Draft guideline

Title	Draft Guideline E-23 – Model Risk Management
Category	Sound Business and Financial Practices
Date	September 30, 2017
Sector	Banks Foreign Bank Branches Life Insurance and Fraternal Companies Property and Casualty Companies Trust and Loan Companies
No	E-23

Table of Contents

> ## Consultation status: Closed
>
> Consultation closed March 22, 2024. We'll keep this draft on the site until the final guideline is released.

The financial services industry is experiencing a rapid rise in digitalization and model application. Organizations are increasingly relying on models to support or drive decision-making. Models now use diverse data and complex techniques, with risks further amplified by the surge in artificial intelligence / machine learning (AI/ML) analytics due to greater data access, digitization, reduced data and storage costs, and enhanced computing power. Increased model risks could expose organizations to financial loss from flawed decision making, operational loses, and/or reputation damage. As such, it's important for organizations to mitigate their model risks by adopting robust model risk management practices and oversight, including adequate controls. OSFI expects organizations to have good understanding of the model lifecycle including associated processes and controls, as well as a sound and prudent model risk management (MRM) framework.

1. Purpose and scope

This Guideline, which is principles-based, sets out OSFI's expectations related to enterprise-wide model risk management (MRM) built on strong model lifecycle principles. It applies to all organizations[1] and to all models,

whether they require formal regulatory approval or not.

OSFI expects model risk to be managed on a risk-based and enterprise-wide basis. The term “enterprise” is used throughout this Guideline, when used in a FRPP context, it refers to the contractual arrangements from which the pension plan is derived, and not to the institution represented by the plan sponsor.

Decisions on how to best manage enterprise model risk are the responsibility of the organization.

2. Definitions

Model

The application of theoretical, empirical, judgmental assumptions and/or statistical techniques, including AI/ML methods, which processes input data to generate results.[2] A model has three distinct components:

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

Model Lifecycle

The subset of stages defining the life of a model. It encompasses all steps for operating, governing and maintaining a model, until it is decommissioned.

Model risk

The risk of adverse financial (e.g., inadequate capital, financial losses, inadequate liquidity, underfunding of defined benefit pension plans[3]), operational, and/or reputational consequences arising from flaws or limitations in the design, development, implementation, and/or use of a model. Model risk can originate from, among other things, inappropriate specification, incorrect parameter estimates, flawed hypotheses and/or assumptions, mathematical computation errors, inaccurate/inappropriate/incomplete data, improper or unintended usage, and inadequate monitoring and/or controls.

Model Risk Management (MRM) Framework

An organization's model risk expectations, governance, including key controls and oversight. MRM frameworks, which are supported by robust model lifecycle management, facilitate sound decision-making within an organization.

Model owner

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

Model reviewer

The independent unit(s)/individual(s) responsible for model validation and reporting its findings and recommendations to the model approver. 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. It is acceptable for the roles of model reviewer and approver to be combined if there is no potential conflict of interest and independence is maintained from the model owner, developer, and user.

Model approver

The unit(s)/individual(s) and/or, depending on size and complexity of the model, committee(s), responsible for assessing the model reviewer's findings and recommendations and approving the use and/or limitation of use of any new model or changes to pre-existing models.

Model user

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

Model stakeholder

The individual(s)/unit(s) that are impacted by the output of the model (e.g., all parties defined above, legal team, compliance function).

3. Outcomes

This Guideline presents the following expected outcomes for organizations to achieve:

1. Models are adequately managed at each stage of their lifecycle.
2. Model risks are managed proportionally to the organizations' model risk profile, complexity and size.
3. Models are well understood within the organisations and associated risks are managed through a well-defined enterprise-wide Model Risk Management framework.

4. Model lifecycle

Principle 1: Organizations develop, approve, and implement processes and controls that define expectations for each of the lifecycle components.

Components of the model lifecycle are illustrated in the following diagram.

