# M RNINGSTAR®

22 West Washington Street Chicago Illinois 60602 Telephone: +1 312 696-6000 Facsimile: +1 312 696-6001

# Submitted Electronically

March 31, 2025

The Secretary Ontario Securities Commission 20 Queen Street West 22nd Floor, Box 55 Toronto, Ontario M5H 3S8

# **RE:** CSA Staff Notice and Consultation 11-348 - Applicability of Canadian Securities Laws and the use of Artificial Intelligence Systems in Capital Markets

Ladies and Gentlemen:

Morningstar welcomes the opportunity to comment on the "CSA Staff Notice and Consultation 11-348 - Applicability of Canadian Securities Laws and the use of Artificial Intelligence Systems in Capital Markets"<sup>1</sup> recently published by the Staff of the Canadian Securities Administrators (CSA). Morningstar is a leading provider of independent investment research and has a long history of advocating for transparency in global markets. Morningstar brings several perspectives to the applicability of Canadian Securities laws to the use of artificial intelligence (AI) in capital markets. Morningstar's mission is to empower investors to reach their financial goals, and our comments reflect this.

This letter contains: 1) a summary of Morningstar, Inc.'s views, 2) the views of Morningstar DBRS on AI usage in the credit rating process, and 3) detailed answers from both Morningstar Inc. and Morningstar DBRS to selected questions posed in the consultation, attached as Appendix A.

### **Executive Summary:**

- Morningstar believes the definition of AI used in the consultation is too broad and encompasses both a) low-risk models/tools/techniques—for example, nongenerative AI models like traditional machine-learning tools and simple neural network models and b) higher-risk models like large language models/transformers/generative-AI capabilities.
- AI is beneficial for capital markets. Both low-risk models (LRMs) and high-risk models (HRMs) have long served investors, and low-risk models are sufficiently covered under current regulation.

<sup>&</sup>lt;sup>1</sup> CSA. 2025. "CSA Staff Notice and Consultation 11-348 – Applicability of Canadian Securities Laws and the Use of Artificial Intelligence Systems in Capital Markets." <u>https://www.osc.ca/sites/default/files/2024-</u>12/csa\_20241205\_11-348\_artificial-intelligence-systems-capital-markets.pdf.

- Overregulating LRMs could stifle tools that have long served consumers, investors, and financial institutions, and there are other ways of mitigating their risks.
- Existing regulation and instruments already sufficiently account for AI usage by Designated Rating Organizations, and National Instrument 25-101 is an effective tool that addresses the concerns around AI governance of credit rating agencies presented in the consultation.
- Morningstar suggests regulators take a risk-based approach to AI regulation that does not subject LRMs to the same degree as HRMs.

# I. Definition of AI Used by CSA Staff Is Too Broad and Captures Models That Should Not Be Considered AI

Morningstar believes that the definition of AI on which the consultation relies is too allencompassing. The definition incorporates both deterministic models and generative AI models, when it should focus exclusively on generative AI models. In this comment letter, we will distinguish between generative AI, which presents some risks that should be managed through some additional guidance, and deterministic models, which we believe are currently well-regulated and do present low risk. Overregulating deterministic models would result in increased compliance burdens for the simplest of technologies and burdensome testing, stifling innovation and harming investors, as many of these models have long served investor interests.

The consultation defines AI as follows:

"An **AI system** is 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. Different AI systems vary in their levels of autonomy and adaptiveness after deployment."<sup>2</sup>

There are a range of AI systems. This consultation definition makes no distinction between LRMs and HRMs. An LRM system that generates outputs based on inputs it receives, whether they be recommendations or decisions, could be deterministic in nature, meaning that for any given input the output is mostly predictable. Some examples include but are not limited to basic algorithms seen in extracting data from files, Excel spreadsheets, research algorithms, autocorrect functions, or even spelling error suggestions. HRMs, on the other hand, produce new and unpredictable outputs and determine what output is best based on a large mass of data and probability rather than logical arguments.

LRMs convert inputs into a predictable output with close replicability. Further, the LRMs can often make their functions explainable and thus inherently pose less risks than HRMs. Since generative AI models are both unpredictable and often so complex they are considered "black box" models – meaning they are less explainable – they inherently pose much higher risks for users.

<sup>&</sup>lt;sup>2</sup> CSA. 2024. "Applicability of Canadian Securities Laws and the Use of Artificial Intelligence Systems in Capital Markets." <u>https://www.osc.ca/sites/default/files/2024-12/csa\_20241205\_11-</u> <u>348\_artificial-intelligence-systems-capital-markets.pdf</u>.

