



Doug Sarro
Assistant Professor
University of Ottawa
Faculty of Law



Martin Petrin
DANCAP Private Equity Chair in Corporate Governance
Western University
Faculty of Law and Department of Management

March 31, 2025

British Columbia Securities Commission
Alberta Securities Commission
Financial and Consumer Affairs Authority of Saskatchewan
Manitoba Securities Commission
Ontario Securities Commission
Autorité des marchés financiers
Financial and Consumer Services Commission, New Brunswick
Superintendent of Securities, Department of Justice and Public Safety, Prince Edward Island
Nova Scotia Securities Commission
Office of the Superintendent of Securities, Service Newfoundland and Labrador
Northwest Territories Office of the Superintendent of Securities
Office of the Yukon Superintendent of Securities
Nunavut Securities Office

Re: CSA Staff Notice and Consultation 11-348 – Applicability of Canadian Securities Laws and the Use of Artificial Intelligence Systems in Capital Markets (the “Consultation”)

We appreciate the opportunity to provide comments on the Consultation. The Consultation, as well as previous publications and research by CSA members,¹ speak to the CSA’s commitment to developing a strong understanding of the opportunities and challenges associated with AI in capital markets. We commend these efforts and your willingness to engage with stakeholders.

As outlined in the Consultation, while AI has the potential to enhance market efficiency and improve investor outcomes, it also introduces significant risks. Canada has a unique opportunity to be at the forefront of these developments by fostering responsible innovation while ensuring robust regulatory oversight. Your efforts to develop thoughtful policy approaches can position Canada as a leader in the intersection of AI and financial markets.

¹ OSC & EY, *Artificial Intelligence in Capital Markets – Exploring Use Cases in Ontario* (2023); AMF, *Issues and Discussion Paper – Best practices for responsible use of AI in the financial sector* (2024); OSC, *Artificial Intelligence and Retail Investing: Scams and Effective Countermeasures* (2024) and *Artificial Intelligence and Retail Investing: Use Cases and Experimental Research* (2024).

Below, we outline several considerations in response to the consultation questions, focusing on selected themes. Generally, our aim is to respond more generally, with different categories of market participants in mind, although we recognize – as outlined in the Consultation – that specific participant categories may require different, or more granular, regulations or guidance.

General Comments

In responding to the risks and opportunities presented by AI, the CSA should be ready to question the assumptions underlying its regulatory framework. For example, dark patterns and digital engagement tactics can influence retail investor behaviour in much the same way that investment recommendations can, yet they are currently regulated in markedly different ways. Also, software and service providers supporting the delivery of regulated activities may come to have as much influence over investor interests and market stability as the registered firms they serve – perhaps even more – suggesting it may become necessary to subject these providers to a level of regulatory oversight similar to that applied to registered firms. When regulatory boundaries and categories no longer serve their intended purposes, they ought to change.

In responding to AI, the CSA will also need to ensure its actions reflect available evidence while at the same time recognizing that it inevitably will need to make decisions based on incomplete evidence. Waiting for a conclusive case for reform before responding to new risks effectively means waiting for these risks to crystalize into concrete harms – possibly in the form of widespread investor losses due to fraud or manipulation, or even a new financial crisis. Regulators should be proactive, acting before crisis occurs. They should also be willing to adapt their approach when it no longer fits the evidence, as we believe is the case with respect to online or automated investment advice. At the same time, rethinking existing approaches supports responsible innovation and unlocks new opportunities for industry, the economy, and businesses.

Use Cases Requiring New or Amended Rules, or Exemptions (Question 1)

Our response to Consultation Question 7 (below) will outline how revised CSA guidance could facilitate the delivery of low-cost, automated investment advice for retail investors with relatively uncomplicated financial circumstances and objectives. Beyond this use case, we expect that AI will facilitate the democratization of a range of financial services that are now largely off-limits to retail investors. For example, AI could support the development of automated proxy advisory services that help small-dollar shareholders exercise their voting rights in a way that reflects their personal preferences and circumstances. This would be a significant step towards the elusive goal of ‘shareholder democracy’. However, if such services arise, they may require more intensive regulatory oversight than non-automated proxy advisory services targeted at institutional investors.

Risks and Governance relating to AI Use in Capital Markets: New/Amended Rules or Guidance, or Adaptation of Existing Frameworks? (Question 2)

Existing regulatory and governance frameworks provide a strong foundation to address risks associated with AI in capital markets. Arguably, AI can to a considerable extent be accommodated within existing rules on disclosure, risk management, governance, etc. However, because AI raises novel questions and exacerbates certain risks, regulators and other market participants should develop or adopt AI-specific enhancements to existing frameworks.

