



**Shalomi Abraham**

Senior Vice President, Head of Legal-Canada

Tel: 416.228.8406

shalomi.abraham@invesco.com

---

**Invesco**

16 York Street, Suite 1200

Toronto, Ontario M5J 0E6

16, rue York, bureau 1200

Toronto (Ontario) M5J 0E6

Telephone 416.590.9855 or 1.800.874.6275

[www.invesco.ca](http://www.invesco.ca)

March 28, 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

**VIA EMAIL TO:**

The Secretary  
Ontario Securities Commission  
20 Queen Street West  
22<sup>nd</sup> Floor, Box 55  
Toronto, Ontario M5H 3S8  
[comments@osc.gov.on.ca](mailto: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  
[consultation-en-cours@lautorite.qc.ca](mailto:consultation-en-cours@lautorite.qc.ca)

Dear Sirs/Mesdames,

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

**Introduction**

We are writing to provide our comments on the CSA Consultation 11-348 – Applicability of Canadian Securities Laws and the Use of Artificial Intelligence Systems in Capital Markets (the “**Consultation**”). Thank you for the opportunity to submit comments.

Invesco Canada Ltd. (“**Invesco Canada**” or “**We**”) is a wholly-owned subsidiary of Invesco Ltd. (“**Invesco**”). Invesco is a leading independent global investment management company, dedicated to delivering an investment experience that helps people get more out of life. As of February 28, 2025, Invesco and its operating subsidiaries had assets under management of nearly USD \$1.9 trillion. Invesco operates in more than 20 countries in North America, Europe and Asia. Invesco Canada operates Invesco’s Canadian business and maintains offices in Toronto, Montreal, Vancouver and Charlottetown.

Our comments below provide analysis on the Consultation’s general and specific guidance, following by a detailed responses to each consultation question.

## **General Comments**

We appreciate the CSA’s efforts to provide clarity relating to the application of existing securities laws to the use of AI systems in Canada. We also appreciate the CSA’s work relating to the addressing of risks presented by the use of AI in the capital markets. Given the significant pace of change and innovation, guidance for registrants and market participants is welcome and helpful, and we expect that as further technological advancements are made in this space, additional guidance and industry consultation will be desirable.

### Flexible Approach

As a general comment, we strongly encourage the CSA to maintain a flexible approach to risk mitigation as it relates to AI. Specific rules that may be drafted today might not be applicable tomorrow, and as such, a principles-based approach continues to make the most sense in the Canadian regulatory environment – especially relating to AI. We also strongly agree with the CSA’s assertion that it is the conduct, rather than the technology itself, that should be regulated under our regulatory framework.

### Definition of AI

The CSA defines an AI system as a “machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.” We believe that this definition captures both current and future uses in the market, and we are generally supportive of the breadth of this definition. However, we note that the definition’s generality potentially risks capturing non-AI technologies, such as deterministic algorithms (for example, an algorithm that rebalances a portfolio when specified allocations drift by certain pre-selected amounts, such as 5%). While such examples may lack inferential or adaptive capacity, we suggest refining the definition of AI to specifically exclude static, rule-based algorithms that lack learning capacity, so as to avoid confusion.

### Explainability concept

We generally agree with the CSA’s view that market participants and registrants must balance the need for explainability with the need for using advanced capabilities of AI systems. We also agree that transparency, while almost always helpful, is not the same as explainability, which is a concept that, by definition, helps users better understand the underlying assumptions and decisions made by the AI model.

In our view, it will likely be the case that the more complex an AI system, the less explainable it will inherently be. We also believe that explainability becomes more relevant as usage becomes more proximate to the client. For example, an AI model that makes suitability suggestions directly to clients will require more explanation, in our view, than an AI model that provides registered portfolio managers with a summary of general economic data which may or may not be used in making investment decisions. This is true regardless of the sophistication of the model. As such, in balancing explainability with usage, we would urge the CSA to not inadvertently stymie innovation by requiring a certain enumerated yet practically indeterminable level of explainability when using AI. However, we would be comfortable and likely supportive of guidance that makes clear that certain activities that are proximate to clients require more explainability relative to activities that are not proximate to clients.

This does not, of course, mean that registrant *accountability* wanes with less explainability – a registrant can and should be held fully accountable for decisions it makes (or it allows an AI model to make on its behalf), whether or not it can explain those decisions.

### Portfolio Management

We agree with the CSA's view that currently, it would be challenging for a registrant using AI systems to demonstrate proper compliance with securities laws for a platform that designs portfolios or executes portfolio management autonomously on a fully discretionary basis. However, we also appreciate the caveat that this may change in the future, given the potential of AI to provide transformative services.

