COGNITIVE SECURITY
SÉCURITÉ COGNITIVE
CANADA

Perceive • Percevoir | Analyze • Analyser | Integrate • Intégrer

RESEARCH REPORT

Cognitive Liberty in Canada: Protecting Mental Autonomy in the Digital Age

A public-interest research report on cognitive autonomy, digital influence, and governance gaps in Canadian law and policy

Report number	COGSECCAN RR-001-EN
Publication date	May 2026
Prepared by	Cognitive Security Canada
Information class	Public-interest research

Document Control

Document type	Research Report
Report number	COGSECCAN RR-001-EN
Title	Cognitive Liberty in Canada: Protecting Mental Autonomy in the Digital Age
Publication status	Public Release
Prepared by	Cognitive Security Canada
Version	v1.0-EN
Notes	Research Report Master Template

Executive Summary

Canadian law does not yet contain an explicit, operational protection for cognitive liberty: the ability of a person to think, interpret, decide, and process information without coercive, hidden, or technologically mediated interference. The Canadian Charter of Rights and Freedoms refers to freedom of thought, belief, opinion, and expression under section 2(b), but Canadian doctrine has largely developed around expression rather than the inner conditions required for independent judgment. Section 7 protects life, liberty, and security of the person, including serious psychological integrity in some contexts, but it remains tied to state action and a high threshold of direct harm. These legal doctrines were built before algorithmic recommendation systems, psychographic advertising, dark-pattern design, generative AI persuasion, and neurotechnology became ordinary features of public and private life [1][2].

The central public-interest problem is not that persuasion exists. Democratic societies depend on persuasion, advocacy, debate, education, and free expression. The problem arises when the architecture of influence becomes opaque, adaptive, data-driven, and asymmetrical: when systems are designed to infer vulnerabilities, shape perception, narrow available choices, and steer behaviour without meaningful awareness or consent. This report therefore treats cognitive liberty as a governance issue at the intersection of law, privacy, consumer protection, platform accountability, AI safety, human rights, and civic resilience.

Current Canadian legal tools only partially address this terrain. Privacy law regulates the collection, use, and disclosure of personal information, but it does not squarely regulate the cognitive effects of using that information to influence belief or behaviour. Competition and consumer protection laws can respond to false or misleading representations, hidden fees, and some unfair design practices, but they do not yet provide a clear remedy for subtle influence that is not technically false. Contract and tort doctrines recognize capacity, consent, negligence, undue influence, and intentional mental harm, but those doctrines are poorly adapted to diffuse, mass-scale algorithmic persuasion. As a result, the person affected by manipulative digital

systems may experience real loss of autonomy while still lacking a clear legal category for the injury [3][5][6].

International developments show that this gap is increasingly visible. Policy discussions on freedom of thought, mental privacy, and neuro-rights have emerged in the United Kingdom, Australia, the European Union, and international human-rights scholarship. These developments do not provide a complete model for Canada, but they show a growing recognition that mental autonomy and cognitive integrity need practical governance mechanisms in the digital age [4][7][8].

This report recommends that Canada develop a staged cognitive-liberty framework. The first stage is definitional: establish public, legal, and regulatory language for cognitive liberty, cognitive integrity, mental autonomy, and cognitive manipulation. The second stage is governance: expand privacy, consumer, competition, and AI oversight tools to recognize manipulative design, psychological profiling, and influence-impact risks. The third stage is evidentiary: develop protocols for documenting influence, including content logs, algorithmic audits, psychological assessment, expert review, and secure access to platform data. The fourth stage is institutional: assign clear regulatory responsibilities and create pathways for remedies, audits, injunctive relief, and public-interest research.

Key finding: Canada has constitutional language, privacy law, consumer protections, and emerging AI governance tools, but no integrated legal framework that clearly protects mental autonomy from non-consensual technological manipulation.

Content Map

Definition and Scope

For this report, cognitive liberty means the autonomy of a person to think, believe, perceive, interpret, and decide without non-consensual manipulation of their mental processes. Cognitive integrity refers to the protective dimension of that autonomy: the safeguards that prevent unreasonable interference with attention, perception, judgment, belief formation, memory, and

decision-making. Mental autonomy is used as a plain-language equivalent where the report addresses policy and public education rather than formal legal doctrine.