Model lifecycle

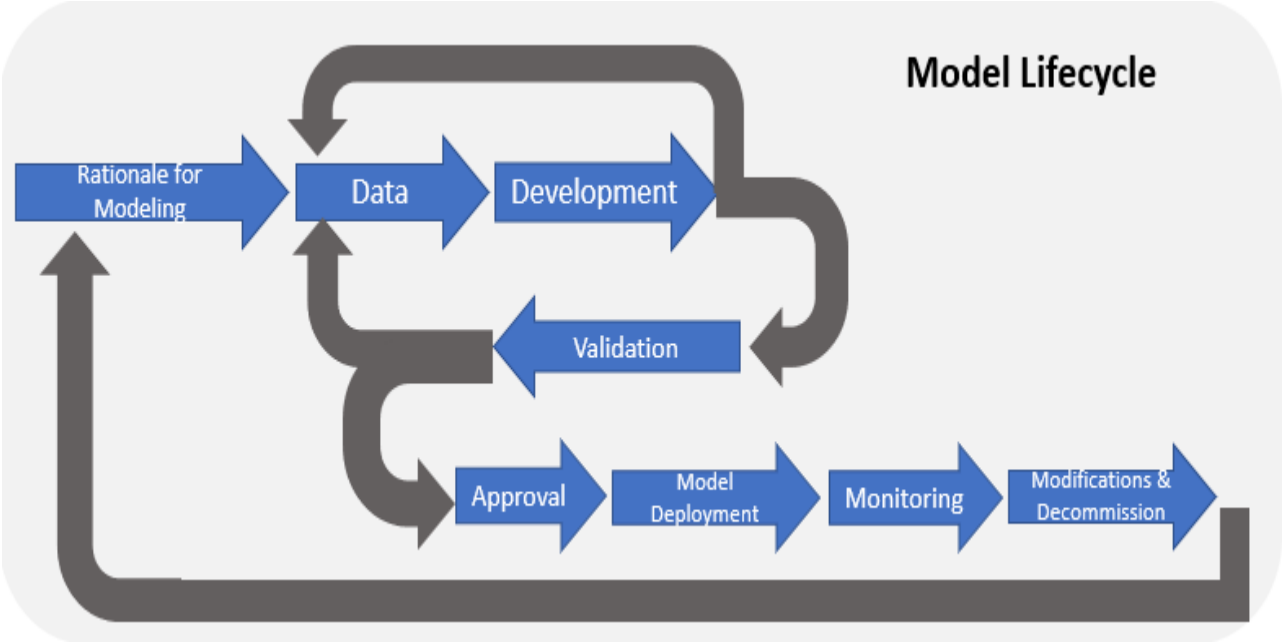

Figure 1 Model lifecycle - Text version

The flow chart illustrates the different stages of the model lifecycle. It consists of a series of arrows connecting various stages, starting from "Rationale for Modeling" and progressing through "Data," "Development," "Validation," "Approval," "Model Deployment", "Monitoring", "Modifications and Decommission". The chart includes the double arrows from "Validation" to "approval" and back to "data" depicting cases where approval is not received.

The chart also includes the double arrows from "development" to "validation" and back to "data" for cases where development is ongoing. Furthermore, the chart includes arrows from "Modifications and Decommission" back to "Rationale for Modeling" to depict that this is a continuous process.

The paths depicted in the graph are illustrative; alternative paths are possible.

Principle 2: Processes and controls consider the size, complexity of the organization and the model's usage.

Each phase of the model lifecycle should be proportionate to an organization's complexity,[4] size and its model usage. For individual models, the application of the model lifecycle requirements should depend on the assigned model risk rating[5]. For newly developed models, where a model risk rating is not yet assigned, organizations may apply requirements of the model lifecycle based on a provisional rating[6] or identify unique requirements (e.g., by model type).

Organizations should ensure documentation supporting the model lifecycle is current, maintained for each phase of the model's lifecycle, and commensurate with the risks of the model. Models using more complicated modelling techniques, with significant assumptions, extensive use of expert judgement, large capital and/or customer impacts should generally be supported by more comprehensive documentation.

4.1 Rationale for modeling

Prior to development, the model owner should identify a rationale for modelling, which clearly articulates the model's purpose, how its outputs will be used, and scope of its coverage. For previously approved models subject to modifications, the rationale should address why changes are needed.

