



NorthstarDAO

March 31, 2025

Via Electronic Submission

To:

Canadian Securities Administrators
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

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

Dear Secretary,

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

We appreciate the opportunity to provide feedback on the CSA's Consultation 11-348 and to provide feedback on how to responsibly support adoption of AI systems in capital markets, to potentially inform whether new rules are required, or that existing rules should be changed.

NorthstarDAO is a federally-incorporated non-profit organization whose mission is to mobilize technology entrepreneurs, investors, and community to build the bridge between emerging technologies (such as artificial intelligence and quantum computing), Web3 (blockchain and cryptocurrencies), and the mainstream through research, education, and advocacy. A significant proportion of our grassroots constituency are the startups pioneering emerging technologies, but are disadvantaged by resource constraints to stay on top of regulation and be engaged in policy development at the individual level. Our organization aims to be a coordination, communication, and engagement vehicle for industry with key stakeholders, including regulators such as the CSA and OSC so

that innovation can happen responsibly, and that its promise can be fully realized here in Canada and abroad.

We wholeheartedly support the Administrator's and constituent regulators commitment to explore the implications of this transformative technology while balancing investor protection and believe a robust regulatory framework is crucial for fostering responsible innovation and building trust. I, a fintech AI founder myself, have drafted the response as a synthesis of direct engagement and consultation with founders from our constituency in the fintech, AI, capital markets, and decentralized finance spaces.

Preamble

We are at a critical juncture whereby a transformational, exponential technology has become accessible by both institutions and retail investors alike – enabling everything from hyper-personalized financial planning to automated order systems with near-autonomous operation. These technologies hold massive promise: lowering barriers to entry, broadening investor access, and enhancing efficiency. As broadly a community of techno-optimists, we recognize these potential benefits pose risks that are equally high. AI can exacerbate systemic risks, embed hidden biases, or be exploited for malicious ends if left unchecked. The CSA rightly recognizes that protecting investors, maintaining market integrity, and reducing systemic risk must go hand in hand with enabling responsible innovation.

While the individual questions posed by the consultation are interesting, we believe it's most important to first outline the potential for AI to empower investors and institutions and how a framework-based approach creates a flexible model for accountability – in establishing this framework, the answers can be more systematically addressed from the top down as an industry consensus and bottom-up by helping all players understand how to operate and adapt technologies within regulations.

Within this letter, we:

1. Outline a NIST-based, risk framework approach to AI alongside regulators rules and principles-based approaches (in addition to exploring a "SOX-like" approach to AI accountability;
2. Highlighting blockchain as a complementary technology to address the challenges and risks associated with AI;
3. Provide specific answers and insights regarding the questions asked in CSA Notice 11-348;
4. Suggest recommendations on innovating responsibly going forward, including ways the CSA can work with other agencies and stakeholder groups to enable this future

1. A NIST AI Framework-based Approach to Risk Management

The [National Institute of Standards and Technology has developed an AI risk management framework](#) built in collaboration with the private and public sectors to better manage risks to individuals, organizations, and society associated with artificial intelligence, while enabling its inclusive responsible adoption. We would suggest that the CSA and constituent bodies consider using the framework and roadmap as a means by which to structure its efforts in responsible AI adoption in capital markets.

The framework is comprised of the following:

- Map: establish the context and scope of AI risk. This includes understanding how an AI system is intended to work, where it might fail, and who or what might be impacted;
- Measure: analyze and track identified risks. This can involve using metrics, audits, and testing methods to measure factors like bias, model drift, or security gaps;
- Manage: prioritize risks and take action to address them. Activities might include changing processes, updating data sets, adding protective tools, or putting response plans in place;
- Govern: oversee the risk management process at an organizational level. This includes setting policies, assigning responsibilities, and continually reviewing and improving the organization's AI governance.

This could conceivably be established at a point in time and updated periodically to adapt to market conditions. This would be administered by a working group struck at the federal and/or provincial level composed of builders in AI models and applications, institutions, retail investors, academics and researchers in AI, for example. This consultation exercise can help form part of the “Map” portion of the exercise, while the working group could holistically map the risks beyond the scope of the questions posed within the consultation document. The working group would then be responsible for *measurement* and *management* of the risks identified in the landscape, which would be collaboratively addressed in the *govern* portion with the regulators and industry.

