Guideline

Title	Guideline E-23 – Model Risk Management (2027)
Category	Sound Business and Financial Practices
Date	September 11, 2025
Sector	Banks Foreign Bank Branches Life Insurance and Fraternal Companies Property and Casualty Companies Trust and Loan Companies
Effective date	May 1, 2027

Table of Contents

A. Overview

- A.1 Purpose
- A.2 Scope
- A.3 Application
- A.4 Key Terms
- A.5 Outcomes and expectations

B. Enterprise-wide model risk management

- B.1 Organizational enablement
- B.2 Model risk management framework
- B.3 Use of models

C. Risk-based approach to model risk management

- C.1 Model identification
- C.2 Model risk rating

- C.3 Risk management intensity

D. Model lifecycle management

- D.1 Policies, procedures, and controls
- D.2 Components of the model lifecycle

Appendix 1: Information tracking for models

A. Overview

The financial services industry is experiencing a rapid rise in digitalization and model applications amplified by the surge in artificial intelligence / machine learning (AI[1] / ML) models. Institutions are increasingly relying on models to support or drive decision-making including in business areas that traditionally did not rely on models. Models now use more diverse data sources and complex techniques that heighten different aspects of model risk. Increased model risks could expose institutions to financial loss from flawed decision making, operational losses, legal implications, or reputation damage. Furthermore, models based on rapidly progressing technologies, like AI, can exacerbate these risks and others. Models characterized by dynamic self-learning and autonomous decision-making may become prevalent.

Institutions should be cognizant of how the use of models in their business can impact their risk profile and should have effective risk management practices to mitigate the risks. Model risk management should be conducted with integrity, at all times, particularly in a world where newer use cases, including those powered by AI, play a greater role in day-to-day operations.

A.1 Purpose

This principles-based guideline sets out our expectations for effective enterprise-wide model risk management (MRM) using a risk-based approach.

A.2 Scope

This guideline applies to all federally regulated financial institutions, including foreign bank branches and foreign insurance company branches, to the extent it is consistent with applicable requirements and legal obligations related to their business in Canada as set out in Guideline E-4 on Foreign Entities Operating in Canada on a Branch Basis.

A.3 Application

This guideline applies on a risk-basis, proportional to the institution's:

- size
- strategy
- risk profile
- nature, scope, and complexity of operations
- interconnectedness, such that disruptions could harm other financial institutions, the financial system, or the broader economy

A.4 Key Terms

Model

An application of theoretical, empirical, judgmental assumptions or statistical techniques, including AI/ML methods, which processes input data to generate results. A model has three distinct components:

1. data input component that may also include relevant assumptions,
2. processing component that identifies relationships between inputs, and
3. result component that presents outputs in a format that is useful and meaningful to business lines and control functions.

Model risk

Model risk involves the risk of adverse financial impact (for example, inadequate capital, financial losses, inadequate liquidity, operational, or reputational consequences) arising from the design, development, deployment, and/or use of a model. This is the inherent risk of using a model and refers to the fundamental characteristics of the model and materiality to the institution. That is, the potential impact on the risk categories described in our Supervisory Framework[2].

Residual model risk

Residual model risk for the purpose of this guideline refers to the risk that remains after institutions have implemented controls, validation processes, monitoring, and other risk-mitigating measures. Thus, residual model risk captures the portion of risk that continues to exist despite institutions' best efforts to identify, measure, and mitigate model risk.

Note: From here onwards, model risk refers to inherent model risk unless otherwise stated.

Model Risk Rating

A categorical model risk tier associated with the inherent level of model risk based on quantitative and qualitative criteria reflecting intrinsic model vulnerabilities and materiality of model impacts from usage.

Model Lifecycle

The components that define the life of a model. It encompasses all steps for designing, operating, and maintaining a model until it is decommissioned. Model lifecycle components are model design (including model rationale, data, and development), model review, model deployment, model monitoring, and model decommission. These components are not necessarily sequential and may vary according to the type of the model, organisational structure, and use case.