Regulatory considerations

Regulators should consider developing a combination of use of existing instruments, targeted guidance, and AI-related disclosure requirements, which may have to be expanded and adjusted on a continuous basis as technology, market practices, and risk scenarios evolve. This could include the following measures, some of which are also mentioned in the Consultation:

- **Disclosure requirements:** Market participants in general should disclose the use of AI; specific AI-related risks (e.g., business risks, fund or investment risks) as well as broader industry and systemic risks; policies on AI oversight by board and management and their implementation; AI governance policies (including data and privacy considerations); details concerning third-party AI arrangements; and AI risk management practices. Disclosures could also be used to increase the transparency of AI-driven decisions (on which see below). However, disclosure requirements should aim to balance transparency with cybersecurity concerns and competitive integrity.
- **‘AI washing’ and similar issues:** Regulators should provide guidance on acceptable market participant statements concerning AI use and capabilities. The Consultation already provides guidance in this regard for both Investment Funds and Non-IF issuers. Regulators should actively monitor these AI statements and address misleading claims.
- **Third-party AI:** In addition to registrants’ responsibilities concerning outsourced services under Part 11 of 31-103CP,² all market participants should be required to disclose details on due diligence and monitoring of, and arrangements with, third-party AI providers. This may include risk assessments and specifics such as audit rights, service levels, and redress mechanisms. Regulators should consider providing guidance on materiality in relation to AI outsourcing, what types of and under what circumstances third party data may be used, details of due diligence and monitoring requirements of AI providers, and guidance on 4th and 5th party risks in supply chains.
- **Regulatory sandboxes:** Regulators should consider time-limited exemptions and accompanying terms and conditions, such as those delivered via the CSA Financial Innovation Hub, to allow AI-driven solutions, particularly in the FinTech sector, to be

² A broader question concerns the allocation of responsibility for financial market risks, i.e. to what extent it should fall on individuals and/or entities. While we have specific views on this issue, it appears to be beyond the scope of the present Consultation.

tested. Outcomes from this testing should be shared with the public, who also should be given an opportunity to provide feedback on resulting revisions to rules or guidance.

Regulatory measures should align with or consider broader AI regulations, both domestic and international, in order to prevent unnecessary regulatory fragmentation and duplication. For instance, Canada's draft Artificial Intelligence and Data Protection Act (AIDA) contained various obligations for AI systems, including pertaining to public information, risk assessment and mitigation, accountability frameworks, human oversight, and record keeping, while the European Union's AI Act sets forth requirements for fundamental rights impact assessments, conformity assessments, certification, and registration systems,³ which may be relevant for certain market participants whose activities reach the level of sufficient connections to the EU.

Governance practices

Market participants, including advisers, dealers, and reporting issuers, should adapt or introduce governance policies and practices to address AI-specific risks. These policies and practices should be aimed at achieving the following objectives:

- **AI literacy and oversight:** Boards should ensure that, as a collective, they possess a level of understanding of AI that allows them to exercise adequate oversight of AI use and its various impacts. Similarly, below the board level, management needs to be equipped to deal with issues relating to the use, implementation, and oversight of AI. This can be achieved in different ways, including through individual directors and managers with in-depth AI knowledge and skills, dedicated AI governance or ethics committees and/or Chief AI Officers (CAIOs). In addition, or alternatively, it is also possible to delegate AI-related responsibilities roles to existing roles or bodies, such as governance committees, risk committees, and Chief Information Officers (CIOs).
- **AI (risk) management, explainability, transparency:** Depending on their specific circumstances and industry, market participants should consider adopting frameworks such as the ISO 42001 AI Management System Standard, the OSFI's E-23 Guideline on Model Risk Management, the Treasury Board Secretariat's Algorithmic Impact Assessment Tool, NIST's AI Risk Management Framework, the Monetary Authority of Singapore's Principles to Promote Fairness, Ethics, Accountability and Transparency in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector, or the IEEE's standards on autonomous and intelligent systems, among others. Additionally, participants can look to AI-related policies adopted by governments and NGOs, including Ontario's directive on the responsible use of AI, which is similar to principles developed originally in the more broadly applicable (including the private sector) European Union Ethics Guidelines for Trustworthy AI, or the OECD's AI Principles. Among others, AI management should include internal mechanisms for appropriate explainability in AI models, ensuring organizations can interpret and justify AI-driven decisions, including in trading and investment contexts.

³ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), PE/24/2024/REV/1.