In establishing a rationale for modeling, the model owner should ensure that all stakeholders of the proposed model are identified. Relevant stakeholders may include, but are not limited to, the model users, enterprise data office, compliance, and legal departments. The decision to proceed to the next step of the model lifecycle should include input from relevant stakeholders, particularly noting qualitative criteria that inform the model development process, as appropriate.

4.2 Data

To manage model risk arising from data, organizations should ensure that data leveraged for model development has the following properties:

- Accurate and fit-for-use (e.g., free from material errors, bias is understood and managed)
- Relevant and representative (e.g., reflects the intended target population of the model)
- Adequately complete for its intended purpose
- Traceable (e.g., lineage and provenance are understood and documented)
- Timely (e.g., updated with a frequency aligned with its intended use).

Also, organizations should be aware of the potentially amplified presence of data quality issues when unstructured data is used. Where synthetic data elements are used along with empirical data for model development, organizations should have controls to delineate such data and should conduct appropriate assessments to ensure that the data meets the above properties.

4.3 Development

An effective model development process should comprise the following activities:

- identify suitable data (that meets the above properties), critical assumptions, and the quantification of key parameters (i.e., calibration)
- data cleansing
- develop a conceptually sound methodology to arrive at desired outputs
- identify appropriate performance measures to assess the model's quality, including defining acceptable performance boundaries
- develop the code
- assess model robustness (stable output from small changes in input values) in relation to changes in model risk drivers
- understand and communicate model outcomes and how they will be achieved
- develop a format for model outputs so that model users can effectively make sound business decisions and model owners can monitor ongoing model performance.

Models that rely on expert judgement to inform key components of model development should reflect the above activities, where appropriate. To support expert judgement as part of the development process, organizations should consider the experts who were consulted, the process used for gathering and analyzing expert perspectives, how a consensus was developed, and the applicability of data indicators.

As the prevalence of dynamic model calibration[7] increases, organizations should identify when a model recalibration event has occurred.

4.4 Validation

Organizations should independently review model development outputs, whether developed in house or by external parties. Organizations can use the work of internal objective reviewers (e.g., at the parent or home office) and/or objective third-party expert resources. Model validation should take place independently of model development and ensure that models are conceptually sound, fit for their intended purposes, and understandable to relevant stakeholders. Additionally, outcomes of model validation should be documented.

Model validation activities should occur at various points during a model's lifecycle, including:

- development and implementation stages of new models
- when models are modified in response to monitoring outcomes
- in response to other internal requirements (e.g., inclusion of new data)
- periodically, to review model performance and affirm the model remains fit-for-purpose.

In consideration of a model's purpose, risk rating, and position within its lifecycle, validation should consider some or all the following activities:

- Evaluation of the purpose, scope and use of the model output
- Evaluation of the quality and appropriateness of model data
- Assessment of the model risk rating, conceptual soundness, model limitations, and corresponding mitigants
- Assessment of explanation provided for how a model produces outcomes
- Evaluation of the reasonability of model outcomes, model performance, and monitoring
- Verification that supporting documentation is complete.

Organizations that use third-party libraries, platforms, and/or automated development processes should be subject to independent validation, commensurate with the risks these elements present.

4.5 Approval

Approval requirements should apply throughout a model's lifecycle, including for modifications and periodic reviews.

Organizations should ensure that model validation occurs prior to approval. The model reviewer is responsible for providing to the model approver(s) the outcomes of its review, along with a recommendation. 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.

The model approval decision typically involves two components:

- assessing whether the model is suitable to be implemented into production (or continued to be used) based on its intended use.

- affirming the assigned model risk rating.

4.6 Model deployment

The deployment process should be a collaborative effort among model developers, model owners, model users, enterprise technology and operations partners responsible for managing staging and production environments.

Prior to deployment, organizations should ensure that the model outputs can be replicated in the production environment. Additionally, the organization should test the functionality and robustness of the production environment and the associated infrastructure.[8] Consistency between data used to develop the model and the production dataset should be confirmed prior to the model's release. Contingent actions should be developed for situations where the model is unavailable for periods of time, there is significant deterioration in the model's predictive properties, or if the model fails.

