भारतीय रिज़र्व बैंक
RESERVE BANK OF INDIA
www.rbi.org.in

DoR.ORG.REC.xxxx/xx-xx-xxxx/2026-27 June xx, 2026

Guidance on Regulatory Principles for Model Risk Management, 2026

Table of Contents

CHAPTER – I - Preliminary

A. Introduction

1. Regulated Entities (REs) are increasingly using models to reap efficiency gains, transform their business processes, improve customer services as well as enhance risk management capabilities, including to defend against cyber attacks. This reliance has grown rapidly over the past few years owing to growing scale and complexity of financial activities, digitalisation of financial services, advances in analytical and computational capabilities, advent of new technologies like Artificial Intelligence (AI) and Machine Learning (ML), and provision of such models by third-party service providers. The significant benefits that arise from use of models, however, usually come with additional model risks. If not managed effectively, such risks may lead to inaccurate outcomes, flawed decisions, financial losses, operational disruptions, compliance failures and other adverse consequences for REs, consumers and the financial system. It is, therefore, necessary that the use of models is supported by appropriate governance, risk management and controls and continuous oversight.

2. Accordingly, this Guidance lays down a broad set of regulatory principles for the management of risks arising from use of models. It covers key aspects of model governance and risk management including model risk tiering, lifecycle management including validation, continuous oversight, change management and business continuity. It also includes specific broad principles for management of risks arising from third party models and models involving AI / ML. As indicated in paragraph I.10 of Utkarsh 2029, further requirements, if any, applicable to AI models may be issued later.

3. The Guidance is intended to assist REs in strengthening their model risk management framework in a manner commensurate with the nature, scale and complexity of their operations, and materiality and risk of the models used by them.

B. Applicability and Scope

4. This Guidance is applicable to following REs of the Reserve Bank of India:

(i) Commercial Banks (including Foreign Banks)

Wherein ‘Commercial Banks’ means all banking companies, corresponding new banks, and State Bank of India as defined under subsections (c), (da), and (nc) of Section 5 of the Banking Regulation Act, 1949.

(ii) Small Finance Banks;

(iii) Payments Banks;

(iv) Local Area Banks;

(v) Regional Rural Banks (as defined under Clause (ja) of Section 5 of the Banking Regulation Act, 1949);

(vi) Urban Co-operative Banks;

Wherein, Urban Co-operative Banks mean Primary Co-operative Banks as defined under section 5(ccv) read with Section 56 of Banking Regulation Act, 1949.

(vii) Rural Co-operative Banks;

Wherein, Rural Co-operative Banks mean State Co-operative Banks and Central Co-operative Banks, as defined in the National Bank for Agriculture and Rural Development Act, 1981.

(viii) Non-Banking Financial Companies in Base Layer (NBFC - BL), Middle Layer (NBFC - ML), Upper Layer (NBFC – UL), and Top Layer (NBFC - TL);

(ix) All-India Financial Institutions, viz., Export Import Bank of India ('EXIM Bank'), National Bank for Agriculture and Rural Development ('NABARD'), National Bank for Financing Infrastructure and Development ('NaBFID'), National Housing Bank ('NHB') and Small Industries Development Bank of India ('SIDBI');

(x) Asset Reconstruction Companies registered with the Reserve Bank under Section 3 of the Securitisation and Reconstruction of Financial Assets and Enforcement of Security Interest Act, 2002; and,

(xi) Credit Information Companies as defined under clause (e) of Section 2 of the Credit Information Companies (Regulation) Act, 2005.

5. The Guidance shall be read in conjunction with relevant Directions issued by RBI, as amended from time to time, or issued in substitution or succession thereto. In case of any inconsistency, the applicable Directions shall prevail.

6. An RE should apply these regulatory principles to all models used by it, whether developed internally, sourced from third-parties, or a combination thereof.

C. Definitions

7. In this Guidance, unless the context otherwise requires,

(1) '**Decommissioning**' means the process of retiring a model from active use.

(2) '**Explainability**' means property of a model to express important factors influencing its results, in a way that is understandable.