While the adoption of AI governance policies and practices within market participants should be encouraged, including via disclosure requirements, we do not believe that securities regulators should prescribe the adoption of specific frameworks at this time. Given the rapid pace at which AI is evolving, practices that sound reasonable at the time of their adoption may quickly prove obsolete. Market participants should have scope to adapt as circumstances change. Furthermore, it is important to keep in mind that businesses need to ensure compliance with current and future AI regulation, such as the draft AIDA in Canada and in certain cases laws of other jurisdictions, such as the EU's AI Act. It is important for securities regulators to consider other (general, industry-specific, and international) AI regulations to avoid setting out contradictory requirements or guidance or imposing unnecessary administrative burdens on market participants.

Human Involvement and Oversight (Question 4)

To ensure AI systems work in ways that facilitate human decision-making rather than frustrate it (e.g., through manipulation or deception), these systems need to work with and under the supervision of humans. The idea of 'human-centric' AI can serve as a useful guiding principle in this regard.⁴ However, the level and nature of human involvement required to achieve this goal can vary.⁵ They should depend on the nature of the activities being undertaken by AI and the extent of uncertainty about its outputs. Close human oversight, such as a 'human-in-the-loop' model that requires human review and approval of AI decisions before they are acted upon, may be appropriate in the case of systems engaged in activities where the costs of error are especially high or that engage individuals' legal rights (e.g., the determination of investor complaints), or that operate in ways that make it especially difficult to predict its outputs.

In appropriate cases, however, it can be sufficient to adopt a 'human-on-the-loop' model, under which an AI system is allowed to run autonomously, subject to the following elements of human intervention:

- **Design:** Beyond generally ensuring there is human involvement in the design cycle for new AI systems, it is important that the personnel involved has competencies relevant to the activities that a system will undertake. For example, firms looking to deploy AI systems that will recommend investments or otherwise influence investor behaviour should have proficient individuals who are aware of relevant regulatory obligations involved in the design of these systems. These individuals should also be involved in selecting and vetting any data used to train these systems.
- **Testing:** AI systems should be tested prior to deployment to ensure that the outputs they produce under different scenarios are in line with expectations. The personnel reviewing these outputs should be in a position to judge whether the outputs are in line

⁴ See High-Level Expert Group on AI, *Ethics Guidelines for Trustworthy AI* (2019).

⁵ See M. Petrin & E. Hickman, *Trustworthy AI and Corporate Governance – The EU's Ethics Guidelines for Trustworthy Artificial Intelligence from a Company Law Perspective*, 22 Eur. Bus. Org. L. Rev. 593 (2021); cf. R. P. Buckley et al., *Regulating Artificial Intelligence in Finance: Putting the Human in the Loop*, 43 Sydney L. Rev. 43 (2021) (arguing for deeper human oversight).

with expectations. For example, a proficient individual would be best placed to assess whether automated investment recommendations produced by an AI system during testing reflect regulatory obligations as well as the investment philosophy and risk tolerance underlying the system's design.

- **Monitoring:** After deployment, some or all of the system's outputs should be reviewed periodically to ensure that they remain aligned with expectations. Any significant errors or other problems discovered during monitoring may require that the system be taken offline until the cause of the error is discovered and resolved. If the relevant AI system is consumer-facing, firms should ensure they have mechanisms for collecting and acting on consumer complaints that might point towards problems with that system. It may also be prudent to build circuit breakers into the system that prevent it from acting on outputs that seem likely to be erroneous or inappropriate.

We note that other standard-setters have developed more accommodating expectations for automation, including the US Securities and Exchange Commission and the European Securities and Markets Authority with respect to automated investment advice,⁶ FINRA with respect to automated trading systems,⁷ and the UK Financial Conduct Authority with respect to both algorithmic trading and automated investment advice.⁸

Transparency and Explainability (Question 6)

AI models, particularly deep learning systems, introduce challenges for securities law due to their complexity and opacity. As the Consultation mentions, some AI systems may function as opaque 'black boxes' with decision-making processes that are difficult to interpret. Although new methods for tracking the inner workings of AI models are now emerging, this poses issues for regulatory compliance, especially in areas requiring documentation and traceability of decisions, such as trading activities, portfolio allocations, and detection of illegal activities. These difficulties are amplified when third-party AI is involved, further reducing transparency.

Despite these challenges, requiring full explainability in every case is neither necessary nor practical. There is typically a trade-off between explainability and performance, and different AI applications require different transparency levels. The Consultation therefore correctly suggests that a balance is needed. Consequently, instead of insisting on full interpretability, regulators may where appropriate focus on a combination of outcomes, procedural safeguards/standards, and independent verification to ensure compliance and accountability.