HRMs also may pose other risks such as privacy and security concerns, as many are known to use inputs as training data. If someone inputs personal information into a public generative model, it is possible that information could be presented to or retrieved by a third party.

By blanketing both LRMs and HRM technologies under the same definition, two problems emerge.

- 1. This broad definition will stifle existing LRM technologies that have long served investors.
- 2. The lack of distinction will hinder the future progress of LRM model advancement and increase compliance costs.

To avoid overregulation, a definition of AI should center on generative AI models, such as large language models, where output is not a predictable result of inputs. Such models present fundamentally different risks than deterministic ones.

In the alternative, if the CSA does not agree with changing focus only onto HRMs, it should consider different levels of regulation for the varying degrees of risk that AI models pose. Regulating AI commensurate with the risk will prevent overregulation of low-risk models, while adequately addressing the dangers of high-risk models.

# II. AI Is Beneficial for Capital Markets

LRMs and HRMs both have many benefits, such as improving accessibility of information for all customers and reducing operational and compliance costs. Generative models can increase productivity by as much as 14% on average, measured by the number of issues resolved per hour, with even greater gains for less experienced workers.<sup>3</sup> Generative models can also upskill junior workers, improving their efficiency and quality of work. An MIT study found that access to AI enabled less-experienced workers to produce higher-quality work faster.<sup>4</sup> Further, in a survey of marketing and compliance leaders in financial institutions, 85% of respondents said the implementation of AI would save them money.<sup>5</sup> Compliance functions can do much more with what they are allocated if given access to AI.

LRMs specifically can reduce human errors and detect anomalies, allowing for more-efficient data analysis. Sifting through tons of data is burdensome for people to do without the assistance of technology, but through fine-tuned algorithms, nongenerative models can sort through, categorize, and analyze that data in much less time with far more accuracy. Further, nongenerative models can operate continuously with little oversight and little risk, improving accessibility to important information. By using LRMs, people can work through more data, save time, and, by extension, streamline processes, leading to a decrease in labor costs. These models are effective, simple, and explainable, which is exactly why regulators should subject them to less regulatory scrutiny.

<sup>&</sup>lt;sup>3</sup> NBER. 2023. "Generative AI at Work." <u>https://www.nber.org/papers/w31161</u>.

<sup>&</sup>lt;sup>4</sup> MIT. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." <u>https://economics.mit.edu/sites/default/files/inline-files/Noy\_Zhang\_1.pdf.</u>

<sup>&</sup>lt;sup>5</sup> Saifr. 2024. "AI insights survey: Adopters, skeptics, and why it matters." <u>https://insights.saifr.ai/request-ai-insights-survey-cw.</u>

Morningstar's own ratings system uses algorithms to simultaneously expand and streamline research and access to site resources, utilizing both internal and third-party data. With algorithms, we can streamline our own internal processes and provide useful information to investors more efficiently. Morningstar Investment Management also provides plan sponsors with managed accounts for retirement plans. These are LRMs that help investors make investment allocation decisions. Despite their nongenerative nature, both the ratings and managed accounts algorithms would fall under this consultation's definition of AI, even though they are both low-risk, have served investors for a long time, and are sufficiently regulated under current laws and regulations.

### III. Blunt Regulation of Low-Risk Models Will Negatively Affect Capital Markets

Morningstar believes blunt regulations, or regulations that apply evenly to all models under the given definition of AI without consideration for varying risk, would stifle the benefits of LRMs. LRMs benefit capital markets by improving efficiency, reducing costs of labor, and providing objective information such as quantitative investment ratings to investors. LRMs are inherently low-risk, so restricting them to the same degree as generative models is counterproductive.

If regulation stifles LRMs and limits their accessibility, then it will negatively affect investors. Investors rely on LRMs, such as robo-advisors, to support their financial literacy and conduct data analysis. Such algorithmic tools can increase services for investors and help them plan for their future. Broad definitions and blunt regulations should not be the path forward.

There are other ways of addressing AI risks. Strategies include but are not limited to:

- Enhanced disclosure requirements to inform users on the risks of generative AI.
- More-robust oversight mechanisms such as oversight committees, internal generative AI use policies, and human-in-the-loop models.
- Disclaimers on chatbots that ensure users are aware that the bot is not intended to give advice on particular topics, such as investments, and warn users not to provide personal information.

# IV. Existing Technology-Neutral Frameworks and Regulations Are Sufficient for Governing AI Use in Designated Rating Organizations

National Instrument 25-101 (NI 25-101) Designated Rating Organizations<sup>6</sup> is a principles-based and technology-neutral framework that regulators can apply to AI in addition to the technologies it already covers. NI 25-101 already covers the usage of AI in governance and oversight, explainability, disclosure, and conflicts of interest—all of which the consultation focuses on.