Critically, this approach of applying a risk-based framework with ongoing industry collaboration enables a flexible and adaptable approach to an ever-evolving technology and risk environment domestically and globally. Reporting on this framework can support public awareness of the risks identified, tying them to regulatory guidance and by empowering institutions to factor these into internal processes.

Combining with a Sarbanes-Oxley-like Framework for Accountability?

This risk-based framework with direct accountabilities, transparency, and monitoring could be considered as a similar framework to the Sarbanes-Oxley Act in the United States, and subsequent Canadian-mirroring provisions like C-SOX (Bill 198) – in this case, applied for the use of AI in Capital Markets.

- Scoping risk: in order to make compliance not overly burdensome in light of the expected ubiquity of AI across all businesses, a “materiality threshold” would help determine whether compliance and auditability are required for those deploying AI.
 - Scoping risky use of AI: identify and inventory all AI/ML systems across the organization, mapping out where they are used, the data flow, and the potential harms (e.g., bias, security breaches, ethical concerns)

- Materiality: classify AI-driven risks by potential market impact (e.g., liquidity disruptions, reputational risks, regulatory penalties). Establish thresholds (e.g., trading volumes, market share, systemic importance) at which AI-related failures could be “material” to investors or the overall market.
- Independent controls testing and audit: this would include mandating comprehensive controls governing each AI lifecycle stage in capital markets, including:
 - Data integrity checks: validate market data feeds, reference data, and order flow information.
 - Model validation and stress testing: require rigorous backtesting, stress testing, and scenario analyses for algorithmic trading models.
 - Bias and fairness assessments: For investor-facing or credit-scoring tools, ensure no unintended discrimination or partiality.
 - Auditability: maintain detailed logs of AI model development, data sourcing, training processes, parameter tuning, and production deployment. This record-keeping is especially important for investigating anomalies or trading irregularities that might require regulatory scrutiny.
 - Ongoing controls: implement real-time or near-real-time surveillance for AI-driven systems. Look for signs of performance drift, unauthorized changes to algorithms, or suspicious trading patterns. Any issues must trigger immediate notifications and remediation measures.
 - Change management: control and monitor who can view or modify AI models, especially those directly linked to high-value trading activities or market surveillance. Track changes to source code, data sets, and training configurations to ensure integrity and guard against insider manipulation.
- Accountability: we have seen other jurisdictions where AI-generated content must have human accountability tied to it (for example, China’s regulations with AI-generated media requires a person to be tied to its creation). Accountability for AI can follow accountability for internal controls:
 - Building on the CEO/CFO certification model in Bill 198, executives would personally certify that AI systems pivotal to financial reporting or trading activities meet rigorous governance standards.
 - Firms could optionally designate a Chief AI Officer with sign-off responsibility, but ultimate legal accountability still rests with top executives.
 - Boards of directors already have a fiduciary duty to oversee major risks. An AI-SOX structure would formally expand that duty to include AI governance, requiring periodic board-level reviews of AI risk assessments, internal controls, and audit findings.
 - Disclosures: if AI systems used in trading, risk management, or client services have material weaknesses (e.g., susceptibility to manipulation, inaccurate risk assessment, or unmitigated bias), firms would be required to disclose these risks in public filings (e.g., annual reports, 10-Ks, 20-Fs, or equivalent). Further, providing clarity on the firm’s AI-related strategies and controls during earnings calls or investor day presentations would also provide greater transparency.

While the NIST’s AI risk management framework lacks the legal mandate that SOX carries, it shares conceptual similarities in how it structures risk management, oversight, and continuous improvement. Over time, especially if regulatory bodies cite or adopt the NIST’s language, there is potential for a more formal “AI compliance” model to emerge that might be considered analogous to SOX. However, it should

be noted that the introduction of such a regime should not be burdensome from a cost of compliance perspective, which would make use of AI more inaccessible and create disparities in unlocking the benefits of the technology for smaller firms, startups, etc.