⁶ US Securities and Exchange Commission, IM Guidance Update No. 2017-02: *Robo-Advisers* (2017), at pp. 6–8; European Securities and Markets Authority, *Guidelines on Certain Aspects of the MiFID II Suitability Requirements* (2023).

⁷ US FINRA, Regulatory Notice 15-09: *Guidance on Effective Supervision and Control Practices for Firms Engaging in Algorithmic Trading Strategies* (2015).

⁸ UK Financial Conduct Authority, *Algorithmic Trading Compliance in Wholesale Markets* (2018), at pp. 4–6; UK Financial Conduct Authority, FG17/8: *Streamlined Advice and Related Consolidated Guidance* (2017), at pp. 7–12.

- **Outcome-based oversight:** This involves regulators and users prioritizing the assessment of AI outputs rather than the systems' internal workings. Thus, AI-driven decisions would be evaluated for fairness, reliability, and adherence to securities laws. AI outputs must align with regulatory obligations and avoid introducing systemic risks.
- **Procedural safeguards and standards:** Market participants should be required to implement robust oversight mechanisms for AI models, ensuring that they meet minimum standards for fairness, accuracy, and robustness. This may include establishing explainability benchmarks tailored to different AI use cases. For example, a model used for trade execution may require less transparency than one flagging suspicious transactions for regulatory reporting. Regulators may also require disclosures that explain how AI-driven decisions are made, monitored, and adjusted.
- **Independent verification:** To mitigate risks associated with opaque AI models, regulators may require independent audits and stress-testing requirements. AI applications deemed high-risk or high-impact could be subject to external review to assess their reliability, bias, and compliance with regulatory frameworks. Similarly to bank stress tests, market participants could be required to test AI systems under various conditions to evaluate resilience against bias, errors, and unusual market conditions. Additionally, regulators could require certification or adherence to established industry standards for AI fairness, transparency, and security to ensure regulatory compliance.

While full explainability and a perfect understanding of every system is not always feasible, regulators may utilize the above-mentioned approaches to balance innovation and compliance.

AI-Powered Investment Services (Question 7)

In our opinion it should be permissible for algorithmic or AI systems to carry out registerable activity provided that it is disclosed and proficient individuals are involved in system design, testing, and monitoring, and a registered firm has assumed responsibility for a system's output.

It is not efficient to require human review and approval of investment recommendations directed at clients with uncomplicated financial objectives and circumstances who are looking to invest relatively small amounts, especially if the system generating these recommendations itself is relatively uncomplicated and has been appropriately tested. A firm carrying out discretionary investment management takes on fiduciary obligations to its clients, creating powerful incentives on the part of that firm to ensure any algorithmic or AI systems it deploys to help undertake these functions are appropriately designed or tested. Indeed, automated and AI systems may not only lead to lower costs, but also better outcomes for investors, including by offering more tailored solutions; reduced human biases, errors, and conflicts of interest; and access to certain investment strategies that might otherwise be unavailable to retail investors.⁹

⁹ For an overview of practical and legal-regulatory considerations on robo-advice, see I. H.-Y. Chiu, *Financial Advisory Intermediaries and AI* and E.C. Chaffee, *Regulating Robo-Advisers*, both in the Cambridge Handbook of Private Law and Artificial Intelligence (E. Lim & P. Morgan, eds., 2024).

We are unaware of any provision under securities legislation or National Instrument 31-103 that would prohibit an appropriately registered firm from providing automated advice so long as it has taken reasonable steps to ensure the advice provided is suitable and the firm otherwise complies with its regulatory obligations. Accordingly, it seems to us that CSA staff could resolve this issue by publishing revised guidance outlining the circumstances under which they would view automated systems as being capable of providing suitable advice. This shift in approach would bring the CSA into alignment with their peers in the United States, United Kingdom, Australia, and other jurisdictions, which – as in part mentioned above – have all concluded that automated systems, developed by proficient individuals and for which regulated firms accept responsibility, are capable and should be allowed to provide automated services.¹⁰

In addition to revisiting their approach to automated investment advice, regulators should also clarify their expectations with respect to digital engagement tactics employed by online discount brokers, which can significantly influence investors’ decision-making but are not currently regulated as recommendations.¹¹ The principle behind these expectations should be that the design features employed by these firms should work to serve the interests of clients rather than those of the firm. Regulators should make clear, for example, that the use of ‘dark patterns’ – design features that manipulate clients into making decisions that do not serve their preferences and interests – is not in line with registered firms’ obligations to retail clients.¹²

Investment Product Offerings (Question 8)

Other responses to this Consultation may shed light on how widely AI systems are currently being used in the context of investment analysis and strategies. If AI systems are already widely adopted to analyze potential products for retail clients, it may be appropriate for the CSA to reconsider its expectations with respect to product shelf offerings. It could then assess whether firms should be allowed, or maybe even required, to incorporate a broader range of alternatives when making investment recommendations, leveraging AI’s capacities to better serve client interests. However, as long as AI adoption remains relatively limited in this context, a more measured approach may be warranted. Instead of immediate regulatory revisions, it may then be sufficient to clarify that firms choosing to employ AI investment recommendations must design these systems with their regulatory obligations in mind. In particular, a firm’s AI system should avoid conflicts of interest and not favour proprietary investment products over others.