(3) '**Model**' means a system, whether developed internally, sourced from third-parties, or a combination thereof, that incorporates data, applies theoretical, empirical, or judgement-based assumptions (input component), uses statistical, mathematical, economic, financial, or such other cognitive techniques (including Artificial Intelligence (AI) / Machine Learning (ML)) to analyse, interpret relationships and process inputs (processing component) and produce results that are used for business or any other operations and decision making (output component). It includes algorithms, analytics, interfaces, applications, decision-based rules, and other computational tools which, by virtue of their use, have a material impact on decision-making in various business processes, irrespective of whether such tools are recognised as models by the RE.

Illustration - A spreadsheet-based loan pricing calculator deployed or used by an RE may be considered as only a basic mathematical tool. However, if the RE uses this tool to derive lending rates, customer margins, or credit terms, such that it takes inputs (borrower type, tenor, credit score, collateral value), applies processing logic (interest rate grids, risk-weighted spreads, margin formulas), and produces an output (final lending rate or price) which then affects business decisions, then it should be considered as a model.

(4) **'Model Approver'** means individual or function, responsible for undertaking approval process and granting approval for model deployment.

(5) **'Model Developer'** means individual or function responsible for designing, developing, testing and training, and documenting the model's methodologies.

(6) **'Model Owner'** means an individual or function, responsible for ensuring that the model's design, assumptions, methodologies, and documentation, are aligned with its intended use, regulatory requirements, and internal policies of the RE, and for coordinating across various stages of the model lifecycle.

(7) **'Model Risk'** means the risk of adverse outcomes of a model arising from, *inter alia*,

(i) Model errors such as inappropriate specification, incorrect parameterisation, flawed hypotheses and / or assumptions, computational errors, inaccurate /

inappropriate / incomplete data, inadequate controls, or other issues in development and inadequate validation.

(ii) Misapplication (improper or unintended usage / misinterpretation of outputs).

(iii) Time-suitability issues (models becoming less fit / unsuitable over time).

(8) **'Model Validator'** means an individual or function, independent of model development or ownership or use, responsible for carrying out model validation, to ascertain whether it is fit, efficient, and serves the intended purpose.

CHAPTER – II - Governance

8. An RE is accountable for the outcomes of all models used by it, irrespective of whether the models are developed internally, sourced from third-parties, or a combination thereof.

A. Model Risk Management Framework (MRMF)

9. An RE should put in place a Board-approved MRMF applicable to all models, including AI / ML models, irrespective of whether such models are developed internally, sourced from third-parties, or a combination thereof.

10. The MRMF should, *inter alia*, cover taxonomy for models; governance structure; scope of model usage; model risk tiering methodology; inventory and documentation standards; and concomitant policies covering the model lifecycle including standards and procedures for model selection and development, validation, approval structure including exceptions and risk mitigants thereof, deployment and monitoring, change management, business continuity management, and decommissioning.

B. Role of the Board

11. The **Board** should be responsible for the oversight of an RE's MRMF, which should, *inter alia*, cover–

(1) approval and periodic review of the MRMF including delegation to the Risk Management Committee of the Board (RMCB) and any other Committee, as required;

(2) approval of the RE's risk appetite and tolerance for model risk and ensuring that they are forward-looking and informed by scenario analysis / stress testing; and,

(3) approval of the RE's policies for model risk management including for model risk tiering.

C. Role of Committees of the Board

12. The RMCB should oversee the implementation of the MRMF and its ongoing compliance. It should:

(1) review validation reports of models with 'high' or equivalent risk as per the model risk tiering and approve their deployment;

(2) review model risk tiering reports periodically as per MRMF, but at least annually;

(3) oversee monitoring of models approved with exceptions, third party models, and models involving AI; and,

(4) review reports of breaches and other material concerns, if any.

D. Role of Senior Management

13. The Senior Management should, *inter alia,*:

(1) establish procedures / processes and ensure allocation of human and technical resources for operationalisation of the MRMF and compliance thereof;

(2) implement the risk based tiering structure for models;

(3) ensure maintenance and regular updation of model inventory and documentation; and,

(4) ensure periodic review of the policies and procedures / processes under MRMF and report to the RMCB.