2. Blockchain as a Complementary Technology to AI in Capital Markets

While the consultation does not explicitly address emergent and adjacent technologies, we believe that blockchain is a complementary technology which helps ameliorate critical concerns about AI: from tracking data provenance, to making AI models more accessible and customizable, to making them more auditable for compliance purposes. We highlight these applications and benefits below.

Tracking Data Provenance

AI systems rely on large, high-quality datasets. Yet, many organizations struggle to trace where their data came from, who collected it, and how it was processed. Gaps in data lineage can lead to compliance issues, unintentional bias, or even the use of fraudulent data in model training. Blockchain enables traceability of data through:

- **Immutable Ledger:** Blockchain creates an immutable record of each dataset's source, transformations, and usage. Each transaction (e.g., a dataset being added or updated) is time-stamped and cryptographically secured, preventing unauthorized alterations.
- **Verifiable Credentials:** Parties providing data can attach verifiable credentials (digital signatures or proofs) that substantiate its authenticity. Other stakeholders (e.g., regulators, auditors) can independently confirm the data's lineage;
- **Distributed Trust:** Because the ledger is maintained by multiple nodes, no single entity can unilaterally modify or delete data records. This decentralized trust mechanism reduces the need to rely on a single third-party "gatekeeper."

Blockchain's ledger can improve judgment of data quality, as stakeholders can trace data origins and detect manipulations more quickly. We anticipate this could enable enhanced regulatory compliance (e.g., for privacy laws), as the chain of custody for personal data is more transparent. Further, this enables clearer accountability, because each participant's activity is logged and visible to authorized parties.

Enabling Greater Data and Model Diversity

AI models often suffer from limited or skewed datasets, potentially leading to biased outcomes. Further, greater commoditization of models can lead to amplification of biases or "one-size-fits-all" approaches to model use; blockchain can enable the creation and licensing of models and datasets for model training and inference. Additionally, many organizations have difficulty accessing diverse data and advanced models due to cost or competition barriers. The following are ways how blockchain can enable greater customization of models:

- **Data-Sharing Marketplaces:** Blockchain-based platforms allow multiple parties (individuals, institutions, or consortia) to share datasets securely. Smart contracts can automatically enforce usage rights, compensation, and access controls (e.g., differential privacy).
- **Tokenized Incentives:** Providers of unique or high-quality data can be rewarded with digital tokens or other on-chain assets. This incentivizes more organizations and individuals to contribute previously siloed data.

- Federated Learning Consortia: Participants can collaboratively train AI models without directly sharing raw data. Blockchain can store model updates in a secure manner, ensuring all updates are traceable and no single party can tamper with the model parameters.

Through this, AI models can be trained on a broader or more curated dataset, leading to better generalizability, fairness, or customization. This can further enable lower barriers to advanced data and model access, as open, blockchain-based ecosystems reduce overhead and gatekeeping (and creating marketplaces for these models). Finally, the creation of these licensable models and datasets enable more competition and innovation, because smaller players can contribute specialized datasets or niche models and still receive value for their contributions.

Auditability of AI Models

AI systems, especially complex deep learning models, can behave like “black boxes” as mentioned by the CSA’s consultation request. Stakeholders often find it difficult to understand how a model reached a particular decision, complicating compliance audits and governance. Blockchain, by nature of being an immutable ledger can enable greater accountability and auditability through:

- Model Version Control: Each version or update of an AI model can be hashed and recorded on a blockchain. Auditors or regulators can confirm exactly which model version was in production at a specific time;
- Transparent Training Logs: Hashes of training datasets, hyperparameters, and model checkpoints can be periodically stored on-chain, enabling a tamper-evident record of the training process;
- Smart Contract-Enabled Audits: Automated or on-demand audits can be triggered via smart contracts. These audits can cross-reference logs, data provenance records, and user consent receipts, ensuring that the model’s lineage and usage adhere to prescribed policies.

With blockchain, it will be easier to demonstrate regulatory compliance by showing verifiable proof of model training, testing, and deployment steps. Further, faster forensic investigations when model performance issues arise (e.g., biased predictions or unexpected anomalies) – in the same way balances are accessible at any time contained within on-chain wallets. Finally, there can be increased trust among stakeholders (clients, investors, regulators) who can access a transparent log of model lineage and usage.