Organization should exercise adequate oversight over models obtained from third parties, including consideration of the third-party's development environment and model architecture.

4.7 Ongoing monitoring

Once a model is deployed, it should be subject to monitoring and periodic validation that is commensurate with its model risk rating. Monitoring intensity should dependant on characteristics of the model, with some requiring more frequent or customized monitoring.

Model owners have primary responsibility for monitoring and gathering input from other stakeholders, as appropriate. Where vendors are involved, model owners should ensure the products they deploy have adequate monitoring controls.

Monitoring results should be shared with model users in a timely manner and reflected in the model's risk rating as appropriate. Unsatisfactory results arising from monitoring should be subject to an escalation process whereby appropriate stakeholders are notified (including downstream parties impacted by the model). Moreover, such results should prompt a remediation plan.[9] Modifications initiated to remediate performance deficiencies should be subject to model validation and approval requirements, where appropriate.

4.8 Modifications and decommission

The modification[10] process reflects the iterative nature of the model lifecycle whereby models may undergo several rounds of revision before ultimately being decommissioned.

Informed by model risk ratings and potential impact of the changes, organizations should segregate modifications into different classes of importance and scale the intensity of redevelopment, revalidation and reapproval activities in a risk-based manner. Organizations should maintain adequate change records that align with approval records to prevent a divergence between the version most recently approved and the one used in production.

Organizations should maintain a catalogue of changes to the model calibration and define thresholds for what signals a material modification to the model. When a modification threshold is breached, the organization should re-evaluate the model, to determine if the model remains fit-for-purpose under the criteria of its most recent approval.

When a decision is made to decommission a model, the model owner should notify relevant stakeholders, including downstream model owners and users, about the upcoming decommission. The decommissioning of a model should not be considered the end of its lifecycle as a decommissioned model could act as a benchmark or might need to be recommissioned if the new model fails to be implemented properly (or perform up to minimum risk tolerances). For business continuity purposes, an organization should maintain the decommissioned model in its inventory for a reasonable period.

Organizations should be aware of and take risk-based actions for modifications and decommissions of any third-party models. For third-party model modifications, organizations should comply with the above minimum requirements as though the model were internally developed. Also, contingency plans should be developed for models with higher risk ratings in the event that vendor support is deemed inadequate or were to cease altogether.

5. Model risk management framework

Principle 3: Organizations establish an MRM framework that provides an enterprise-wide view of their exposure to model risk.

The MRM framework should reflect the organization's risk appetite for model risk and define the process and requirements to identify, assess, manage, monitor, and report on model risk throughout the lifecycle of models employed throughout the organization.

Organizations should regularly review and update their MRM framework to ensure it remains relevant and appropriate, and to make continuous improvements considering lessons learned. The MRM framework should, at a minimum, adhere to all principals included in this guideline.

Components of the MRM and its' interaction with Model Lifecycle are illustrated in the following diagram.

