orders, and enforce meaningful penalties.

Conclusion

This report demonstrates three foundational realities. First, Canadians possess sophisticated understanding of the risks that AI chatbots pose to young people, with over 70% concerned across all surveyed harm categories and clear recognition that risks are most acute where adolescents increasingly rely on AI guidance: mental health, social relationships, and emotional wellbeing. Second, Canadians assign clear responsibility to entities best positioned to mitigate these risks: the companies that design, deploy, and profit from these systems. Fifty-six percent hold AI companies primarily accountable, providing unambiguous democratic mandate for mandatory corporate obligations. Third, Canadians support specific, operationalizable interventions mapping onto proven regulatory frameworks: 69% want stricter rules, 74-80% back data limits and content restrictions, and 60% support independent oversight.

Canada possesses the institutional capacity to respond. Privacy law can protect minors' data by default and constrain exploitative profiling. AI legislation can impose safety-by-design obligations requiring that developers identify and mitigate foreseeable risks before deployment. Independent oversight can be established by empowering regulators with audit authority, transparency requirements, and meaningful enforcement tools.

Endnotes

1 AI chatbots are online platforms designed to simulate human conversations using artificial intelligence (AI). Some are specifically designed to serve as companions, designed to simulate social conversations and provide personal support.

2 The data was sampled and weighted by age, gender, and province, according to population estimates provided by Statistics Canada.

3 [California's Age Appropriate Design Code](#) established the foundational framework adopted by subsequent state legislation. Connecticut prohibits endless scrolling; [Maryland](#) requires guardian monitoring practices and best-interest design; [Nebraska](#) mandates chronological feeds, nighttime/school-day notification pauses, and time-limiting options; [New Mexico](#) would disable unknown user contact and 10 PM–6 AM notifications; [Vermont](#) focuses on reducing addictive and harmful design features.

Contributors

Mathieu Lavigne, Ph.D.

Mathieu Lavigne is the Analytical Lead of the Media Ecosystem Observatory.

Helen A. Hayes

Helen A. Hayes is the Associate Director (Policy) of the Centre for Media, Technology, and Democracy.

Esli Chan

Esli Chan is a senior research fellow at the Media Ecosystem Observatory.

Christopher Ross

Christopher Ross is a senior analyst at the Media Ecosystem Observatory.

The work of the Media Ecosystem Observatory and the Canadian Digital Media Research Network is funded in part by Heritage Canada's Digital Citizen Initiative.