While the maturity of blockchain-based AI solutions varies, early implementations show promise in tackling AI’s “trust gap.” Combined with robust governance frameworks (e.g., risk management, privacy-by-design principles), blockchain-driven innovation could significantly elevate the trustworthiness, inclusivity, and accountability of AI systems.

3. Responses to the Questions within Consultation 11-348 (*Paraphrased Question*)

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?

By our high-level analysis, most AI use cases can be addressed through existing technology-neutral securities laws, supplemented by interpretive guidance and principle-based frameworks (especially if tailored per the NIST framework as suggested earlier). However, two use cases arose that may require

new or amended rules, or targeted exemptions – these namely arise due to the autonomous and rapid decision-making that can be enabled by AI:

- Fully autonomous portfolio management: we have observed the building of fully automated portfolio management solutions by builders in Canada and abroad (where no individual ultimately signs off on decisions). If an AI acts as an advising representative, it cannot be certified as required by a registered individual. A new rule or exemption might define a ‘supervisory AI management’ category (possibly at the controls-level, as suggested in the NIST-SOX framework), requiring enhanced oversight.
- High-frequency, self-executing trading: we have also observed builders in Canada and abroad building tools which enable the development, execution, and adaptation of strategies without direct human intervention. Existing algorithmic trading rules may be insufficient for truly learning systems and its autonomous actions. The CSA could consider new guardrails around adaptation frequency, mandatory drift detection, and “circuit breakers”.

In each case, a combination of updated guidance on risk management (per the framework we mentioned), human accountability (at the Executive level), and potential new categories of controls testing would be more appropriate and adaptable for the circumstances.

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?

As suggested in our section on a NIST-based framework, we suggest that an adapted version of the framework, alongside an execution and accountability framework enabled by a SOX-like protocol will support risk management.

3. Data plays a critical role in AI. What considerations should market participants keep in mind when determining which data sources to use (e.g. privacy, accuracy, completeness)? What measures should market participants take to account for unique risks tied to data sources (e.g. privacy, accuracy, security, quality)?

While market participants should be aware of the data sources used in the models they are relying upon, it is currently highly opaque and being revealed through various lawsuits occurring between publishers and content creators and the builders of these foundational models. We however want to highlight other influences of generative AI which may mislead investors:

- AI-generated deepfakes in advertising: there are already reports of deepfake videos spoofing the likenesses and voices of prominent figures including Tesla CEO Elon Musk, Investor Kevin O’Leary, and Canadian figures such as former Prime Minister Justin Trudeau and Omar Sachedina, CTV News Anchor to sell investment advice groups, investment products. These technologies require mere seconds of reference content and are becoming increasingly high-fidelity, making it difficult for users to distinguish legitimate endorsements vs. deepfake content.
- AI-generated media’s influence on capital markets: on [May 22, 2023, an AI-generated image of an attack on the Pentagon](#) that was reposted by verified accounts across social media triggered a temporary, but significant impact to the markets. Around 10 a.m. Eastern time, as the photo circulated across accounts, the Standard & Poor’s 500 declined about 0.3% to a session low. As reputable news sources refuted the claim and it was revealed that image was a hoax, the index

rebounded. It is conceivable as the accessibility of these image and video generation technologies, the increased number of spoofed accounts, and the difficulty in discerning fact from fiction in an increasingly rapid news cycle that AI-generated content can impact capital market sentiments.

- AI will further allow for the aggregation and processing of greater amounts of data. We have already observed projects that are combining historical financial data at the company and markets level, but also other macro data including news articles, social media sentiments, and other geopolitical inputs which might influence capital markets. This more holistic approach may provide greater insights to decision-making, however, it may further complicate and introduce greater noise to decision-making.

We also believe that the blockchain-based solutions we mentioned in the previous section outline real-world solutions for ensuring data provenance, accuracy, and auditability.

As outlined in our NIST-SOX hybrid framework, we believe that a framework that takes a risk-based approach to scope potential impacts of AI, creates testing and audits of key applications of AI, and ensuring accountability at the leadership and board level will ensure that a system to ensuring risks are mitigated in the deployment of AI.

4. What role should humans play in AI oversight (e.g., 'human-in-the-loop')? Are there certain uses of AI where direct human involvement is especially crucial?