Model risk management framework - Model lifecycle

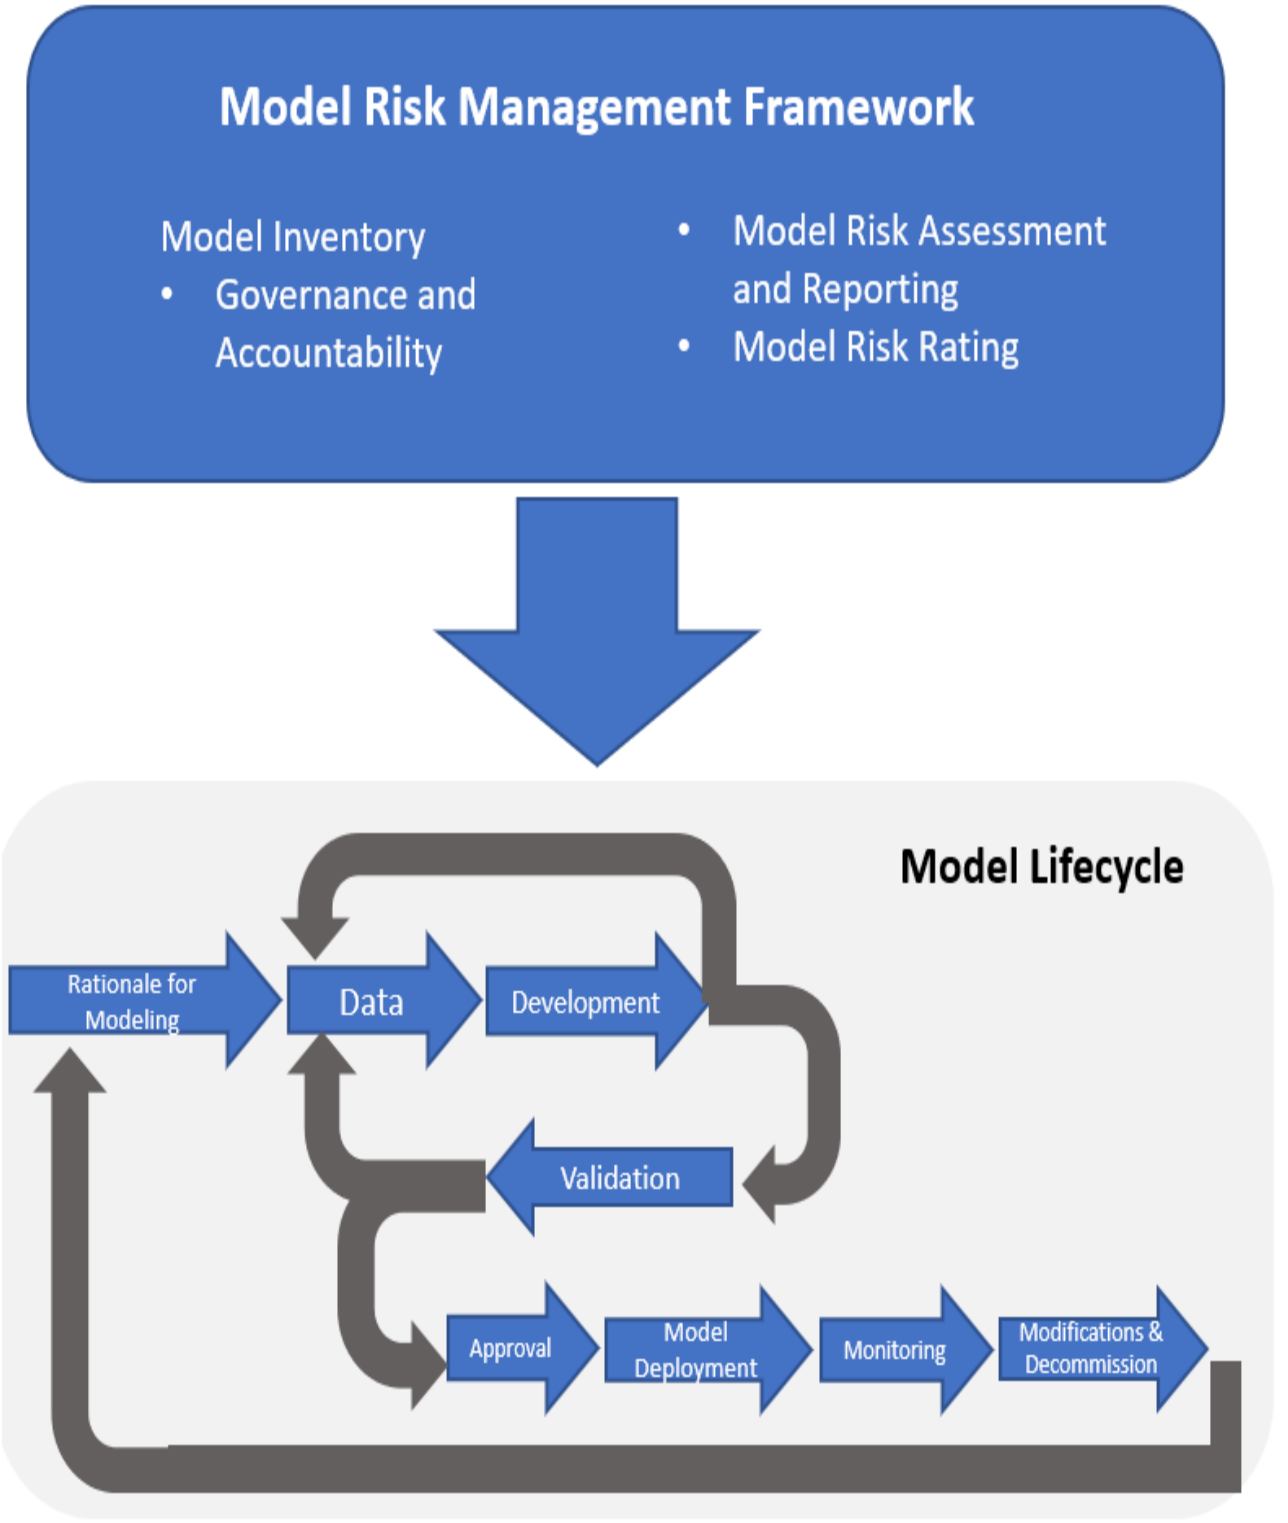

Figure 2 Model risk management framework - Model lifecycle - text version

The flow chart illustrates the different aspects of the Model Risk Framework: model inventory, governance and accountability, model risk assessment and reporting, model risk rating.

An arrow is pointing towards the model lifecycle.

Model lifecycle chart: The flow chart consists of a series of arrows connecting various stages, starting from "Rationale for Modeling" and progressing through "Data," "Development," "Validation," "Approval," "Model Deployment", "Monitoring", "Modifications and Decommission". The chart includes the double arrows from "Validation" to "approval" and back to "data" depicting cases where approval is not received. The chart also includes the double arrows from "development" to "validation" and back to "data" for cases where development is ongoing. Furthermore, the chart includes arrows from "Modifications and Decommission" back to "Rationale for Modeling" to depict that this is a continuous process.

5.1 Model inventory

Principle 4: Organizations maintain a centralized inventory that is the authoritative record of all models in use and recently decommissioned. The inventory should be evergreen and be subject to robust controls.

The inventory, maintained at the enterprise level, should also serve as a basis for management reporting. Updates to the inventory should by timely, including changes to reflect modifications to models, their risk classification, or updates to their performance status. An organization should implement risk-based controls to confirm the accuracy of its model inventory.

Information related to model inventories in Appendix A must be made available to OSFI on request.

5.2 Governance and accountability

Principle 5: Organizations have policies, procedures, and governing authorities for each phase of the model lifecycle, where expectations are established based on model complexity and importance.