To illustrate some of what NI 25-101 accounts for that is applicable to AI:

- DROs are required to develop internal control mechanisms.
- DROs must monitor the effectiveness of those mechanisms.

<sup>&</sup>lt;sup>6</sup>OSC. 2015. "Unofficial Consolidation: National Instrument 25-101 Designated Rating Organizations." <u>https://www.osc.ca/sites/default/files/2020-09/ni\_20150505\_25-101\_unofficial-consolidation.pdf</u>. (NI 25-101).

- DROs must prevent misleading information.
- DROs must take reasonable steps to ensure data is reliable.
- DROs must assess whether the methodologies and models used in determining credit ratings appropriately account for risk.

NI 25-101 accounts for much more than just these considerations of interest to CSA staff, and Morningstar DBRS details many of them in more detail in Appendix A. This national instrument is just one example of currently available tools with robust applicability to AI governance. Morningstar believes that the existing regulatory landscape has already long since accounted for governance of advanced technologies and the insurance that their development and oversight is sufficiently monitored, and it is Morningstar DBRS' belief that NI 25-101 is an effective tool.

### V. Risk-Based Approach

Morningstar suggests that regulators should take a risk-based approach when regulating AI, as not all models pose the same level of risk. By categorizing the threat of each model, then acting on the threat with appropriate concern by applying strict regulations only to those models that pose the greatest risk for harm, low-risk, nongenerative models can continue to serve investors and strengthen capital markets. Regulations that govern AI models commensurate to the risk will allow for a more principles-based and effective approach to AI regulation. A principles- and riskbased approach will allow AI to develop and not increase compliance costs and burdens for market participants disproportionately to the risk involved.

The magnitude of governance activities (that is, manual, hybrid, automated, and, in some instances, incorporating artificial intelligence tools) should reflect the nature, scale, complexity, and overall risk profile of a supervised entity. For supervised entities, there should be a higher degree of automated controls as well as greater integration between the systems of control functions to optimize monitoring activities and a supervised entity's reporting of management information to executive senior management and the management body.

General expectations for companies using high-risk models should include an internal control framework mature enough to assess and manage the risks of AI and to be integral to the AI lifecycle within a company; the establishment of a supervised entity's AI strategy, ethics, and principles; an appropriate governance and risk-management framework; sufficient disclosures and system documentation; and controls around the design criteria, modeling, training, evaluation, and deployment of AI systems.

# Conclusion

In summary, Morningstar acknowledges that regulatory action should be taken in regard to highrisk generative AI models, but we believe that low-risk models should not be defined as AI and are sufficiently governed under the current regulatory regime. Further, given the higher risks posed by generative AI, regulators could provide guidance on setting up an internal control framework. The definition adopted by the consultation is too broad. The definition would require cumbersome testing on existing deterministic models that have been serving investors for a long time.

LRMs carry many benefits for capital markets such as increasing efficiency, reducing labor costs, and improving accessibility to objective information on investments, among a bounty of others. Furthermore, Morningstar encourages the staff of the CSA to carefully consider the concept of making compliance commensurate with the risk posed by particular investor interactions as a means of achieving its goals without discouraging beneficial services for investors.

We thank the staff of the CSA for the opportunity to comment on the consultation. Should you wish to discuss any of the comments in this letter, please do not hesitate to contact any of us as indicated below:

Elizabeth Collins at elizabeth.collins@morningstar.com Adam Leber at adam.leber@morningstar.com Melissa Gentili at melissa.gentili@morningstar.com

Sincerely,

Elizabeth Collins Global Head of Credit Operations and Standards DBRS, Inc. – US

Adam Leber Global Chief Compliance Officer, Wealth, Retirement, and Research Morningstar Inc.

Melissa Gentili Regional Compliance Officer DBRS Limited - Canada

# **Appendix A:**

Questions:

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.

Answer:

Higher-risk models, such as probabilistic generative AI models, should be subject to stricter regulatory scrutiny than lower-risk nongenerative models. Existing rules already sufficiently cover LRMs, particularly in managing conflicts of interest in their usage. Any new regulation or guidance should focus primarily on high-risk models, as a risk-based

approach is the most effective way to guide AI advancement without stifling the development and use of valuable lower-risk models.

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)?