Especially in the nascent stages of technology deployment such as AI, greater human oversight is critical. A human-in-the-loop model not only helps in creating a feedback loop for models (training and inference), but is also a vital check for high-stakes decisions (e.g. discretionary portfolio management, large-scale trading) or AI systems using 'black box' approaches with low explainability.

Applying a risk-based approach, humans should be involved to:

- Review critical outputs before they impact markets or clients (this can apply certain risk or materiality thresholds);
- Override or intervene in suspicious model outputs (using variance analysis - similar to how radiologists use AI to assist in diagnostic imaging, where the AI will flag anomalies for radiologist approval);
- Monitor for bias or model drift (flagging where possible review of data, weights, or other tuning of the model is necessary);
- Provide feedback to improve models (reviewing the quality of the outputs and results of the insights of the model)

AI is currently not in a position to be fully autonomous or trusted to directly engage retail investors. The inclusion of humans-in-the-loop helps with escalation, monitoring, and improvement of the deployment of AI.

5. Monitoring AI systems on a continuous basis: can we adapt processes from algorithmic trading to AI?

Assessment of AI can be done in the model training and in the live deployment of the AI. We outline applications of both below. If we can monitor rules-based algorithms in real time, we can similarly keep tabs on AI models, with different means for algorithmic oversight, including in-development, stress tests, post-trade reviews, spot checks, and corrective actions.

- Model training: we have observed companies that are using historical data to test models being used for predictive pricing and trading strategies. This retroactive analysis vs. historical actuals allows models to be fine-tuned;
- Stress tests: build “mini sandbox environments” replicating real-world conditions. For AI-based order execution systems, for instance, run simulations that subject the model to extreme market volatility. Through this, you can evaluate how the AI’s trade recommendations shift under stress to detect “model drift” (subtle performance changes over time).
- Post-trade reviews & spot checks: borrowing from existing “best execution” reviews, we’d examine the rationale behind each AI-initiated trade post-hoc: “did it align with the system’s stated objectives and risk parameters?” This can ensure ‘alignment’ of the model to its intended purposes. As part of a controls framework, you could randomly sample a set of trades or recommendations to confirm they’re consistent and not the product of model drift.

Firms should, however, recognize the dynamic learning capability of AI. Periodic revalidation or re-certification of AI models may be necessary as part of its risk management and model improvement activities.

6. Certain aspects of securities law require detailed documentation of decision-making. How can we ensure transparency and explainability in less-explainable AI?

As described in our section on the use of blockchain in AI model training, inference, and data sourcing, we believe this will create greater transparency and explainability. This, combined with independent audits (which allow for review of proprietary models and a controls framework to ensure its proper deployment) and a risk-based framework can all in collaboration support these objectives.

7. Regulatory accommodations for FinTech solutions that provide KYC, onboarding, advice, and discretionary management in substitution of proficient individuals?

If AI-based solutions seek to replace tasks that currently require human proficiency (which we have observed in other fintech platforms, whereby automation helps improve the speed, capacity, and quality of processing functions), the CSA could consider:

- Conditional exemptions: Where a firm can demonstrate robust AI oversight (as part of its controls and risk-based assessment framework), it might receive relief from certain proficiency requirements, subject to ongoing audits and record-keeping;
- ‘Hybrid’ models: Mandate that an advising representative periodically validate outputs or sign off on changes to the AI model. This can be part of the controls and accountability framework we have outlined at the executive, controls, and operational staff levels;
- Risk-based escalation: as alluded to earlier, the role of the professionals may evolve to an orchestration role of AIs - as has been observed in the healthcare industry. Where radiologists leverage AI for massive, rapid interpretation of diagnostic images, the AI flags anomalies for review and approval by a radiologist to apply their judgment and to determine best courses of action. A similar analogue may apply to high-value, high-risk actions, thereby requiring qualified individuals to review, approve, and execute;
- Limited product scope: Early-stage AI solutions might be restricted to simpler, lower-risk product shelves until proven stable. This can be incorporated as part of the risk-assessment framework to identify which products and functions should have AI onboarded, whereas those

higher risk areas should be more patient until there is greater reliability and more robust controls frameworks to support its deployment.