We believe that if/when AI systems do become advanced enough to provide fulsome, repeatable and verifiably high-quality portfolio management services, the CSA should issue guidance with additional regulatory expectations surrounding the use of such models.

### Disclosure Obligations

The Consultation states that if an investment fund's use of AI systems is "marketed as a material investment strategy, this should be disclosed as an investment objective to which Part 5" of NI 81-102 applies.

In light of the guidance relating to limiting the use of AI in providing autonomous portfolio management services, we are unclear as to when the CSA would expect the use of AI to be included in a fund's investment objectives, and we would appreciate additional guidance or clarification in this regard. While we share the concern relating to so-called "AI-washing" and we certainly agree that the CSA is correctly referencing the language in NI 81-101F1 to describe the obligation, we can also foresee circumstances where an IFM wishes to market the use of AI as part of an investment strategy where ultimate decisions are left to human (registered) portfolio managers. We would not expect these activities to be referenced in the objectives as "AI-driven", for example, and so would urge the CSA to clarify its expectations.

### **Responses to Specific Consultation Questions**

**Q1.** 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.

We are not currently aware of any use-cases for AI systems that cannot be accommodated without new or amended rules. We note that this is primarily a factor of both the principles-based approach of our securities regulatory regime and the (correct) philosophy set out in the Consultation that conduct, rather than technology, should be regulated. To the extent that there are indeed existing rules that cannot accommodate specific use-cases, our view is that exemptive relief is likely to be the most appropriate mechanism to deal with such issues at this time.

That being said, we acknowledge that it is certainly possible that use-cases come up in the near future where exemptive relief would not be an efficient way to address regulatory concerns. For example, if a rapid increase in autonomous trading systems or AI-driven client advisory platforms introduces systemic risk into our capital markets, a broader regulatory review would be helpful or perhaps even necessary. We also appreciate that the CSA is considering whether current governance and oversight models are sufficient for the most prevalent use-cases for AI, particularly if AI begins to introduce systemic risk into our capital markets.

To the extent that any new rules are being considered now or are considered in the future, we urge the CSA to continue to use a flexible, principles-based approach.

**Q2.** 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)?

We do not believe that new or amended rules are currently necessary in Canada. Guidance in the form of commentary, which sets out of regulatory expectations (such as in the Consultation) is more helpful to registrants, particularly as the pace of change and innovation in this space accelerates.

Existing frameworks can be used by registrants to satisfy risk management requirements, but our view is that formal adoption of one standard over another (or of a CSA-specific model) is unnecessary and would potentially create confusion and additional regulatory burden. Registrants should maintain adequate policies and procedures and should be audited on the basis of having sufficient policies, but requiring the adoption of a specific governance framework would be suboptimal.

That being said, we strongly urge the CSA to not inadvertently use guidance as a form of informal rulemaking. While we appreciate that guidance sets out the CSA members' expectations, it is not the law, and regulatory reviews or enforcement actions should generally not be based on a strict interpretation of issued guidance.

**Q3.** 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)?

We believe that market participants should take adequate measures to ensure sufficient accuracy of data for use in AI systems. However, given the many use-cases of AI, and the spectrum of available sophistication of AI models, we believe that data integrity should be managed through current or revised risk management frameworks as opposed to specific new regulations set out by the CSA. Firms should, to the extent necessary and appropriate, have adequate policies and procedures to ensure that data-related considerations are addressed as part of a broader AI governance framework.

We also note one area of potential systemic risk which ought to be considered by the CSA relating to data - namely AI data homogeneity, which increases the risk of herding if bad data is used in multiple AI systems (or a single system with bad data is used by multiple registrants). As such, if the CSA believes that policies and procedures are insufficient to guard against this potential systemic risk, following testing and public consultation, we would welcome specific recommendations or guidance relating specifically to data integrity.

**Q4.** 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.

AI sophistication can be measured on a spectrum. In our view, it is reasonable to argue that the extent to which humans must oversee AI systems should depend on the nature of the particular system in question and the sophistication of the model. To the extent that there is proximity to clients, additional scrutiny is desirable as well. For example, having a human-in-the-loop (“**HITL**”) for simple, explainable AI models that review and “clean” old data packs may not be as important as HITL when using AI to draft disclosure documents that are received by investors. As we have previously argued, a principles-based regulatory model is the most flexible and appropriate framework in this regard.