Model Risk Management (MRM) Framework

A set of policies and procedures, that reflects the institution's risk appetite for model risk and defines the governance requirements to manage model risk. It includes roles and responsibilities as well as defined processes to identify, assess, manage, monitor, and report on model risk, both at an enterprise level and throughout the model lifecycle.

Model Inventory

An institution’s system of record for storing key information related to models and serving as a basis for reporting. It contains all models whose inherent risk is determined to be non-negligible to the institution.

Model Owner

The unit(s) or individual(s) responsible for coordinating model development, implementation and deployment, ongoing monitoring and maintaining the model's administration, such as its documentation and reporting. The model owner may also be the model developer or user.

Model Developer

The unit(s) or individual(s) responsible for designing, developing, evaluating, and documenting a model's methodology.

Model Reviewer

The unit(s) or individual(s) responsible for reviewing the model’s conceptual soundness, inputs, methodology, assumptions, and performance. Other responsibilities might include providing the model developer and user with guidance on the appropriateness of models for defined purposes and assessing model monitoring results as a part of periodic or ongoing validation.

Model Approver

The unit(s) or individual(s) or committee(s) with the authority to approve the use of a model within the institution, typically after considering the findings reported by the model reviewer and other governance requirements. This

role is often part of a higher-level governance body (for example, a Model Risk Committee or senior management function) that ensures each model aligns with the institution's risk appetite and strategic objectives.

Model User

The unit(s) or individual(s) that rely on the model's outputs to inform business decisions.

Model Stakeholder

The unit(s) or individual(s) that are involved in the model lifecycle, use and governance of the model (for example, all parties defined above, legal team, compliance function).

A.5 Outcomes and expectations

The following are the expected outcomes of effective MRM:

1. Model risk is well understood and managed across the enterprise.
2. Model risk is managed using a risk-based approach.
3. Model governance covers the entire model lifecycle.

B. Enterprise-wide model risk management

Outcome 1: Model risk is well understood and managed across the enterprise.

Institutions recognize model risk is transverse in nature, and senior management holds an enterprise-wide view of the risks. Model stakeholders are aware of a given model's intended use, inherent limitations, and potential negative outcomes to their business. There is adequate governance to manage and control model risk. Institutions are accountable for their use of models.

B.1 Organizational enablement

Principle 1.1: Effective reporting structures and proper resourcing should enable sound model governance.

Consistent with our Corporate Governance Guideline, senior management is responsible for:

- defining and applying the roles, accountabilities, and expectations for effective MRM across the enterprise,
- ensuring appropriate MRM personnel are in place with the requisite skills and experience, particularly for novel technologies, like AI,
- ensuring appropriate communication and reporting of model risk to the board of directors.

MRM should involve a multi-disciplinary team representing a wide range of expertise and functions from across the organization, including legal or ethics professionals as appropriate. This comprehensive approach is particularly critical as institutions adopt advanced technologies like AI/ML, which can significantly increase model complexity and potential impact.

Appropriate resources should be allocated to MRM and those resources should be allocated optimally towards managing the identified risks. Institutions should be able to provide evidence that they are structured and resourced to support a sound governance framework.

B.2 Model risk management framework

Principle 1.2: The MRM framework should align risk-taking activities to strategic objectives and risk appetite.

Institutions should establish an MRM framework that:

- fits into an institution's broader risk and governance framework as outlined in our Corporate Governance Guideline,
- reflects the institution's risk appetite for model risk,
- defines the processes and requirements to identify, assess, manage, monitor, and report on model risk,
- has clear guidelines for the major components of their MRM framework such as: model identification, model inventory, model risk ratings, and requirements for model lifecycle governance,
- defines how an institution provides transparent and consistent reporting of model risk at different levels of the enterprise,
- is subject to periodic review, especially as new technologies emerge, and

- covers models or data sourced from external sources like foreign offices or third-party vendors (pursuant to our Guideline B-10 Third-Party Risk Management Guideline).

