



March 31, 2025

VIA EMAIL

To:

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 of 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

The Secretary
Ontario Securities Commission
20 Queen Street West
22nd Floor, Box 55
Toronto, Ontario M5H 3S8
E-mail: comments@osc.gov.on.ca

Me Philippe Lebel
Corporate Secretary and Executive Director, Legal Affairs
Autorité des marchés financiers
Place de la Cité, tour PwC
2640, boulevard Laurier, bureau 400
Québec (Québec) G1V 5C1
Email: consultation-en-cours@lautorite.qc.ca

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”)*

The Canadian Advocacy Council of CFA Societies Canada (the “CAC”)¹ appreciates the opportunity to provide the following general comments and responses to the specific questions set out below.

Artificial Intelligence (“AI”) and its applications in capital markets and investing are rapidly evolving in a number of ways, not all of which are easy to map to current regulatory paradigms. However, we believe that regular periodic examination of

¹ The CAC is an advocacy council for CFA Societies Canada, representing the 12 CFA Institute Member Societies across Canada and over 21,000 Canadian CFA charterholders. The council includes investment professionals across Canada who review regulatory, legislative, and standard setting developments affecting investors, investment professionals, and the capital markets in Canada. Visit www.cfacanada.org to access the advocacy work of the CAC.

As the global association of investment professionals, CFA Institute sets the standard for professional excellence and credentials. The organization is a champion of ethical behaviour in investment markets and serves as the leading source of learning and research for the investment industry. CFA Institute believes in fostering an environment where investors' interests come first, markets function at their best, and economies grow. With more than 200,000 charterholders worldwide across 160 markets, CFA Institute has ten offices and 160 local societies. Find us at www.cfainstitute.org or follow us on LinkedIn and X at @CFAInstitute.



regulation and related guidance considering latest developments is the best possible approach, maintaining a balance between addressing underlying regulatory concerns in rules while allowing for latitude for market participants to develop and adopt new technologies – some applications of which may serve to address these same underlying regulatory concerns. We believe Staff Notices such as this one from the Canadian Securities Administrators (“CSA”) are one such tool in seeking to maintain this balance in Canadian securities regulation, and support the principles-based approach to regulation, relying on existing securities laws/regulation to the greatest extent possible in combination with accompanying guidance. We believe this provides the degree of regulatory flexibility necessary to deal with AI’s changing risks to regulatory purposes, while providing the required space for innovation. We would encourage the CSA to maintain this approach in keeping with the notion that securities laws are intended to be technology neutral.

Below are our responses to the specific questions highlighted.

1. **Are there use cases for AI systems that you believe cannot be accommodated without new or amended rules, or targeted exemptions from current rules? Please be specific as to the changes you consider necessary.**

With respect to registrants, we believe that progressive automation of certain aspects of decision-making in the investment process and advice generally by AI systems may cause (either now or in the near future) to revisit existing rules and guidance as to their regulatory purposes. This could start with identifying needs for targeted exemptions (perhaps as part of a sandbox-style learning exercise) and potentially proceed to new rules and/or guidance where needed as to the appropriate oversight and controls required to effectively regulate the inputs, outputs, configuration and key functions of utilized AI systems, particularly when utilized for automation of decision-making in a parameterized and risk-managed way, consistent with regulatory objectives. While (at least for the time being) most registrants may utilize AI systems solely in support of decision-making with human oversight, we believe there is a need to be forward-thinking in regulation to maintain the principle of technological neutrality, and to keep AI-related innovation by registrants (incumbent or new) inside the regulatory perimeter. As such, we agree with the provided guidance on decision-making support and limited automated decisions in a limited way only insofar as it is a temporary measure toward a broader examination of regulation in this area, or as part of direction or an invitation to registrants exploring novel applications in this area to explore novel exemptions or regulatory sandboxes, primarily to facilitate learning by regulators in this area and consideration vis-à-vis needs for regulatory change and consistency with regulatory purposes. However, it is our view that the use of AI systems to automate KYC and onboarding functions is relatively low risk given that human oversight should occur at specific stages of the broader process, such as determining suitability and subsequent investment decisions or recommendations. In this regard, we suggest that a stronger position be taken by the CSA to focus on needs for regulatory adaptation at the higher risk stages of the investment decision cycle, and perhaps more obviously permit the use of AI systems by registrants to automate lower-risk routine actions such as information collection to satisfy KYC obligations.