This definition is intentionally grounded in observable systems. It does not rely on speculative or conspiratorial claims. The report is concerned with documented influence mechanisms, including algorithmic recommendation systems, targeted advertising, dark-pattern user interfaces, psychographic profiling, behavioural nudges, generative AI persuasion, misinformation campaigns, and emerging neurotechnologies. These systems may operate across private platforms, public communication environments, consumer interfaces, and institutional decision systems.

The report includes four areas of analysis: Canadian constitutional and civil-law gaps; the digital influence environment; comparative and emerging governance approaches; and practical recommendations for legal, regulatory, and evidentiary development. It excludes clinical diagnosis, individualized legal advice, cybersecurity incident response, and claims that cannot be traced to observable mechanisms or documented practices.

The public-interest focus is the protection of human perception and judgment. Cognitive Security Canada's human-systems lens treats the individual not as a weak link but as a person embedded in systems of information, design, incentives, institutions, and social pressure. The question is therefore not only whether a user made a choice, but whether the surrounding system preserved the conditions required for meaningful choice.

Research Questions and Method

This report is guided by four research questions:

What legal protections currently exist in Canada for freedom of thought, psychological integrity, consent, privacy, and decision autonomy?

Where do current legal doctrines fail to address private-sector, algorithmic, or technologically mediated influence?

What international and comparative developments suggest possible pathways for Canada?

What practical reforms would help make cognitive liberty a usable public-interest principle rather than only an abstract concept?

The method is analytic synthesis. The report draws on public legal sources, Canadian Charter doctrine, privacy and consumer protection policy, international human-rights discourse, and selected technology-governance literature. It uses the existing RR001 source base, including Canadian legal commentary on technology-facilitated mind hacking, Justice Canada's Charterpedia discussion of section 7, CIGI work on freedom of thought, Government of Canada consumer guidance on dark patterns, United Kingdom online-safety materials, Australian Human Rights Commission material on neurotechnology, and Canadian AI governance materials [1]-[16].

The report treats uncertainty conservatively. It avoids unsupported causal claims about individual mental states and distinguishes between influence, manipulation, coercion, and

harm. It also recognizes that persuasion can be legitimate, protected, and socially valuable. The concern is not ordinary persuasion; it is covert, exploitative, asymmetrical, or non-consensual interference with the conditions of autonomous judgment.

Conceptual Framework

The conceptual framework rests on three linked ideas: autonomy, environment, and evidence. Autonomy is the person's capacity to form beliefs and decisions. Environment is the surrounding architecture that shapes what the person sees, feels, considers, and rejects. Evidence is the practical requirement that any legal or policy response must be traceable, reviewable, and capable of being tested.

Traditional law often presumes that a person's decision is autonomous unless a recognized vitiating factor is proven: incapacity, misrepresentation, duress, undue influence, unconscionability, or serious state-caused psychological harm. Digital influence challenges that assumption because it can operate through personalization, repetition, emotional salience, information scarcity, social validation, and interface design rather than direct threats or obvious deception.

A human-systems model therefore asks five questions:

Perception: What information was made visible, hidden, amplified, or suppressed?

Interpretation: What emotional, social, or cognitive cues shaped the user's meaning-making?

Choice architecture: Were alternatives available, understandable, and realistically selectable?

Consent: Did the person understand the influence system and have a meaningful way to refuse it?

Accountability: Can the system's design, data use, and outputs be audited after the fact?

This framework connects cognitive liberty to existing legal values without requiring every concern to become a constitutional claim. It can inform privacy design, platform duties, procurement standards, public education, competition enforcement, AI impact assessment, and civil litigation. The goal is not to criminalize influence; it is to identify when influence becomes unreasonable interference with autonomous judgment.

Table 1 summarizes the main constitutional doctrines that anchor the discussion while also showing why they are incomplete for digital influence.