CHAPTER – III - Model Risk Management

14. An RE should assess model risk at both individual and enterprise-wide level, on an ongoing basis. In case, the assessed risk of a model exceeds the RE's risk appetite, it should initiate timely action (e.g., enhanced controls, restrictions on use, remediation, decommissioning) and a report to this effect should be placed before the RMCB.

15. It should implement the three lines of defence with model owners being first line of defence, an independent model risk management and validation function being second line of defence, and a robust and independent internal audit function being third line of defence.

16. It should undertake ongoing performance testing using backward-looking and forward-looking approaches, including AI specific evaluations where applicable, and benchmarking, as appropriate.

A. Risk based Model Tiering

17. An RE should establish and implement a risk-based model tiering structure, for classification of all models in the inventory. It should review the risk tier of all models at least annually, or earlier, as specified in the MRMF or in response to specific triggers.

18. The risk tier of a model should be used to guide processes including, inter alia, –

(i) validation prioritisation, its intensity, frequency, and methods and techniques;

(ii) approval structure, i.e., models with 'high' or equivalent risk should require approval from the RMCB; and other models may be subject to delegated approvals;

(iii) risk mitigation and controls;

(iv) scope of monitoring, reporting, and review;

(v) details in the inventory and documentation; and,

(vi) business continuity planning.

19. The model risk tiering should, inter alia, be based on materiality of the model (e.g., significance of model to the RE's business processes, its impact on the RE's financial and operational outcomes, and its potential implications for consumers), complexity of the model (e.g., the degree of difficulty in understanding and exercising effective oversight, use of unstructured data, challenges relating to its explainability), and other relevant factors (e.g., regulatory or supervisory considerations).

20. The RE should ensure that the integration of multiple factors does not result in one factor offsetting or diluting the other, and the model tier should represent composite risk profile of the model (e.g., a low complexity should not result in a disproportionate reduction of the overall risk tiering of a highly material model).

B. Model Inventory and Documentation

21. An RE should maintain accurate, comprehensive, and up-to-date inventory of all active, inactive (including under development), and decommissioned models to enable it to have an overview of individual and enterprise-wide model risk, serve as basis for management reporting, and help to identify model inter-dependencies. It should ensure that no model is used, relied upon, or deployed unless it is part of inventory.

22. The inventory, at minimum, should include key details such as model owners, developers, validators, and approvers; risk tier; intended use; dependencies with upstream and downstream models; and key observations from validation, monitoring, and audit.

23. The decommissioned models should be a part of the inventory for at least ten years from the date of decommissioning or the date they cease to serve as backup or benchmark reference, whichever is later, or such longer period as required under applicable law.

24. An RE should develop comprehensive documentation for all models, including third-party models. The minimum documentation period for all models should be in alignment with their retention period in inventory.

C. Consumer Protection and Grievance Redressal

25. An RE should not use any model that harms consumer. Its grievance redressal mechanism should also address grievances arising from consumer facing models used by it.

CHAPTER – IV - Model Lifecycle Management

A. Model Selection and Development

26. (1) An RE should, prior to initiating model development, define and document the rationale and objectives of using the model, and its scope of application.

(2) It should consider the costs (e.g., additional risks, likelihood of potential adverse outcomes, fairness, ethical considerations, and biasness) and benefits of introducing or replacing existing processes with the model.

27. The model development should follow a structured and systematic process, aligned with its intended use and output, which should, *inter alia*, include, collection, pre-processing, and transformation of data, assessment of assumptions and limitations, design of the model, evaluation and refinement of model performance.

28. The data used for model development, empirical or synthetic, should be as per data governance processes of the RE.

B. Model Validation

29. An RE should ensure that all models, including third-party models, are subject to independent validation by the RE.

30. The validation should be conducted prior to and after deployment of a model, following modification, on internal or external triggers, and periodically as specified in the MRMF.

31. It should include, *inter alia*, assessment of inputs (e.g., data, assumptions and limitations), soundness (conceptual and design), performance, and alignment with intended use.

32. The model validation outcomes should be documented as per documentation standards laid down under MRMF / its policies.

33. Validation reports, including key findings, and recommendations, should be placed before RMCB, or delegated authority as specified in MRMF, within three months of completion of the validation.

C. Model Approval

34. An RE should have an approval structure, including exception approval that, *inter alia*, covers approval authorities, thresholds, additional requirements for models approved with exceptions, and remediation timelines.