- 2. Should there be new or amended rules and/or guidance to address risks associated with the use of AI systems in capital markets, including related to risk management approaches to the AI system lifecycle? Should firms develop new governance frameworks or can existing ones be adapted? Should we consider adopting specific governance measures or standards (e.g. OSFI's E-23 Guideline on Model Risk Management, ISO, NIST)?**

As we noted above, we believe a nimble approach characterized by periodic reviews with a focus on regulatory purposes will best position regulators to identify needs for change in regulation and guidance. We believe regulators should actively focus on discovery by encouraging permissive engagement (such as through sandbox approaches) on novel applications and leading-edge technologies. Given the fast pace of evolution of AI systems and applications, we believe that dynamic responsiveness is the appropriate approach. We also believe it would be prudent for market participants to be mindful of existing requirements for technology and process oversight, and to generate governance frameworks specific to the AI systems, applications, and workflows that they use. We believe this will create an impetus for market participants to engage with the risks of AI in a meaningful manner, which they may not otherwise do, if they passively rely on unspecific governance frameworks or policies and procedures designed to target other obligations. This also creates an opportunity for learning for regulators, through review of registrant oversight practices and frameworks.

Furthermore, we believe that market participants should have flexibility in designing their governance measures and as such, we would not yet support the adoption of a prescriptive model, as we believe appropriate standards and expectations for AI governance models are still developing, and have not yet fully crystallized for domains like investing or capital markets. We acknowledge that it may be useful for market participants to be provided with examples of developing governance frameworks and standards through guidance, which they could consider and consult in creating their own oversight framework and procedures. We would encourage the CSA to provide links to sample frameworks and relevant developing international standards such as through organizations like ISO, FINOS², and NIST, including the types of considerations market participants should be considering in tailoring these approaches to their specific uses of AI systems, and to their specific business models. Guidance should also continue to include any minimum expected oversight or governance elements for registrants to consider.

- 3. Data plays a critical role in the functioning of AI systems and is the basis on which their outputs are created. What considerations should market participants keep in mind when determining what data sources to use for the AI systems they deploy (e.g. privacy, accuracy, completeness)? What measures should market participants take when using AI systems to account for the unique risks tied to data sources used by AI systems (e.g. measures that would enhance privacy, accuracy, security, quality, and completeness of data)?**

² "FINOS Releases First Draft of AI Governance Framework for Financial Institutions..." Online: [Link](#)



We acknowledge concerns in giving AI systems unbounded access to personal information, and believe that systems should be configured to use data elements stripped of personally identifiable information to the greatest extent possible, with a strong focus on disclosure and consent as to incremental data usage in AI systems, and with strong governance as to data not becoming widely available to commercial providers of AI systems for model training or other purposes. While we are sensitive to concerns on data accuracy and completeness, we're also mindful that imperfect datasets are abundant and not without utility, particularly in combination with data science techniques and use of AI systems to account for these imperfections and to correct where possible for their shortcomings, without necessarily compromising on quality or utility of outputs, or related regulatory obligations.

In addition to privacy, accuracy and completeness, market participants using AI systems should also consider bias and fairness in AI systems and their inputs, particularly vis-à-vis employee or client data. That is, market participants should guard against collecting and using datasets in a manner that may result in AI systems' outputs that discriminate, prejudice or disproportionately apply to certain groups. AI systems should also be subject to regular ongoing monitoring, performance tracking and testing in these regards, as there is potential for drift in system outputs to occur.