Provision	Protected interest	Doctrine / examples	Gap for digital influence
Charter s.2(b)	Freedom of thought, belief, opinion and expression [1]	Primarily developed through expression cases and limits on government action.	Does not directly bind private platforms and has not been developed as an inner-freedom doctrine for

Provision	Protected interest	Doctrine / examples	Gap for digital influence
			algorithmic influence.
Charter s.7	Life, liberty and security of the person, including serious psychological integrity [2]	Cases recognize bodily and psychological integrity in direct state-action contexts.	High threshold; requires state action and serious harm. Diffuse private influence is unlikely to qualify.
Charter s.1	Reasonable limits	Justification analysis applies after a Charter infringement is established.	Not reached unless a recognizable Charter violation exists.
Related rights	Conscience, equality, privacy-adjacent values	Potentially relevant by analogy.	No settled doctrine for algorithmic belief interference or mental autonomy.

Evidence and Case Studies

The evidence base for cognitive-liberty concerns is strongest where the influence mechanism is observable. The most relevant case studies are not single dramatic events but recurring design and governance patterns.

Algorithmic recommender systems

Recommendation engines on social media, video platforms, search environments, and newsfeeds profile users by behaviour and serve content that is optimized for engagement, retention, relevance, or revenue. The user may experience this as neutral discovery, while the system may be selecting content according to hidden objectives. The autonomy concern is that repeated exposure can reshape attention, emotion, and perceived reality while the person lacks a clear view of the selection process [4][11][12].

Targeted advertising and psychographic profiling

Targeted advertising uses behavioural, demographic, inferred, and contextual data to deliver messages to specific users or groups. In ordinary commerce, targeting can be convenient. In high-impact domains such as elections, health, finance, youth environments, or crisis situations, targeting can exploit vulnerabilities and informational asymmetries. The concern is not that messages are persuasive; it is that the persuasion may be optimized around inferred psychological traits without meaningful consent.

Dark patterns and user-interface manipulation

Dark patterns are design choices that steer users toward actions they may not otherwise choose, such as making cancellation difficult, hiding privacy controls, using false urgency, or nudging users toward more disclosure. Government of Canada consumer materials recognize dark patterns as a public concern, and such practices may overlap with misleading or unfair conduct in some contexts [5][6]. For cognitive liberty, dark patterns matter because they show how autonomy can be impaired through design without overt coercion.

Misinformation, deepfakes, and synthetic persuasion

Coordinated misinformation and synthetic media can alter the information environment that citizens use to form beliefs. Generative AI may increase the volume, personalization, and plausibility of persuasive content. Cognitive liberty does not require the state to decide all truth claims, but it does require serious attention to systems that make it harder for individuals to identify source, intent, authenticity, and manipulation.

Emerging neurotechnology

Neurotechnology and brain-computer interfaces raise a more direct version of the same problem: the boundary between data about the person and data from the person's mental processes. The Australian Human Rights Commission has warned that these technologies challenge traditional boundaries around thought and rights [8]. Even if widespread consumer neurotechnology is still developing, policy choices made now will shape future standards for mental privacy and cognitive integrity.

Threat / Risk Analysis

This report uses a risk model based on mechanism, exposure, vulnerability, harm, and accountability. A cognitive-liberty risk is higher where the mechanism is opaque, exposure is repeated, the user is vulnerable, the domain is high-impact, the influence is hard to refuse, and the system cannot be audited.

Risk mechanism	Observable indicators	Potential public-interest harm	Governance concern
Opaque personalization	Hidden ranking, limited explanation, no easy reset.	Narrowed perception, echo chambers, reduced informed choice.	Transparency and user control.
Psychological profiling	Inferences about beliefs, emotions, vulnerabilities, or mental states.	Exploitation of vulnerable users or groups.	Sensitive-data classification and consent.
Dark-pattern design	Obscured opt-outs, false urgency, forced continuity, confusing consent.	Invalid or degraded consent; financial or privacy loss.	Consumer protection and competition enforcement.

Risk mechanism	Observable indicators	Potential public-interest harm	Governance concern
Synthetic or coordinated influence	Bots, deepfakes, undisclosed AI content, coordinated campaigns.	Distorted public understanding and trust.	Authentication, disclosure, and platform accountability.
Neural or biometric inference	Collection of brain, biometric, affective, or attention data.	Loss of mental privacy and self-determination.	High-risk data governance and rights protections.