35. The decision-making process and exception approval for a model should be documented and should include the rationale for decision / approval.

D. Model Deployment and Ongoing Monitoring

36. The deployment of a model should be undertaken in coordination with all relevant stakeholders, including the RE's IT and data functions. An RE should ensure that model outputs are replicated and stable in production environment.

37. All deployed models, including third-party models, should be subject to ongoing monitoring, to ensure alignment with the intended outcomes. The monitoring should, *inter alia,* include requirements, if any for model modification, replacement, or extension beyond original scope. The models approved with exceptions should be subject to enhanced monitoring by RMCB.

E. Change Management

38. An RE should have a structured process for change management, which should cover roles and responsibilities for carrying out and approving changes.

39. It should ensure that any change is implemented in a controlled manner and at enterprise level, with necessary mechanisms to recover from failed changes or unexpected results.

40. Any change to a model, should be preceded by a documented impact assessment including the continued suitability of the model for its intended use.

41. An RE should maintain a comprehensive record / log of changes and versioning, and approvals.

42. An RE should define the threshold / criteria of what constitutes a material change, a breach of which should re-initiate the process for validation and approval.

F. Business Continuity Management and Decommissioning

43. The continuity planning for models should form part of an RE's overall Business Continuity Planning policy / document. It should, *inter alia,* include potential disruptions (e.g., model unavailability, performance degradation, or failure) and fallback mechanisms (e.g., manual interventions, substitution, or back-up arrangements).

44. An RE should ensure that all relevant stakeholders are informed of decommissioning of a model in a timely manner for enterprise-wide transition.

CHAPTER – V - Specific Models

A. Third-Party Models

45. An RE acquiring, using or relying upon third-party models at any stage of the model lifecycle is accountable for its outcomes.

46. All provisions of the MRMF should apply *mutatis mutandis* to third-party models. These should additionally be subject to:

(i) independent validation by the RE in accordance with paragraphs 29 to 33 notwithstanding any validation, certification, or assurance provided by the third-party provider; and

(ii) enhanced oversight by the RMCB, irrespective of their risk tier.

47. Prior to acquisition or use of a third-party model, an RE should undertake due diligence which should, *inter alia,* include credibility of the service provider, methodological soundness of the model and its limitations, and the suitability and quality of data used.

48. An RE should ensure that contractual arrangements governing third-party models include provisions, *inter alia,* for access to minimum technical documentation that should give reasonable understanding on design, configuration, assumptions and operation of the model and be sufficient to validate the model as per RE's MRMF; audit rights for the RE and its supervisory authority either directly or through external experts engaged by them, and continuity and exit arrangements.

B. Models Employing Artificial Intelligence / Machine Learning

B.1 Risk Management

49. An RE should define the scope of AI / ML model, including for foundational AI models and frontier AI models, and put in place additional controls, commensurate with its potential impact on customers, business operations, and financial outcomes.

50. It should assess whether the risks arising from such models can be adequately identified, measured, monitored, and managed. It should ensure that such AI / ML models are deployed only in the business processes / use cases where commensurate risk can be effectively managed.

51. In cases where the third-party provider does not disclose adequate information regarding the AI / ML model, the RE should identify risks that arise from such constraints, and put in place the necessary mitigants, such as limiting the usage.

52. For assessment of risk and assigning risk tier to an AI model, an RE should in addition to the principles laid down in paragraphs 19 and 20, consider the extent of reliance and the level of autonomy placed on the model outputs for decision-making.

53. For material third-party AI models, datasets, and dependencies deployed by the RE, it should consider additional risks arising from dependence on a limited number of model providers including supply chain risk, limitations in independent validation, and changes in model behaviour or capabilities resulting from provider-driven updates.

54. An RE should identify and address risks arising from the behavioural characteristics of AI models and implement appropriate safeguards. It should test the model behaviour under atypical or stressed scenarios to ensure vulnerabilities do not arise under edge cases, abnormal inputs, manipulations, and adversarial conditions.

(1) (i) It should define the explainability and transparency thresholds for all AI models and ensure that their outputs are explainable to the extent required for the business process. It should apply higher thresholds for explainability to models which are relied upon for material decision-making or have significant impact on customers or its operations.