#### General Answer:

Governance standards, such as NIST's framework, can be effective, but any new regulatory standards should specifically focus on the highest-risk models. Much of the risk associated with AI stems from a model's predictability— the less predictable a model's outputs, the greater the risk, and the greater the risk, the more oversight is required. Conversely, the lower the risk, the less need for oversight, and in such cases, overregulation could impose unnecessary burdens on useful and reliable tools.

#### Credit Ratings:

NI 25-101 is generally principles-based and technology-neutral, allowing for regulatory requirements to be interpreted and adapted to varying activities. Morningstar DBRS believes that the existing principles-based regulatory framework governing designated rating organizations addresses the use of AI systems, including in the context of the overarching themes articulated by the CSA in its consultation paper: governance and oversight, explainability, disclosure, and conflicts of interest. For example, as the CSA has indicated in its guidance for DROs, subsections 2.2, 2.7, and 4.8 of Appendix A to NI 25-101 may apply in the context of a DRO's use of AI.

In addition, given that different AI systems may present differing levels of risk, Morningstar DBRS believes that a risk-based approach to AI governance is appropriate and supports equilibrium between compliance and innovation. A risk-based approach also aligns with the existing regulatory framework governing designated rating organizations in Canada. More specifically, subsections 2.26 and 2.27 of Appendix A to NI 25-101 provide the following:

2.26 The designated rating organization will design reasonable administrative and accounting procedures, internal control mechanisms, procedures for risk assessment, and control and safeguard arrangements for information processing systems...

2.27 The designated rating organization will monitor and evaluate the adequacy and effectiveness of its administrative and accounting procedures, internal control

mechanisms, procedures for risk assessment, and control and safeguard arrangements for information processing systems, established in accordance with securities legislation and the designated rating organization's code of conduct, and take any measures necessary to address any deficiencies.<sup>7</sup>

In line with ss. 2.26 and 2.27, Morningstar DBRS has adopted an internal control framework with three lines of defense for the oversight and management of risk and internal control processes. Furthermore, Morningstar DBRS' internal control framework is monitored and evaluated on a periodic basis, including for the purpose of identifying and, as necessary, managing new or emerging risks. To the extent AI systems are used in Morningstar DBRS' credit ratings process, our existing internal control framework requires the identification, assessment, and appropriate management of any corresponding risks.

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)?

General Answer:

Biases are often introduced during the training phase of a model, as the data on which the model is trained can influence its outputs. It is crucial that when overseeing model training through a human-in-the-loop approach, the data provided to the model is unbiased. An AI oversight committee with members of differing perspectives can play a key role in establishing considerate and inclusive data standards. Additionally, to maintain user privacy, models can be adjusted to reject inputs containing personal information. For example, some chatbots, like Morningstar's Mo chatbot, reject questions containing phone numbers, credit card numbers, addresses, Social Security numbers, or dates of birth. These models also often include persistent disclaimers reminding users not to share personal information and offering guidance on how to use the model safely.

Credit Ratings:

Morningstar DBRS believes that market participants should continue to consider and apply existing privacy, confidentiality, data protection, and data integrity policies, procedures, and practices when determining what data to use as an input for any given AI system.

<sup>&</sup>lt;sup>7</sup> NI 25-101. P. 11

Morningstar DBRS has adopted policies and procedures designed to comply with the following provisions of Appendix A to NI 25-101, which may apply in the context of AI. In addition to ss. 2.26 and 2.27 listed in our response to Question 2, they are:

2.6 The designated rating organization, its ratings employees, and its agents must take all reasonable steps to avoid issuing a credit rating, action, or report that is false or misleading as to the general creditworthiness of a rated entity or rated securities.

2.7 A designated rating organization will adopt all necessary measures so that the information it uses in assigning a rating is of sufficient quality to support a credible rating and is obtained from a source that a reasonable person would consider to be reliable.

2.9 The designated rating organization will assess whether the methodologies and models used for determining credit ratings of a structured finance product are appropriate when the risk characteristics of the assets underlying the structured finance product change significantly. If the quality of the available information is not satisfactory or if the complexity of a new type of structure, instrument, or security reasonably should raise concerns about whether the designated rating organization can provide a credible rating, the designated rating organization will not issue or maintain a credit rating.<sup>8</sup>

4.16 The designated rating organization and its DRO employees will take all reasonable measures to protect the confidential nature of information shared with them by rated entities under the terms of a confidentiality agreement or otherwise under a mutual understanding that the information is shared confidentially. Unless otherwise permitted by the confidentiality agreement or required by applicable laws, regulations, or court orders, the designated rating organization and its DRO employees will not disclose confidential information.