B.3 Use of models

Principle 1.3: Models should be appropriate for their business purposes.

Institutions should deploy models only when they meaningfully contribute to decision-making, risk assessment or business purposes and deliver reliable outcomes consistent with their intended use. Effective MRM aids institutions by supporting the safe adoption and use of models to advance business goals.

Developing, deploying and using a model depends on having a clear purpose along with adequate data, systems, and technology. These elements should, therefore, be part of any discussion on model risk.

As organizational needs evolve, models that are no longer fit for purpose should be modified, replaced, or decommissioned.

C. Risk-based approach to model risk management

Outcome 2: Model risk is managed using a risk-based approach.

A risk-based approach is implemented and ensures MRM requirements are proportional to the level of model risk identified by the institution. A risk-based MRM framework is documented and implemented in a way that enables consistency across functional and business units. Institutions identify sources of model risk, and ensure adequate resources are allocated to manage, mitigate, or accept[3] those risks as appropriate. A model inventory keeps an active record of models with non-negligible risk and key information about those models, including model risk ratings. The model risk rating approach reflects all material dimensions of inherent model risk. Model governance and other elements of MRM is commensurate with the identified level of model risk.

C.1 Model identification

Principle 2.1: Institutions should identify and track all models in use or recently decommissioned.

Institutions should have defined processes to periodically identify models used throughout the enterprise, including vendor and third-party models. These processes should include:

- surveying the institution to identify new models and updating the status of existing ones.
- applying triaging, including identifying whether a model has non-negligible inherent model risk and should therefore be subject to model lifecycle governance requirements (Note: we recognize that not all identified models, as per the Key Terms section A.4 above, carry model risk. Hence, models that do not bear model risk, may not require full model lifecycle governance.),
- assigning a provisional rating for new models or updating an existing model's risk rating where it has materially changed in substance or usage.
- storing in the institution's model inventory all models deemed to carry non-negligible inherent model risk.

An institution's model inventory should be:

- a comprehensive inventory of models whose inherent risk is determined to be non-negligible to the institution.
- maintained at the enterprise level, storing key information related to models (see Appendix A) and serving as a basis for management and regulatory reporting.
- accurate, evergreen, and subject to robust controls.
- updated in a timely manner, including model modifications and changes in model use, risk rating, or performance status.
- inclusive of decommissioned models for a period the institution considers reasonable.

C.2 Model risk rating

Principle 2.2: Institutions should establish a model risk rating approach that assesses key dimensions of model risk.

A risk rating approach should be implemented based on inherent model risk, thereby reflecting model vulnerabilities and materiality of model impacts. This enables institutions to ensure that their most critical or complex models receive enhanced scrutiny, while less risky models are subject to fewer requirements.

The risk rating approach should be supported by clear, measurable criteria for each risk dimension and incorporate both quantitative and qualitative factors:

- Quantitative factors include the importance, size and growth of the portfolio that the model covers (as applicable), or potential operational, security or financial impacts.
- Qualitative factors include business use or purpose, model complexity or level of autonomy, reliability of data inputs, customer impacts, or regulatory risk.

Institutions may organize risk factors according to other risk dimensions relevant to the institution's context and practice (for example, "vulnerability and materiality" or "uncertainty and impact").

Each model should be assigned a model risk rating. Institutions should have defined processes to do this. The processes may include having a temporary risk rating that is confirmed in the model review or model approval stage of the model lifecycle. Model risk ratings should be regularly reviewed and updated as appropriate, including when a trigger event occurs (for example, a decrease in performance or a material change in the model's use, data, or infrastructure).

Institutions may include a risk rating category that implies a negligible level of inherent model risk and therefore exempts such models from the full model lifecycle governance requirements. There should be a robust process to approve and track such exemptions.

Externally developed models should be assessed for model risk ratings on a standalone basis.