(ii) Where full explainability is not achievable, the RE should ensure that such models are subject to enhanced risk management measures and controls, including enhanced validation and testing, mechanisms to verify and corroborate model outputs prior to their use, frequent validations and continuous monitoring, usage restrictions, and other compensating controls necessary to mitigate risks arising from limited explainability.

(2) It should put in place appropriate control boundaries through system-level controls or model design features to mitigate risks of hallucinations, particularly in models capable of generating content (e.g., generative AI models) and use cases where the model outputs directly or indirectly drive customer interaction or decision-making.

(3) It should identify the risk of bias, and discriminatory outputs, specifically in use cases such as unfair treatment of certain customer groups. Further, for such type of models, it should conduct fairness assessment and implement appropriate mitigants, including recalibration or redesign. For complex models, the RE may

consider constraining complexity (e.g., regularisation of AI models) and limiting feature selection to mitigate such risks.

(4) It should ensure that models are not overfitted to training data and are capable of appropriate generalisation. The RE should assess the model performance with out-of-sample data and varied scenarios and ensure model's ability to perform reliably in real-world and evolving conditions.

(5) It should ensure that models do not rely on spurious correlations or unintended relationships that may adversely affect outcomes.

(6) It should ensure that model outputs under similar inputs should not exhibit excessive or unexplained variation. The risks arising from variability in outputs, stochastic behaviour, and model uncertainty should be appropriately managed by measures such as confidence scores and probability outputs.

(7) An RE should establish appropriate mitigants to address the data risks such as data quality, non-representativeness, incompleteness, breach of intellectual property rights. Changes in data distribution, including data drift and concept drift, should be monitored and addressed on an ongoing basis.

55. An RE should put in place structured challenge processes, including red-teaming or equivalent testing, particularly for models involving customer interaction or generative capabilities.

56. An RE should implement enhanced controls for models with dynamic or automatic updates, including defining a clear scope of what can be updated automatically, strict justifications for enabling automatic updates, enhanced data quality checks, and more stringent and frequent monitoring.

57. An RE should have enhanced documentation for the AI models considering their complexity, self-adapting nature, and huge reliance on training data, to enable traceability, reproducibility, and auditability.

B.2 Model Deployment Controls

58. An RE should ensure that deployment of AI models do not introduce vulnerabilities in the model or the RE's production environment and should implement appropriate safeguards, which should, *inter alia,* cover:

(1) access controls to prevent unauthorised access, use or modification;

(2) safeguards against cyber risks; and

(3) controls for risks arising from external interfaces or APIs or integration pipelines with third-party components or systems.

59. In respect of models, including generative AI models, having interface with customers or external users, -

(i) additional cyber security controls should be implemented, such as controls against prompt injection and adversarial inputs, limitations on session and context persistence, and detection of anomalous usage patterns.

(ii) appropriate disclosures and warnings should be provided to the users that they are interacting with AI / ML based system, along with limitations of such systems.

(iii) option should be provided to the customer to switch to human assistance when requested for.

B.3 Human Oversight

60. An RE should establish robust human oversight for AI models including use cases involving automated decision-making by models. It should establish appropriate risk mitigants which *inter-alia* include:

(i) Human-in-command arrangements (e.g., human-in-the-loop / human-on-the-loop / other human oversight mechanisms);

(ii) override, suspension, or deactivation mechanisms, including kill-switch arrangements; and,

(iii) periodic review of model outputs and model-driven decisions by humans to identify anomalies.

61. The oversight mechanism should also consider risks arising from automation bias, over reliance on model outputs and decision fatigue.

62. It should ensure that personnel involved in oversight possess adequate expertise and understanding of model functioning and are able to effectively challenge, override, or escalate issues / concerns in model outputs where required.

63. It should ensure that human oversight arrangements, including decisions, interventions, overrides, incidents and near misses, are periodically reviewed and strengthened based on experience.

CHAPTER – VI - Other Provisions

64. The final Guidance on 'Regulatory Principles for Model Risk Management', following public consultation, would supersede Chapter-3 on Credit Risk Models of the 'Guidance Note on Credit Risk Management, dated October 12, 2002.