4.18 The designated rating organization and its DRO employees will take all reasonable measures to protect all property and records relating to credit rating activities and belonging to or in possession of the designated rating organization from fraud, theft, or misuse.<sup>9</sup>

For example, Morningstar DBRS has in place policies and procedures that set out the measures Morningstar DBRS has adopted so that the data and information it uses in assigning and maintaining a credit rating are sufficient, in terms of quality and quantity, to support a credit rating.

Morningstar DBRS also has in place policies and procedures that describe the process for model development and the requisite approvals for the use of models in determining ratings and other opinions. The procedures also describe the controls in place that are

<sup>&</sup>lt;sup>8</sup> NI 25-101. P. 8-9.

<sup>&</sup>lt;sup>9</sup> NI 25-101. P. 15.

intended to mitigate model error, including assessing the availability, quality, and relevance of data to design, develop, and validate a model..

With respect to generative AI in particular, Morningstar's and Morningstar DBRS' policies and procedures outline responsibilities when entering data into a generative AI tool. Note that Morningstar DBRS does not currently use generative AI in the credit ratings process.

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.

#### Answer:

Human-in-the-loop is a necessary element of AI oversight and does not require a significant amount of human capital, making this requirement not burdensome. AI oversight committees are also an effective approach to managing AI, particularly when their members are people with differing perspectives, as this helps mitigate biases. Additionally, companywide training on the ethical and safe usage of AI, alongside the establishment of an AI code of ethics, can play a crucial role in curbing the risks associated with the use of generative models.

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?

#### Answer:

Regular testing to monitor the effectiveness of AI models is both possible and encouraged, as it is important to maintain the quality of the model over time. However, testing should be weighted toward higher-risk models, as these require more oversight to ensure they continue to function safely and effectively. 6. Certain aspects of securities law require detailed documentation and tracing of decisionmaking. 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)?

#### Answer:

If an institution does not fully understand the processes of its own model, it cannot reliably trust the model's outputs. Modern LLM and generative AI models present unique risks with respect to explainability due to their nondeterministic nature. Nondeterministic behavior can potentially produce biased results, even in models trained with unbiased data and subject to well-constructed alignment processes. Therefore, transparency and explainability of a model are vital. LRMs are inherently more explainable and, as such, should be subject to less regulatory scrutiny.

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?

#### Answer:

For Morningstar in particular, nongenerative algorithms enable investors to use tools and analytics to research securities, assess their risk, and monitor their portfolios. When managing securities, there are too many variables and calculations for humans to handle effectively without the use of nongenerative algorithms, some of which may fall under the proposed AI definition. However, most of these models are not generative and, therefore, provide predictable and explainable outcomes. Additionally, Morningstar and other financial institutions can implement protections on the back-end to prevent investment advice from being given via a chatbot or other generative models.

9. Should market participants be subject to any additional rules relating to the use of thirdparty 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?

General Answer:

Third-party models are valuable to institutions and market-party AI systems because they reduce the need for in-house technical expertise. However, relying on third-party models can raise privacy and confidentiality concerns. Caution should be exercised when considering requirements for third-party models, especially those that rely on open-source code. Open-source models may sometimes lack the industry-specific or market-idiosyncratic contexts necessary for producing accurate results. There are also challenges in determining who is responsible for vetting open-source code in various scenarios, and responsibility may not always fall on the party utilizing the open-source code. Rather than regulating third-party providers directly, we encourage a principles- and risk-based approach that targets high-risk models.

### Credit Ratings:

Regulatory requirements related to managing third-party risk are addressed in Appendix A to NI 25-101, including in ss. 2.28, which precludes the outsourcing of activities if doing so materially impairs the effectiveness of the DRO's internal controls.<sup>10</sup> Morningstar DBRS believes that the existing regulatory framework encapsulates third-party risk related to AI use, and DROs should continue to consider and apply, using a risk-based approach, existing outsourcing and procurement policies, procedures, and practices in the context of a third-party product or service provider that uses AI.

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.

<sup>&</sup>lt;sup>10</sup> NI 25-101. P. 11.

#### General Answer:

Some high-risk AI models may introduce a range of potential risks, which is why it is important for regulators to establish clear AI principles and adopt a risk-based approach to regulation. Blanket regulations that apply to all models within the scope of the consultation's definition may waste resources on models that pose little risk, potentially stifling those models, particularly more-explainable and nongenerative ones. It is generative models and large language models that should receive the majority of regulatory scrutiny. Additionally, the existing regulatory regime can effectively address the risks posed by high-risk models, as many of these risks align with the concerns that the current regulatory framework is designed to handle.