The most difficult issue is proof. A person may feel manipulated, but legal systems require evidence. Demonstrating cognitive interference may require content logs, platform data, psychological expertise, design analysis, and comparative testing. Courts and regulators will need standards that distinguish ordinary persuasion from manipulative interference. Without evidence protocols, cognitive liberty may remain rhetorically powerful but practically unenforceable.

The risk analysis also needs proportionality. Not every nudge is wrongful. Many public-interest interventions, such as public-health notices or accessible design, intentionally influence behaviour. The difference lies in transparency, legitimacy, consent, accountability, and the user's ability to understand and resist the influence. A governance framework should therefore focus on high-risk, hidden, exploitative, and non-consensual practices rather than ordinary communication.

Governance and Law

Canadian constitutional gaps

The Charter provides important language but limited operational protection. Section 2(b) includes freedom of thought, belief, opinion, and expression, but the jurisprudence has largely focused on expressive activity. The Charter also applies primarily to government, leaving most private-platform influence outside direct constitutional review [1]. Section 7 can protect psychological integrity, but only in serious state-action contexts. It does not provide a clear route for challenging diffuse private digital influence [2].

Civil law and private-law gaps

Contract law recognizes consent, capacity, misrepresentation, unconscionability, and undue influence, but these doctrines were not designed for mass-scale digital choice architecture. Tort law recognizes some forms of mental harm, but liability usually requires a recognized duty, proximity, foreseeability, and proof of harm. Privacy law regulates personal information but does not yet squarely regulate the influence effects produced by that information. Competition and consumer protection laws can respond to deception, but many manipulative designs may not be framed as explicit falsehoods [3][5][6].

Comparative developments

Internationally, freedom of thought and mental autonomy are receiving renewed attention. The United Kingdom's online safety framework includes language connecting platform duties with freedom of thought and conscience [7]. Australian human-rights work on neurotechnology has highlighted the need to revisit rights boundaries when technologies can infer or affect cognition [8]. European and international discussions on AI, biometric data, neural data, and neuro-rights provide additional reference points. These developments do not settle Canada's approach, but they show that cognitive liberty is becoming a practical governance topic rather than a purely philosophical concept.

Governance gap summary

Domain	Current protection	Main limitation	Potential reform path
Constitutional law	Freedom of thought/expression and security of the person.	State-action limits and high harm threshold.	Use as a guiding value; develop statutory protections.
Privacy law	Consent, accountability, data safeguards.	Focuses on data handling more than cognitive effects.	Classify mental-state inferences and psychological profiles as sensitive data.
Consumer / competition law	False, misleading, unfair, or deceptive practices.	Subtle manipulation may not be treated as deception.	Define manipulative design and dark patterns as enforceable harms.
AI governance	Emerging high-impact AI oversight.	Cognitive manipulation may be under-specified.	Add influence-impact assessments and user-control requirements.
Civil litigation	Capacity, consent, negligence, mental harm.	Poor fit for diffuse algorithmic systems.	Develop evidentiary presumptions and disclosure tools.

Recommendations for Canada

The following recommendations translate the report's analysis into a staged public-interest reform agenda. They are designed to be practical, incremental, and compatible with existing Canadian legal and regulatory institutions.

Recommendation	Lead actors	Timeline	Complexity

Recommendation	Lead actors	Timeline	Complexity
Define cognitive liberty and cognitive integrity in federal policy language, including mental autonomy, thought privacy, and non-consensual technological manipulation.	Justice Canada; Heritage; ISED; privacy and human-rights bodies; civil society.	Short term	Medium
Treat psychological profiles, inferred mental states, neural data, affective data, and high-impact persuasion profiles as sensitive personal information.	Federal and provincial privacy lawmakers; privacy commissioners.	Short to medium term	Medium
Require high-impact platforms and AI systems to complete influence-impact assessments for recommender systems, targeted advertising, and persuasive interfaces.	ISED; AI regulators; platform regulators; procurement authorities.	Medium term	High
Create user-facing transparency and control rights: personalization labels, reset buttons, opt-outs, explanation tools, and data-access rights for influence systems.	Platforms; privacy regulators; consumer protection agencies.	Short to medium term	Medium
Recognize dark patterns and manipulative design	Competition Bureau; provincial consumer ministries; courts.	Short term	Medium