Even a well-structured MRM framework cannot eliminate model risk: some elements of uncertainty or limitations will persist even after controls and mitigants are applied. We do not expect residual model risk to drive the primary governance and oversight of models. Visibility of both the inherent and residual model risk, however, provides a view of the extent of risk mitigation and is useful for reporting to senior management.

In cases where model risk falls outside the institution's risk appetite, the institution should establish appropriate remediation actions.

C.3 Risk management intensity

Principle 2.3: The scope, scale, and intensity of MRM should be commensurate with the risk introduced by the model.

The institution's MRM framework should establish the scope, scale, and intensity of model governance requirements through the model lifecycle based on the inherent model risk rating.

The inherent model risk rating should drive the:

- frequency, intensity, and scope of model review.
- documentation requirements.
- level of authority required to approve the model, and any exemptions as needed.
- frequency, intensity, and scope of model monitoring.
- interval at which the risk rating is re-assessed.

The model risk rating should also, with information from the model review assessment, and depending on the risk appetite, determine the:

- limits or constraints on model usage (for example, limiting the model's scope, implementing additional safeguards, or even pursuing a different modeling approach)
- intensity of monitoring approach (for example, higher-risk models with limitations may require increased scrutiny in monitoring)

- controls and mitigants used to manage residual model risk (for example, allocating capital reserves to cover potential losses, enhancing model's conservatism, and contingency planning)

Unique MRM requirements may also be applied using other criteria (for example, by model type) but these should be in addition to the primary risk-based requirements. Institutions should further ensure their MRM capabilities are appropriate relative to overall complexity levels. For example, the extensive use of advanced AI/ML techniques should have correspondingly mature governance and oversight.

D. Model lifecycle management

Outcome 3: Model governance covers the entire model lifecycle.

Institutions manage risks arising from each component of the model lifecycle. Depending on the business case and the nature of the model, the overall model lifecycle may vary. Components include design—which encompasses model rationale, data acquisition, and model development—, review, deployment, and monitoring. Organizations have robust and reliable processes, procedures, and controls implemented throughout the model lifecycle.

D.1 Policies, procedures, and controls

Principle 3.1: MRM policies, procedures, and controls should be robust, flexible, and lead to effective requirements applied across the model lifecycle.

An institution's policies, procedures and controls, covering the full model lifecycle, are crucial components of the MRM framework. They should:

- apply to all models commensurate with the risk,
- ensure governance structures and practices remain sound and consistent across the enterprise,
- define explicit responsibilities and accountability for model stakeholders (for example, owners, reviewers, and users),

- identify and involve key stakeholders—such as data science teams, business units, compliance, ethics and legal teams, information technology, and risk management—early in the model lifecycle process,
- ensure appropriate independence and objectivity are maintained,
- be sufficiently flexible to accommodate evolving technologies, different model types (especially crucial given the "black box" and autonomous nature of many AI/ML models), varying levels of model risk, and organizational changes,
- be comprehensively and thoroughly documented.

D.2 Components of the model lifecycle

Model design

Model design encompasses:

- establishing clear organizational rationale for a model.
- adhering to standards that ensure data quality and accuracy.
- following appropriate model development processes.

Model rationale

Model owners should establish a clear rationale for deploying a new model or modifying an existing model. This involves:

- articulating a well-defined purpose—including the model's scope, coverage, and how its outputs are to be used (for model modifications, the rationale should explain the need for such changes).
- identifying the specific business use case and assessing the risk of the model's intended usage.

When establishing a rationale for models using advanced techniques, like AI/ML, model owners should consider several additional factors that stem from the nature of these models. These can include the:

- level of transparency and explainability required.
- need for alternative controls, especially for "black box approaches" or autonomous models.
- potential for the model to lead to biased outcomes, negative social and ethical implications, or privacy risks.

Model data

Principle 3.2: Data used to develop the model should be suitable for the intended use.