- 4. What role should humans play in the oversight of AI systems (e.g. "human-in-the-loop") and how should this role be built into a firm's AI governance framework? Are there certain uses of AI systems in capital markets where direct human involvement in the oversight of AI systems is more important than others (e.g. use cases relying on machine learning techniques that may have lesser degrees of explainability)? Depending on the AI system, what necessary skills, knowledge, training, and expertise should be required? Please provide details and examples.**

Please see our comments above. In our view, we see a "human-in-the-loop" as an important part of an AI governance framework to monitor key AI system oversight elements such as inputs, configuration choices, reliability indicators, decision parameters and/or outputs. However, we are also cognizant that the role humans should play and where "in-the-loop" they will function should depend on the business model, the sophistication of the firm and key individuals in deploying and monitoring AI systems, the type of AI technology used, and the risk of the business function on which an AI system is being deployed. For certain AI applications, the purpose of the technology might be to replace humans in the performance of a process or function, or the delivery of a product or service. This should be assessed as to the specific risk level as part of a wider AI systems governance framework at the firm, and risk-informed requirements should be placed on the governance of the AI system, with key factors including monitoring of inputs, outputs, configuration choices and parameters. We don't believe the correct response to all AI system governance concerns is manual/human review of all outputs before enacting decisions, but believe that systems of control should be risk-weighted. We also don't believe that 'human-in-the-loop' is homogenously understood, as oversight frameworks can have many effective elements other than solely review of system outputs. We also believe that AI systems used in monitoring as an element of a governance framework is one potentially effective control which further abstracts the concept of 'human-in-the-loop', but while

potentially increasing the efficacy of governance systems' monitoring and oversight. We don't believe 'human-in-the-loop' monitoring of outputs or other elements is a panacea to poorly conceived or operated AI system governance, and believe that the focus of registrants (and regulators) should be ensuring high quality design and implementation of effective AI system oversight frameworks, regardless of whether it is humans, AI systems, or some combination doing the work.

We believe this presents a challenge in skills and knowledge for both regulators and registrants, as the technical skills required to design and implement effective AI system oversight and monitoring governance are not typically found adjacent to deep domain expertise as to the specific AI system applications in the sector, and accompanying regulatory requirements. As a community of practice develops in this area, we would encourage regulators to be active as a convenor and contributor, in elevating domestic and international best practices and standards/framework developments.

- 5. Is it possible to effectively monitor AI systems on a continuous basis to identify variations in model output using test-driven development, including stress tests, post-trade reviews, spot checks, and corrective action in the same ways as rules-based trading algorithms in order to mitigate against risks such as model drifts and hallucinations? If so, please provide examples. Do you have suggestions for how such processes derived from the oversight of algorithmic trading systems could be adapted to AI systems for trading recommendations and decisions?**

We believe effective monitoring is defined in the context of the complexity of the AI system and should be risk-weighted. The more complex the system, and the greater the risks to clients and broader regulatory goals and purposes, the more requirements should apply, and with it the difficulty of effective monitoring increases. We acknowledge the myriad risks of AI systems, including model drift and hallucination, and believe that some inspiration can be taken from oversight of algorithmic trading systems for low to medium-risk applications, through foundational system oversight paradigms such as stress tests, post-trade reviews, and spot checks against expectations. But for the most complex and risky applications of AI systems, where easy dissection of decisions may not be possible on a reasonable timeline or necessarily objective enough to form definitive system performance conclusions, certain techniques from the AI domain itself may need to be adapted to this purpose, such as the use of orthogonal AI models/systems for monitoring purposes, similar to how powerful orthogonal models are being used in training of foundational AI models and to account for limitations in available training data. All of which is to say, the sophistication of the monitoring system and governance framework should be attuned to the risk level of the AI system application. As noted previously, we would encourage the CSA to take a principled approach to facilitate learning and the development of a community of practice, and to facilitate co-discovery of best practices and learning through sandbox approaches and regular reviews with accompanying guidance (and regulatory change where needed).

- 6. Certain aspects of securities law require detailed documentation and tracing of decision-making. This type of recording may be difficult in the context of using models relying on certain types of AI techniques. What**



level of transparency/explainability should be built into an AI system during the design, planning, and building in order for an AI system's outputs to be understood and explainable by humans? Should there be new or amended rules and/or guidance regarding the use of an AI system that offer less explainability (e.g. safeguards to independently verify the reliability of outputs)?