Recommendation	Lead actors	Timeline	Complexity
as enforceable consumer and competition issues where they impair meaningful choice.			
Develop evidence protocols for cognitive-interference claims, including content preservation, algorithmic audit, psychological assessment, and expert review.	Courts; law societies; regulators; researchers; digital forensics experts.	Medium term	High
Assign a clear regulatory home for cognitive-integrity complaints, whether through expanded mandates for existing commissioners or a specialized digital/cognitive integrity function.	Parliament; federal regulators; provinces and territories.	Medium to long term	High
Integrate cognitive-liberty criteria into public procurement for AI, social media monitoring, digital service design, education technology, and public communication systems.	Treasury Board; federal departments; provincial public-sector buyers.	Medium term	Medium

Implementation Timeline

Implementation should be phased so that terminology, evidence, and institutional capacity develop before heavy enforcement. A realistic Canadian roadmap would move from awareness to definitions, then to standards, then to enforcement.

Phase	Period	Priority actions	Expected output
Phase 1: Awareness and terminology	2026	Publish public-interest briefs; convene law, privacy, AI, psychology, and civil-society stakeholders; refine definitions.	Shared vocabulary and initial policy discussion.
Phase 2: Policy design	2026-2027	Map statutory entry points in privacy, consumer, competition, AI, and human-rights law; draft model clauses.	Discussion paper and legislative options.
Phase 3: Standards and evidence	2027	Develop influence-impact assessment templates, dark-pattern criteria, and evidence protocols.	Regulatory guidance and forensic documentation standards.
Phase 4: Pilot enforcement and procurement	2027-2028	Apply cognitive-integrity criteria to high-risk digital services and public procurement.	Test cases, audits, procurement clauses, and reporting templates.
Phase 5: Institutionalization	2028 onward	Assign formal regulatory mandates and remedy pathways.	Stable governance capacity and public reporting.

The timeline is intentionally staged. Cognitive liberty will not become enforceable through a single phrase or declaration. It requires definitions, institutional responsibilities, evidence standards, and public education. Early work should focus on building credible language and avoiding overstatement; later work can address duties, remedies, audits, and enforcement.

Research Gaps and Next Steps

The main research gap is evidentiary. Legal and policy literature increasingly recognizes the problem of cognitive manipulation, but practical standards for proving interference remain underdeveloped. Future work should focus on influence forensics: the repeatable documentation of content exposure, design architecture, algorithmic selection, inferred profiles, psychological impact, and user decision pathways.

Priority next steps include:

Develop a Canadian cognitive-liberty glossary for public, legal, and regulatory use.

Create a model influence-impact assessment for platforms, AI systems, and public-sector digital tools.

Build an evidence checklist for lawyers, investigators, researchers, and regulators assessing manipulative digital systems.

Compare Canadian privacy and consumer law with UK, EU, Australian, and Latin American neuro-rights developments.

Produce a shorter public briefing for law firms, policymakers, and civic organizations.

Identify ethically appropriate research partnerships for studying user autonomy, dark patterns, recommender systems, and AI-mediated persuasion.

A defensible evidence protocol would preserve user data logs; document content served, timing, and interaction patterns; request platform records where legally possible; conduct design and algorithmic audits; gather relevant witness context; and use qualified psychological or cognitive-science expertise where harm is alleged. The goal is not to reduce human judgment to a technical trace, but to make cognitive-liberty claims reviewable, proportionate, and credible.

Conclusion

Cognitive liberty is an emerging Canadian governance issue because the systems that shape thought, attention, perception, and decision-making are becoming more powerful, more personalized, and less visible. Canadian law contains important values that point toward mental autonomy, including freedom of thought, psychological integrity, informed consent, privacy, and protection from deception. Yet these values remain fragmented across doctrines that were not designed for adaptive, data-driven influence systems.

The public-interest task is to translate cognitive liberty from an abstract right into practical safeguards. That means defining the concept carefully, identifying high-risk mechanisms, strengthening consent and transparency, expanding privacy and consumer protection tools, building evidence protocols, and assigning regulatory responsibility. Canada does not need to treat every influence as a rights violation. It does need a framework for identifying when influence becomes hidden, exploitative, non-consensual, and harmful to autonomous judgment.