It is also important to distinguish between oversight and accountability. Ultimately, registrants should not be able to avoid accountability for the actions of their AI models if those actions negatively impact clients or skirt current regulatory rules. A HITL may mitigate risk, but whether a human is involved should not relieve firms of their ultimate obligations or their fiduciary duties to clients. HITL must not become a ‘safe harbour’ for AI failures or errors.

Finally, we can foresee situations where generative AI inadvertently (or perhaps advertently) prioritizes firm interests over interests of investors, under certain conditions. We do not know whether these potential conflicts of interest will be adequately considered by the gen AI model, and thus believe that HITL may be a critical factor for ensuring compliance with NI 81-107 in the investment funds context.

**Q5.** 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 that there is a spectrum of available tools to effectively monitor AI systems on a continuous basis. These include:

- **Unit Tests:** Developing a comprehensive suite of unit tests for the AI model to ensure that individual components function correctly. For instance, if the AI model includes a feature for sentiment analysis, create tests to verify that the sentiment analysis component accurately classifies text data.
- **Integration Tests:** Ensuring that different components of the AI system work together as expected. For example, test the integration between the data preprocessing module and the model training module.
- **Scenario Testing:** Creating test scenarios that mimic real-world trading conditions. This can include testing the AI model with historical market data to ensure it performs well under various market conditions.
- **Anomaly Detection:** Implementing automated anomaly detection systems that flag unusual or unexpected model outputs for further investigation. This can help identify potential issues early.
- **Random Sampling:** Conducting random spot checks on the AI model's output to ensure it is consistent and accurate. For example, randomly select a subset of trading recommendations and manually verify their validity.

In our view, the choice of which tool to use (if any) should be left to the registrant, given the multiple use-cases for AI. Guidance may be helpful to note that proximity to clients should be a factor in considering the robustness of a testing regime, but ultimately, the CSA should avoid mandating one type of testing over another.

**Q6.** 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)?

We expect that existing rules are sufficient to capture the required detailed documentation of decision-making.

As stated above, one model for balancing the need for sophistication of the AI system with the level of explainability could depend on proximity to clients. That is, the closer an AI process/system is to a client output (such as client disclosure or an investment process), the more explainable it should be, all else being equal. This principle would help registrants design AI models in a manner that is appropriate to the problem they are seeking to solve with AI, but is flexible enough to ensure sufficient accountability.

**Q7.** 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?

We do not believe that AI systems are sufficiently sophisticated to carry out discretionary investment management today, and we do not believe that registered advisors/portfolio managers should be completely removed from the client relationship or investment process.

However, we can foresee a not-too-distant future where the sophistication does catch up to the desire of these FinTech solutions. We do not have a view, yet, as to whether this should be allowed, given potential issues with data, bias, performance etc. If it is allowed, however, we emphasize our view that any AI system that is highly proximate to clients be highly explainable and transparent.

**Q8.** 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?

While we are not currently of the view that expectations relating to suitability should be drastically changed as a result of AI's capacity to analyze product shelf offerings, we are excited by the potential of AI to increase investor choice. Specifically, we believe that AI could be used as a tool to support dealers in increasing the number of products that they are able to offer (including non-proprietary products).

As an independent asset manager, we expect that a robust analysis of products by an adequate AI system could serve investors well by removing barriers that limit product shelves or that create a bias towards proprietary products. This could be used to further the objectives of the Ontario Capital Markets Modernization Taskforce relating to investor choice, and we urge the CSA to consider this further and provide guidance if appropriate. However, we do not believe that AI tools should materially impact regulatory expectations on human registrants, particularly given the proximity to clients of this type of suitability analysis and the importance of adequately conducting a suitability analysis.

**Q9.** 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?

We believe that the regulatory framework relating to oversight of material service providers is sufficient to ensure adequate oversight of third-party products or services that utilize AI. So long as the conduct, rather than technology, is being governed, we do not believe that new rules need to be put in place.

**Q10.** 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.

We are highly appreciative of this question and the concern raised by the CSA of systemic risk (and we will avoid the bad joke about the ultimate systemic risk of AI – ie the “Terminator” scenario). At this time, discussion of systemic risk is too speculative, given the limited use cases in the capital markets today. However, we can see a day where one or all of the enumerated risks become real. Our view is that data must drive an analysis of systemic risk, and as such, would be supportive of efforts by the CSA to collect meaningful data about the use of AI systems in the future, if/when this becomes appropriate.

We would be pleased to discuss any of the foregoing with members of the CSA at your convenience, should that be helpful or desirable.

Yours truly,

**Invesco Canada Ltd.**

Shalomi Abraham  
Senior Vice President, Head of Legal–Canada

cc: Glenn Brightman, Chief Executive Officer  
Invesco Canada Ltd.