In our view, end-to-end explainability may be difficult to achieve with complex machine learning or generative AI systems, though this is improving in generative AI systems with the proliferation of reasoning and hybrid-reasoning models. Multiple parameters may be at play and the interactions between those parameters and the datasets can result in creative outputs that emerge in ways that can be difficult to immediately comprehend. We would encourage the CSA to take a risk-weighted approach to the risks of AI systems in this area, and to provide additional guidance on what is sufficient explainability for more complex AI systems, particularly in riskier application areas. For instance, whether it would be sufficient if a market participant can explain the types of inputs, including configuration parameters, exception management, datasets used, vendor due diligence procedures, along with performing parameter-driven oversight on the outputs.

7. FinTech solutions that rely on AI systems proposing to provide KYC and onboarding, advice, and carry out discretionary investment management challenge existing reliance on proficient individuals to carry out registerable activity. Should regulatory accommodations be made to allow for such solutions and, if so, which ones? What restrictions should be imposed to provide the same regulatory outcomes and safeguards as those provided through current proficiency requirements imposed on registered individuals?

Please see our comments above. We believe the CSA should take a risk-weighted approach to AI system applications, and are generally of a favourable view as to allowing systems to take on certain lower-risk discrete functions of an individual registrant, provided they are implemented and monitored as part of a robust oversight framework involving both proficient registered individuals and those with technical expertise in implementing and governing AI systems, such that quality of outputs is equivalent to or better than (either in terms of quality, detail, or reliability) those generated by the individual registrant operating without the assistance of AI systems.

There are a range of levels of risk as to the spectrum of registerable activity and AI system applications, and believe that in the near future for most firms the applications of AI systems might be limited to lower-risk data collection and processing applications (such as for KYC purposes), or to augmenting an existing investment decision process. These types of lower-risk applications should have a reasonable expectation of AI system oversight and governance communicated through guidance, with reference to emerging global standards and frameworks, that is surmountable for most registrant firms.

Where AI systems go beyond deterministic problems, or where the predictability or evaluation of outputs of AI systems becomes more complex, such as in making suitability determinations or investment recommendations, the system oversight and



governance requirements should correspondingly increase, and may demand that registrant firms add new technical expertise to appropriately govern AI systems.

In the most complex conceptions of the implementation of AI systems where significant portions of client interaction or investment decision-making might be fully automated through the use of AI applications, we acknowledge the challenge to existing regulatory paradigms, but believe that this would be an opportunity for co-discovery of viable risk-appropriate oversight mechanisms, and would encourage regulators to engage through mechanisms like targeted exemptions and sandboxes for learning with otherwise technically adept and proficient firms, rather than leaning on prohibitions that may not age well. We believe such systems could meaningfully improve client experience and some of the variability of quality of decisions in certain registrant segments, and could have important implications for broader challenges such as cost and value of advice, so would encourage regulators to constructively engage and build their own aptitude and knowledge along the way.

8. **Given the capacity of AI systems to analyze a vast array of potential investments, should we alter our expectations relating to product shelf offerings and the universe of reasonable alternatives that representatives need to take into account in making recommendations that are suitable for clients and put clients' interests first? How onerous would such an expanded responsibility be in terms of supervision and explainability of the AI systems used?**

The Know-Your-Product obligation under National Instrument 31-103 - *Registration Requirements, Exemptions and Ongoing Registrant Obligations* (“NI 31-103”) requires that registered firms only make securities available to the registered firm’s clients if they have been approved. Likewise, when conducting a suitability determination, registered representatives are currently required to consider a reasonable range of alternatives available through the registered firm. In our view, this system of approval and due diligence on prospective investment products and securities provides an important investor protection role. We believe that given the finite and observable data points and relatively deterministic processes in each case that AI systems could be solutions for firms that might at times otherwise struggle in certain instances to implement human processes that lead to consistent high-quality decisions/recommendations by representatives, and could lead to a net savings in compliance costs for objectively better (at least in terms of consistency) client recommendations. Particularly for newer obligations such as KYP requirements, we believe there is the potential for regulatory burden to be relieved through the development and adopting of utility-like technologies (some of which might involve AI systems) at scale. We would encourage regulators to engage with these use-cases and technology applications, and their potentially positive implications for regulatory burden, and consider where regulatory contributions to open-source or collective burden-reducing technology solutions might be prudent, particularly were it through a neutral third-party such as an organization like FINOS.