Data can vary in terms of sources, formats, and types. It may be structured, semi-structured or unstructured. It can also be synthetically generated or derived from proxy sources. The consequences of flawed data (for example, due to inaccuracies, bias, or missing records) is significant. This is especially crucial for AI/ML models, since they can easily mirror unwarranted data relationships and preserve them in the outputs.

Institutions should adopt robust data governance with standards for collecting, storing, and accessing data used in models. This should be integrated with enterprise-level data governance, strategy, and management. Data leveraged for model development should be:

- accurate and fit-for-use (that is, free from material errors with biases understood and managed).
- relevant and representative (that is, reflecting the intended target population of the model).
- compliant (that is, adhering to statutory, regulatory, and internal requirements for data ethics and customer privacy).
- traceable (that is, having documented lineage and provenance).
- timely (that is, updated at an appropriate frequency).

One important purpose of the above data properties is to enhance the explainability of the model and associated outputs. Consideration needs be given to the potential for unwanted bias within the data which can translate into unfair model outputs with associated reputational risks.

Institutions should also:

- regularly perform thorough data quality checks (for example, outlier detection, missing value analysis, and consistency evaluations).
- implement controls to ensure quality and document provenance and use of synthetic data and proxy data.
- have controls to ensure appropriate data cleansing operations (for example, handling of inaccurate and missing values).

Model development

Principle 3.3: Institutions should have model development processes that set clear standards for performance and documentation.

Model development involves selecting conceptually sound methodologies, data, and algorithms. In the case of AI/ML models, it also requires proper training and parameter optimization. Institutions should establish clear, consistent, and repeatable practices for developing models. These include:

- standards for model documentation, including elements such as:
    - the set up and running of the model,
    - limitations and restrictions on use,
    - detailed descriptions of processes for creating, accessing, and maintaining the data used to develop the model,
    - detailed descriptions of model assumptions and methodology,
- the role of expert judgment, including which experts are involved and how their input affects the model output,
- analyses and performance tests performed by the developers.
- guidelines for selecting conceptually sound methodologies, data (including the application of variable transformations), and algorithms.
- explainability requirements, which may vary based on the model's purpose, level of autonomy, regulatory requirements, or the potential impact on customers and stakeholders.
- performance and other criteria for model selection.
- standards on how the model outputs are used and reported.
- performance and other criteria for model monitoring.

Model review

Principle 3.4: Institutions should have a process to independently assess conceptual soundness and performance of models.

The model review process should be independent from model development. It should validate that models are properly specified, working as intended, and fit-for-purpose. Institutions can use the work of internal reviewers or objective third parties. The extent and frequency of model reviews should be commensurate with the model risk rating. Events that should prompt a model review include the following:

- development of new models,
- model modifications (including changes to algorithms, parameters, or supporting operational components),
- breaches of model performance (for example, when monitoring highlights material errors, drifts, or threshold violations),
- significant data changes (for example, new data sources, altered data definitions, or shifts in data quality),
- scheduled risk-based periodic reviews.

Model reviews may include:

- confirming or challenging the model risk rating.
- reviewing the model purpose, scope, conceptual soundness, limitations, mitigants, and reasonableness of model outcomes.
- reviewing data quality and appropriateness.
- reviewing novel methodologies, algorithms, tools, and procedures for example, for AI/ML models.
- evaluating the level of explainability for the model workings as per the intended use of the model, including AI/ML models (This includes confirming that model outputs are appropriately explainable and comply with performance expectations).
- reviewing third-party models and platforms or sub-components (including data and libraries) used for model development.
- documenting review assessments and actions taken in response.

- reporting the outcome of the review, including an assessment detailing any findings and recommendations, along with an overall recommendation on approval to the model approver.

Model approval

Model approval processes should occur throughout the model lifecycle. The model approval decision typically involves two components:

- Assessing whether the model is suitable to be implemented into production (or continue to be used) based on its intended use.
- Affirming the assigned model risk rating and residual risk assessment.