¹⁰ IOSCO, *Update to the Report on the IOSCO Automated Advice Tools Survey* (2016), at pp. 5–6; IOSCO, *The Use of Artificial Intelligence and Machine Learning by Market Intermediaries and Asset Managers* (2021), at pp. 17–21.

¹¹ OSC, *Digital Engagement Practices: Dark Patterns in Retail Investing* (2024), at p. 5.

¹² *Ibid*, at p. 3.

Third-Party AI Products and Services (Question 9)

The use of third-party AI products and services is widely recognized as a potential source of significant risks (see also below, Question 10),¹³ which makes it appropriate for the CSA to focus on it and, in the Consultation, emphasize registrants' responsibilities in this regard. As mentioned above (Question 2), there should be specific and relatively extensive disclosure requirements concerning the use of third-party AI. Market participants that use third-party AI can take several steps to manage or mitigate risks, and disclosures should be designed to encourage such steps, including due diligence and arrangements that facilitate explainability.

Another option, if detailed disclosures concerning third-party AI are not feasible or deemed insufficient by themselves, is for regulators to require assessments of AI system reliability, conducted by vendors, users, and/or independent third parties. A variation of this approach is a requirement for certifications attesting to a third-party AI's adherence to specific recognized standards or frameworks. Requirements directed at third-party providers themselves, such as in the form of mandated registration, regulatory approvals, or reporting, are yet another option, although they pose significant monitoring and enforcement challenges. Should vendor accountability or related rules be imposed, their scope may have to be limited to areas involving products or services used in applications that are deemed high-risk. Another challenge stems from the fact that AI models may depend on supply chains of data, algorithms, and computational resources, which adds complexity by introducing fourth- and fifth-party risks. This raises the question on how to design requirements concerning disclosures and allocation of responsibility for (materialized) risks that are tied to certain suppliers further down the chain.

Systemic Risks (Question 10)

The increased use of AI systems in capital markets introduces and exacerbate systemic risks.¹⁴ One of the primary concerns is herding behavior, where multiple market participants make similar decisions based on signals from the same AI models and/or data providers. This phenomenon could amplify market volatility and potentially trigger crashes or financial crises. The concentration of technology providers further heightens these challenges and also creates dependencies that could magnify disruptions if these providers experience failures or cyber incidents. This raises concerns about outsourcing risks, making them a particularly important focal point and priority for securities regulators. To mitigate these risks, market participants should consider diversifying their AI providers and developing in-house AI capabilities where feasible. However, in-house AI systems may not be viable for many organizations and applications, particularly those that require scale and complexity offered by large AI models.

¹³ In addition to work undertaken by CSA members (note 1), see, e.g., OSFI & GRI, *Financial Industry Forum on Artificial Intelligence: A Canadian Perspective on Responsible AI* (2023); Bank of England & UK Financial Conduct Authority, *Artificial Intelligence Public-Private Forum* (2022); ESMA, *Artificial intelligence in EU securities markets* (2023).

¹⁴ See, e.g., Gary Gensler, *Office Hours with Gary Gensler: Systemic Risk in Artificial Intelligence*, <https://www.sec.gov/newsroom/speeches-statements/gensler-transcript-systemic-risk-artificial-intelligence-091924>; H. J. Allen, *Driverless Finance*, 10 Harv. Bus. L. Rev. 157 (2020).

In addition to disclosures, regulators should consider introducing system-wide stress testing to assess correlated AI-driven risks in financial markets, transparency requirements for AI-driven trading models, and possibly frameworks to encourage the use of multiple models/providers.

We appreciate the CSA's leadership in addressing AI and capital market-related regulatory issues. We hope you will find these inputs useful and wish you well with your ongoing work. Please feel free to contact us at either of the e-mail addresses listed below on this or any other issue in the future.


Yours truly,



Doug Sarro
Assistant Professor of Law
University of Ottawa
Faculty of Law – Common Law Section



AND



Martin Petrin
Dancap Private Equity Chair in Corporate Governance
University of Western Ontario
DAN Department of Management & Organizational Studies
School of Law