We believe that were such a responsibility parameterized as to minimum expectations and that utility solutions were co-developed to meet them that this requirement need not be inordinately onerous on most registrants.



9. **Should market participants be subject to any additional rules relating to the use of third-party products or services that rely on AI systems? Once such a third-party product or service is in use by a market participant, should the third-party provider be subject to requirements, and if so, based on what factors?**

In so far as registrants are concerned, we are of the view that the existing regime that requires service provider oversight, as outlined in Part 11 of the Companion Policy to NI 31-103, would be sufficient to cover the use of AI systems when embedded in third-party products or services, rather than directly sourced and integrated into the registrant's core processes and obligations as an AI system (see our prior comments). However, to conduct ongoing reviews on the quality of outsourced services where reliance on AI systems is one component, additional guidance on the types of considerations (and minimum regulatory expectations) registrants should query when conducting such oversight may be helpful given the general lack of registrant sophistication in this area, the nascency of the technology, and the lack of prior relevant guidance on this specific item.

10. **Does the increased use of AI systems in capital markets exacerbate existing vulnerabilities/systemic risks or create new ones? If so, please outline them. Are market participants adopting specific measures to mitigate against systemic risks? Should there be new or amended rules to account for these systemic risks? If so, please provide details.**

Examples of systemic risks could include the following:

- **AI systems working in a coordinated fashion to bring about a desired outcome, such as creating periods of market volatility in order to maximize profits;**
- **Widespread use of AI systems relying on the same, or limited numbers of, vendors to function (e.g., cloud or data providers), which could lead to financial stability risks resulting from a significant error or a failure with one large vendor;**
- **A herding effect where there is broad adoption of a single AI system or where several AI systems make similar investment or trading decisions, intentionally or unintentionally, due, for example, to similar design and data sources. This could lead to magnified market moves, including detrimental ones if a flawed AI system is widely used or is used by a sizable market participant;**
- **Widespread systemic biases in outputs of AI systems that affect efficient functioning and fairness of capital markets.**

In our view, the increased use of AI systems in capital markets carries certain risks and does have the potential to exacerbate existing vulnerabilities and systemic risks. We are also in agreement with the listed systemic risks and are particularly concerned with the herding effect from homogenously configured/trained or acting AI systems that could result in magnified market events and volatility if governance, oversight, and configuration processes fail to appropriately account for this potential behaviour.



However, we believe that insofar as AI systems have potential risks, they are only the amplification of existing risks of poorly configured and governed technology systems, and that need to be appropriately accounted for when designing, implementing, and overseeing these systems. We hope that a highly competitive market develops for the full value chain of inputs, foundational models, and resulting AI solutions for a range of capital markets and investing applications in Canada. We believe in part that there is a need for minimum expectations to be met via close examination of supporting utility applications, while at the same time encouraging robust competition in the most leading-edge of applications. We also believe that AI systems in part may be the most effective solutions to the problems that they exacerbate, and that regulators must be mindful of the systemic fragility effects of increasing systems complexity. As such, we would encourage the CSA to re-envision how to monitor and regulate these systemic risks and second-order fragility effects, and we would not be opposed to the proposal of additional measures to address these concerns.

Concluding Remarks

We thank you for the opportunity to provide these comments and would be happy to address any questions you may have. Please feel free to contact us at cac@cfacanada.org on this or any other issue in the future.

(Signed) *The Canadian Advocacy Council of
CFA Societies Canada*

**The Canadian Advocacy Council of
CFA Societies Canada**