The model may be approved despite identified weaknesses or limitations provided that compensating mitigants are in place, or the model stakeholder group provides justification for using a model with known limitations or weaknesses.

Model deployment

Principle 3.5: Models should be deployed in an environment with quality and change control processes.

Model deployment requires careful coordination to ensure the model is properly configured, tested, and moved into production[4]. Integrating models into operational settings should preserve their integrity and reliability. This is particularly important for AI/ML models that may depend on multiple components, diverse and dynamic data sources, and third-party elements.

Effective model deployment should involve the following, as applicable:

- collaboration among model developers, model owners, model reviewers, model users, and technology or operations teams,
- consistency between data used to develop the model and the production dataset,
- tests demonstrating that the model operates as expected in the production environment,

- clearly documented procedures that outline deployment steps, stakeholder responsibilities, approval hierarchies, change control, monitoring frameworks and exception handling including overlays,
- performance of risk assessments for related risks prior to deployment such as cybersecurity risk, infrastructure vulnerabilities, and other potential operational risks (see OSFI Guidelines B-13 and E-21),
- review of explainability requirements and communication of explanatory outputs to the appropriate stakeholders.

Model monitoring

Principle 3.6: Institutions should have defined standards for model monitoring, and model decommission.

Model monitoring ensures models remain fit-for-purpose and aims to detect performance issues or breaches. This may be particularly challenging for complex AI/ML models, where scalability, performance, and monitoring for model drift are critical. Model monitoring should document and include:

- monitoring standards covering frequency, scope, and evaluation criteria as applied to different risk rating levels and model types,
- evaluating criteria that include both quantitative measures (for example, performance metrics) and qualitative assessments (for example, verifying the model is within its original scope),
- tracking and evaluating operational factors such as changes in: model performance, model usage, input data, external dependencies (for example, version updates), or the characteristics of what is being modelled (for example, added product features),
- defining thresholds for breaches and criteria for material model modifications, including changes in operational factors,
- determining contingency plans for model unavailability, deterioration in model performance, or outright failure along with the escalation procedures for addressing these,
- implementing processes for handling AI/ML's unique challenges, such as autonomous decision making, autonomous re-parametrization, and the elevated potential for model drift,
- ensuring issues are shared promptly with relevant stakeholders, following appropriate escalation procedures.

Model decommission

Model decommission is the formal retirement of a model from active use. Reasons to decommission may include:

- performance issues (for example, a severe performance failure, repeated monitoring breaches, excessive overrides),
- business, regulatory, or strategic changes,
- obsolete data, methodology, or cost–benefit considerations.

Decommissioning should follow a disciplined process and include:

- Alerting all relevant stakeholders of the planned decommission.
- Retaining the retired model and documentation for a set period as a benchmark or fallback.
- Determining what additional actions are needed for decommissions of any third-party models.
- Monitoring downstream effects to ensure no residual impacts.

Appendix 1: Information tracking for models

At a minimum, institutions should maintain the following for each identified model:

- model ID
- model name and description of key features and use
- model risk rating
- model owner
- model developer
- model origin (for example, internally developed or vendor).

In addition, for models deemed to have non-negligible risk and hence stored in the model inventory, institutions should also maintain:

- model version
- date of model's deployment into production

- model reviewer
- model approver
- model dependencies
- data sources and description
- approved uses of the model
- model limitation(s) (including exceptions and additional requirements)
- date of model's most recent model review
- monitoring status (with exceptions as applicable)
- next review date.

Footnotes

1 While there are currently no universally agreed-upon definitions of AI, for the purposes of this guideline, consider the OECD definition of AI (where an AI system can be based on one or multiple models): "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."

2 Supervisory Framework - Office of the Superintendent of Financial Institutions

3 When permitted under the institution's risk appetite.

4 The term production environment should be understood as a generic term for the way in which a model or its outputs are accessible once it has been approved.