For Cognitive Security Canada, the significance is direct. Protecting cognitive liberty supports the broader mandate to promote situational awareness, decision integrity, and human-centered

analysis in complex socio-technical environments. The next research step is to convert this report into two companion products: a concise legal awareness brief for law firms and policymakers, and a practical evidence checklist for documenting cognitive-interference risks in observable digital systems.

References

[1] Laidlaw, Emily. “Technology-Facilitated Mind Hacking: Protection of Inner Freedoms in Canadian Law.” Centre for International Governance Innovation (CIGI), Policy Brief, January 2024. https://www.cigionline.org/documents/2507/FoT_PB_no.5.pdf

[2] Justice Canada. “Charterpedia - Section 7: Life, liberty and security of the person.” Government of Canada. https://www.justice.gc.ca/eng/csj-sjc/rfc-dlc/ccrf-ccdl/check/art7.html

[3] Office of the Privacy Commissioner of Canada. Research and policy materials on privacy, online influence, and personal information protection. https://www.priv.gc.ca/

[4] Alegre, Susie, and Aaron Shull. “Freedom of Thought: Reviving and Protecting a Forgotten Human Right.” CIGI Special Report, September 2024. https://www.cigionline.org/documents/2719/Freedom.of.Thought_SpecialReport.Alegre.Shull.pdf

[5] Government of Canada, Office of Consumer Affairs. “Dark patterns.” Innovation, Science and Economic Development Canada. https://ised-isde.canada.ca/site/office-consumer-affairs/en/dark-patterns

[6] Competition Act, R.S.C., 1985, c. C-34, including provisions related to false or misleading representations and deceptive marketing practices.

[7] United Kingdom Parliament. Online Safety Act materials and parliamentary publications, 2023. https://bills.parliament.uk/publications/50887/documents/3349

[8] Australian Human Rights Commission. “Protecting Cognition: Background Paper on Neurotechnology and Human Rights.” https://humanrights.gov.au/know-your-rights/rights-of-individuals/technology-and-human-rights/protecting-cognition-background-paper-neurotechnology-and-human-rights

[9] R. v. Morgentaler, [1988] 1 S.C.R. 30.

[10] Canadian case law on psychological integrity and mental harm, including section 7 and tort-law contexts summarized in Justice Canada Charterpedia and related legal commentary.

[11] CIGI and related scholarship on algorithmic recommendation, digital persuasion, and freedom of thought in online environments.

[12] Alegre, Susie, and Aaron Shull. Discussion of technology mediation, access to information, and agency in Freedom of Thought: Reviving and Protecting a Forgotten Human Right.

[13] Academic and regulatory literature on echo chambers, recommender systems, and platform-mediated harms, as referenced in the original RR001 source base.

[14] Office of the Privacy Commissioner of Canada. “Hacking the Human Mind: Lessons for Canada’s Democracy.” Completed Contributions Program Projects, 2023-2024. https://www.priv.gc.ca/en/opc-actions-and-decisions/research/funding-for-privacy-research-and-knowledge-translation/completed-contributions-program-projects/2023-2024/p_202324_07/

[15] Government of Canada. “The Artificial Intelligence and Data Act (AIDA) - Companion document.” Innovation, Science and Economic Development Canada. https://ised-isde.canada.ca/site/innovation-better-canada/en/artificial-intelligence-and-data-act-aida-companion-document

[16] Government of Canada. Digital Charter and AI governance policy materials, Innovation, Science and Economic Development Canada.

[17] “Neurotechnology, Cognitive Liberty, and the Law: Building a New Legal Architecture for Mental Autonomy in the Digital Age.” Legal Studies in Digital Age. https://jlsda.com/index.php/lsda/article/view/335

Suggested Citation

Cognitive Security Canada. (2026). Cognitive Liberty in Canada: Protecting Mental Autonomy in the Digital Age. Research Report COGSECCAN RR-001.

Disclaimer

This report is intended for public-interest research and educational purposes. It does not constitute legal, medical, cybersecurity, financial, or professional advice. Findings are based on cited public sources and analytic synthesis at the time of publication.