8. Should the universe of reasonable product alternatives be expanded for AI-driven recommendations?

Given the exponential scale, capabilities, and evolving sophistication of models, it is conceivable that AI can theoretically scan a vast range of assets beyond what humans can cover, leveraging an enhanced array of historical and current data from multiple sources. If harnessed responsibly, this can enhance suitability by factoring in more alternatives. But it also raises complexity in supervising or explaining final choices. A framework-based approach along with the accountability model we have suggested would accommodate these evolving scopes and uses of the technology, regardless of the asset class or data sources; we further expect this to evolve as new asset classes emerge through decentralized finance (e.g., cryptocurrencies and tokenization of real world assets).

9. Should there be additional rules for third-party AI products or services?

While the market participant remains accountable for their decisions and actions, third-party providers should likewise shoulder accountability based on scale, criticality, and the potential impact on markets. Akin to how investment dealers can't blame software companies for bad market decisions, so too must an AI user remain ultimately responsible. However, AI providers are still responsible for being transparent and accountable for the products they offer and the claims they make about their products. This may include the following:

- Due diligence requirements for market participants: before using a vendor's AI system, participants should evaluate reliability, data security, model risk management, etc., but with AI-specific criteria;
- Vendor Registration or Certification: in high-risk/impact scenarios, consider requiring specialized registration/certification or adherence to recognized AI standards. Factors triggering heavier oversight might include market share (how many institutions are relying on the vendor), criticality in trade execution, or broad usage across many firms.

10. Does increased AI use exacerbate systemic risks?

While as techno-optimists we believe that AI presents greater opportunities than risks, we do acknowledge that this technology may exacerbate systemic risks. Critically, in the design of AI models, there is a concept of "alignment", which refers to both internal and external alignment of the model to human values. The external in this case refers to the embedding of values that uphold the best interests of humanity (or in the case of capital markets, the best fiduciary interests of their investors and stakeholders), and internal alignments referring to how closely and persistently does the model adhere to those external values. Strong alignment helps ensure, at a fundamental level, a model identifies, addresses and factors in systemic risks and broader-scale impact of its decisions and outputs.

We envision that potential systemic risks introduced by AI include:

- Increase disparity between institutions and retail investors: while open source and consumer grade models are becoming increasingly commoditized, institutions may invest more heavily in the development of more sophisticated, customized models (which at this point, can take significant resources, time, and costs - inaccessible to most retail investors) to give themselves

an edge. Further, a key element of differentiation in AI will be in data access for model training and inference – which may introduce more walled gardens of data which would be inaccessible only to paying customers or be made proprietary and denied access for model training – increasing the edge of well-resourced, well-connected hoarders of data;

- (Intentional or Unintentional) Coordination of AI systems creating volatility: If multiple AI strategies learn that short-term volatility can be exploited for profit, we could see “self-fulfilling flash crashes.” What may help mitigate this is collective oversight (industry-wide scenario drills), with circuit breakers specifically designed for AI-driven feedback loops at the institution level;
- Concentration risk of vendors: overreliance on one or two major vendors (be it cloud providers or data aggregators or model-builders) means a glitch could reverberate system-wide. Encouraging (or requiring) diversity in data sources and distributed or backup solutions to reduce single points of failure of models, data sources can help ensure accessible, durable use of AI. Greater democratization of data and models (potentially enabled by blockchain), with greater commoditization of consumer models may help mitigate this;
- Convergent/parallel activity: similar architecture and data among popular AI systems could lead them to make the same trades at the same time. This may cause magnified market moves that punish the unsuspecting investor. We have observed the use of multiple models and agents to cross-check the decisions and recommendations of models as a ‘sense check’ or provide alternative views for the user to consider;
- Exacerbation of systemic biases: given most AI’s training on historical data, there is an inherent risk that it exacerbates and perpetuates existing biases. For example, if an AI systematically undervalues certain types of issuers or clients, capital might be choked off from deserving segments, or capital decisions may be sub-optimal in spite of available data. Ongoing bias testing, independent reviews, and back-testing may help identify these biases.