Organizations should develop policies, processes, procedures, and governing authorities for each phase of the model lifecycle. Model developers and reviewers should have the requisite quantitative skills to conduct a model review and have knowledge of the business area for which the model is being used.

Organizations may obtain models or data from external sources like third-party vendors or foreign offices, which might have proprietary information. An organization that acquires a model from an external source (e.g., third party vendor, foreign office) is still expected to establish an MRM framework. Moreover, the organization retains ultimate accountability for outsourced activities and should secure adequate documentation from providers to understand the model's design, calibration and operation consistent with internally developed models.

To support effective MRM, policies should define exceptions[11] and establish thresholds that are aligned with the organization's risk appetite. An exception policy should provide for consistent identification, notification to appropriate stakeholders, and approval across model types. It should also detail circumstances that merit the removal of the model and/or imposing conditions that limit model usage. In granting an exception, the model approver should have the power to impose restrictions on the model's usage.

Principle 6: Organizations recognize the interdependency between data and model risk and have adequate policies and procedures to govern data in models. These policies and procedures should align with the organization's data governance framework and strategy at the enterprise level.

An organization's data governance policies and procedures should be integrated with and leverage, where possible, data governance and management requirements established at the enterprise level. Together, these should provide

a consistent approach to understanding and managing vulnerabilities, challenges, and changes to the data, including bias, fairness, privacy and other relevant considerations, particularly with respect to AI/ML techniques.