We believe that while these potential systemic risks are possible, they can be mitigated by frameworks, by new technologies, and by proper regulation which allows AI to be deployed responsibly. The possibilities and prosperity offered by AI to institutions and retail investors outweighs these risks, in our assessment.

4. Overall Recommendations

- **Development of a Framework to guide AI adoption and regulatory development in capital markets:** this may be established through a NIST-inspired, risk-based framework, supported in execution by a working group (mentioned later). Further, the accountability for compliance of such a framework at the individual institution level could be accomplished through a SOX-like approach which requires accountability and audits – without becoming overly burdensome for compliance and risk in adopting AI. We encourage the continued innovation through the sandbox models, proactive engagement with industry, and tracking of risks posed by AI with the support from other agencies and jurisdictions, including:
 - CSA and OSFI to address systemic financial risk;
 - CSA and Privacy Commissioner to ensure data usage in AI respects privacy laws (as greater data sources will be incorporated into AI models to make more informed and insightful decisions);

- CSA and Competition Bureau to watch for harmful monopolies in AI vendor space (this may be less applicable if models become more decentralized, as outlined in our blockchain-enabled vision);
- CSA and CRTC to monitor advertising and increased media content creation which may influence investors and capital markets;
- Other cross-agency sandboxes, task forces, and bodies to monitor AI developments
- **Formation of an “AI Pulse” Working Group:** made up of industry experts as an advisory group to provide quarterly reports on how the technology is evolving at the macro and application layer. The group will monitor use cases, advances in technologies, and identify any individual or systemic risks that emerge. As the pace of change will be rapid, edge cases can quickly scale into high-impact, systemic risks, a group of professionals “in the trenches” - from the AI community (model-builders, application-layer builders), institutions utilizing AI, retail investors, AI ethicists, researchers and beyond would be appropriate. AI can be used to support the work of this body with ongoing and real-time scanning of the market’s trends, sentiments, and applications of AI. This would include monitoring potential new enabling technologies such as **blockchain** to support concerns with data reliability, model transparency, and AI accessibility.
- **Invest in Academic-Private partnerships:** Canada is home to 3 of the world’s leading institutions in artificial intelligence (created as part of the Pan-Canadian AI strategy - namely Amii, Mila, and the Vector Institute). While largely focused on research in the sciences, computation, and industrial applications - greater collaboration between the financial services industry at the research, application, and policy layer should be encouraged and supported.
- **Public Education campaigns:** as regulators such as the OSC have invested in, and provided significant amounts of resources on educating about cryptocurrencies as investments, so too should regulators emphasize education around AI/ML literacy - being able to identify risk vectors, how to utilize AI in investment decisions, and what inherent biases may be present in AI-supported tooling. This education should also be part of the risk management of staff training programs at institutions, as well as with regulators in understanding the risk landscape.

Conclusion

There has not been a technology that proves to be as transformative as AI in modern history - its impact will be felt in capital markets and beyond. Canada's securities regulators face a pivotal moment in its nascent days to be able to harness it responsibly while not stifling the opportunities it offers to investors and institutions in Canada. The same AI tools that promise broader market access and frictionless operations can also introduce new vulnerabilities, from deepfake fraud to model monocultures. Ultimately, we believe that AI can revolutionize capital markets: improving transparency, enabling new product categories, and personalizing services for more Canadians, creating new opportunities for wealth and prosperity.

By uniting existing technology-neutral principles, risk-based approaches (such as NIST), along with targeted interpretive guidance and close collaboration across agencies including OSFI, the Privacy Commissioner, the Competition Bureau, the Bank of Canada, the CRTC, and the prospective AIDA framework, Canada can develop a comprehensive and coherent framework for AI oversight that strikes the right balance between protection and innovation.

NorthstarDAO believes that a flexible, risk-based approach, collaborative engagement with industry, and education of all parties will be key to achieving the CSA's goals in unlocking the potential of AI for capital markets, but doing so with consumer protection in mind. We stand ready to contribute to this important endeavor and look forward to continued dialogue with the CSA and the constituent bodies.

Sincerely,



–

Randall Baran-Chong, CPA, CA
Lead Steward and Architect, NorthstarDAO Foundation
CEO, BizBridge Technologies Corporation
E: randall@northstardao.com
W: www.northstardao.com