5.3 Model risk assessment and reporting

An organization should define metrics to facilitate the transparent and consistent monitoring of model risk at the enterprise level. It should periodically report, at minimum, the following to model owners, users, reviewers, and senior management:

- Model types
- Performance of individual models over their model lifecycle
- Description of the operating environment in which models are used
- Exceptions from the organization's MRM framework
- Enterprise-level assessment of model risk.

5.4 Model risk rating

> Principle 7: The model risk rating scheme considers both quantitative and qualitative criteria, as well as impacts to downstream processes.

Under the MRM framework, the organization should implement an appropriate model risk rating scheme applicable to all models. It should be designed to allow for consistent application across model types to facilitate the identification, assessment, management and reporting of model risk at the enterprise level.

In devising the model risk rating scheme, the organization should consider model risk using quantitative and qualitative criteria. Ultimately, the application of the model lifecycle requirements, documented in the MRM framework, should be commensurate with a model's risk rating. For example, the model risk rating may drive the level of authority required to approve the model, frequency and scope of model monitoring and independent review, and interval at which the risk rating is re-assessed.

Quantitative factors should include considerations such as the importance, size and growth of the portfolio that the model covers, or potential customer or financial impacts. Qualitative factors should include considerations that give rise to uncertainty such as, business use or purpose, complexity of statistical approaches applied, reliability of data inputs, or conclusions from the model review process.

Assigned model risk ratings should be regularly reviewed, including when a trigger event occurs[12] and updated as appropriate based on experience. In cases where model risk ratings fall outside the organization's risk appetite, the organization should establish appropriate remediation actions.[13] Models supplied to subsidiaries/foreign branches by a parent/home office should be assessed for model risk ratings on a standalone basis. The subsidiary/foreign branch should have access to technical documentation from its parent/home office to assess and manage the model's unique risk profile.

Appendix A

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

- Model ID
- Model version
- Model name and description of key features
- Model risk classification
- Identification of model stakeholders (e.g., owner, developer, reviewer, approver)
- Date of model's most recent validation
- Exception status
- Performance rating resulting from ongoing monitoring
- Model dependency (i.e., instances when the outcome of the model is an input into another model)
- Date of model's deployment into production
- Approved uses of the model
- Model limitation(s)
- Next review date

- Model origin (e.g., internally developed, vendor)

1 In the context of this guideline 'organization' is used to refer to a federally regulated financial institution (FRFI) and a federally regulated private pension plan (FRPP), as applicable.

2 While there are currently no generally agreed-upon definitions of AI/ML, OSFI adopted, for purposes of its Discussion Paper on Developing Financial Sector Resilience in a Digital World: Selected Themes in Technology and Related Risks (PDF, 5.6 MB) (September 2020), definitions of these two terms (see page 21 of this paper).

3 Underfunded defined benefit pension plan.

4 OSFI may consider smaller organizations, who employ complex models on which they have high degree of reliance, as more complex organizations.

5 Refer to section 5.4 Model Risk Rating section of this guideline for OSFI expectations.

6 Organizations should apply conservatism when determining provisional model ratings.

7 A dynamic model calibration is a model that can automatically adjust its own parameters or behavior in a production environment.

8 Depending on the nature of the deployment (e.g., new model versus modification), this may include system integration and/or user acceptance tests.

9 Depending on the nature of the deficiency, this may result in application of a judgemental overlay, acceleration of model modification timelines, granting a temporary exception, or placing restrictions on model usage.

10 Examples of modifications include inclusion of more recent data; introduction of a new data source; change in the technology/infrastructure used to supply the data or determine outputs; change in the underlying methodology; change in the model's operating environment.

In defining exceptions, organizations should recognize the integrated nature of the model lifecycle (for example, if a model does not adhere to requirements of the development policy, it may lead to a finding during the validation phase) and provide for appropriate classification internally. Examples of exceptions include: a model is used without requisite approval or outside its intended purpose; a model that displays

11 persistent breach of performance metrics (including backtesting); a model is overdue for its scheduled revalidation.

12 Examples of triggers that could prompt review are changes in underlying business environment, increases in the size or scope of a business line, asset transfer between pension plans, deterioration in model performance, material model modifications.

13 Examples of actions include modifying the model, increasing its monitoring frequency, increasing the frequency of its grading assessment, limiting its usage.