



TC NOTES

PRACTICAL **LEADERSHIP**
AND **GUIDANCE** FROM
TORONTO CENTRE

SUPERVISION OF BANK MODEL RISK MANAGEMENT

JULY 2021

SUPERVISION OF BANK MODEL RISK MANAGEMENT

TABLE OF CONTENTS

Introduction	2
Model Risk	3
Supervisory Standards for Models	4
Home-Host Coordination of Supervision	6
Overview of Models	6
What is a Model?.....	7
Model Accuracy	7
Model Overfitting.....	8
Rule Sets	8
Three Lines of Defence Against Model Risk	8
Model Development	9
Model Purpose	9
Consideration of Alternative Approaches	10
Data Preparation	10
Model Selection.....	11
Model Monitoring Plan.....	12
Model Documentation	12
Model Development Process.....	12
Model Validation	14
Independent Model Validation.....	14
Vendor Models	16
Independence.....	16
Internal Audit	16
Model Governance	17
Machine Learning Models	17
The Advent of ML	17
Supervision of ML Models	18
Conclusion	20
References	21

Copyright © Toronto Centre. All rights reserved.

Toronto Centre permits you to download, print, and use the content of this TC Note provided that: (i) such usage is not for any commercial purpose; (ii) you do not modify the content of this material; and (iii) you clearly and directly cite the content as belonging to the Toronto Centre.

Except as provided above, the contents of this TC Note may not be transmitted, transcribed, reproduced, stored or translated into any other form without the prior written permission of the Toronto Centre.

The information in this TC Note has been summarized and should not be regarded as complete or accurate in every detail.

SUPERVISION OF BANK MODEL RISK MANAGEMENT

Introduction¹

Banks use models to inform business decisions across every aspect of their operations. They use models to inform credit decisions, such as whether to accept or reject a credit card application, and to inform larger scope decisions, such as asset-liability management. This is particularly true of large and complex banks where financial management is not feasible without the accuracy, efficiency and speed of information provided by models. As a consequence, banks are exposed to model risk, namely the risk that model results lead to incorrect decisions and financial loss to the bank.²

This Toronto Centre Note discusses the supervision of model risk management by bank supervisors. Guidance on model risk management issued by US bank supervisors (Board of Governors et al. 2011), the Basel Committee on Banking Supervision (2012) and the European Central Bank (2019) serve as points of reference. Supervision of model risk can differ across jurisdictions in terms of implementation, but the underlying issues and possible remedies are the same. The principles presented in this Note apply to models and to the supervision of model risk in both developed and developing countries.

Supervision of model risk management is tailored to a model's purpose, varying from *compliance-based approaches* to *risk-based approaches*. Compliance-based supervision is applied to models used to comply with laws and regulations – for example anti-money laundering, regulatory capital requirements and supervisory stress tests. Risk-based supervision of model risk is used for models that inform business decisions regarding individual transactions, portfolio management and strategic decisions that are not prescribed by laws and regulations. The supervision of model risk management should differ with model purpose. Supervisors using risk-based approaches to supervising model risk may issue guidance to banks that communicate supervisory expectations and recommendations for the management of model risk. Compliance-based supervision is considerably less flexible than risk-based supervision because it enforces banking law and regulation. Differences in banking laws and supervisory practices can result in differences in how supervisors approach the supervision of bank model risk management.

Despite the differences between risk-based and compliance-based approaches to the supervision of model risk, both approaches should consider all stages of the *model life cycle*, namely model development, independent validation, implementation, monitoring, change and eventual retirement. The supervision of model risk should also include the governance framework that determines how bank decisions that concern models are made - strategic planning, organizational goals and responsibilities - and the model risk management program that implements the governance framework for day-to-day model risk management activities and resource allocation.

¹ This Toronto Centre Note was prepared by John P. O'Keefe.

² Other types of financial institutions also use models to measure and manage risks. This Note focuses on banks, but many of the concepts and their implications for supervisors are common across all types of financial institution.

Model Risk

Model risk can arise from a wide variety of situations, including:

- **the use of a modelling approach that is not appropriate for the model's intended use;**
- **incorrect assumptions used during model development;**
- **inappropriate and/or inaccurate data used for model estimation;**
- **incorrect standards for model accuracy relative to the model's intended use; and**
- **mistakes in model implementation that cause the implemented model to differ from the initial model.**

The potential for model risk is increased when existing models are applied to new situations. For example, a common problem for banks is how to measure the risks of new products and customers where the bank has no historical experience to use in model development. In such situations, it is unlikely that bank models based on existing products and customers can be safely applied to new products and customers. Rather, banks should develop separate models for new products and customers that rely on external data (for example, loan default and loss data from peer banks and external data on customer creditworthiness such as credit bureau information).

Model risk can also arise when models become outdated, typically because the processes they are trying to represent change over time. A good example of this type of change occurred in the early-to-mid 2000s when banks in the US and in some other countries began lowering the lending standards for residential mortgages. Prior to that time, US mortgage borrowers could borrow up to 90-95 percent of the value of the property, and were required to purchase primary mortgage insurance for the bank until such time as the borrower had at least a 10 percent equity interest in the property (either through loan repayments or price appreciation). But in the 2000s, lenders dropped those requirements, as well as other prudential standards, and no longer required borrowers to put down a deposit on the property or to take out mortgage insurance. The collapse of residential real estate markets during 2007-2009 and the global economic recession led to a high rate of residential mortgage defaults, severe losses on loans and a banking crisis.

Leading up to 2007, bank models for mortgage default risk and loan losses were based on the experience of the 1990s when lending standards were sound, and therefore they underestimated the impact of falling housing prices on loan losses. These same model deficiencies were evident in the models that banks use to establish their capital requirements (Basel II capital requirements), which resulted in the inadequate capitalization of many banks and contributed to bank failures.

While this Note focuses on individual models, supervisors recommend that banks manage their *aggregate model risk* as well - the combined risk due to all the models used by the bank. The use of common data, assumptions and processes for model development across multiple models can increase aggregate model risk if there are errors in these model development resources. In addition, models sometimes use the results of other models as inputs, which can increase aggregate model risk due to the interactions and dependencies among models.

Banks can address model risk in several ways. Risks from individual models can be mitigated by the robust testing of models for their sensitivity to variation in input data and the robustness of their results to new situations and data. Banks can also place limits on how model results are

used when model accuracy is weak. For example, limits on credit lines extended to new credit card customers can help limit risk exposures resulting from accepting potentially risky new card customers. Banks can also reveal the extent of uncertainty in model output by reporting confidence intervals around predicted values and/or a range of predictions rather than a “spot” estimate based on the mean of the confidence interval. *Conservative overlays* or model “add-ons” are another tool for mitigating model risk. Examples of conservative overlays are using model estimates that result from assumptions, data and scenarios that are more adverse or pessimistic to the bank, and adding an additional “buffer” to model results, such as increasing predicted loan default rates by a certain percentage.

Aggregate model risk can be addressed by strengthening bank model risk management and the internal oversight of model risk management, and also through an additional capital charge designed to absorb potential losses due to model risk.

Banks sometimes find that model performance deteriorates over time. This can be due to changes in the underlying data and market conditions that can make the model “out-of-date”. Since the development of a new model can take several months, banks often use *model tuning* as an interim solution to address weaknesses in models. An example of model tuning is adjusting all predictions by a certain fixed amount to account for changes in the average response in the market, for example to adjust all default estimates upward by 5 percent. Model tuning is not the same as a model overlay because the adjustments are based on actual outcomes experienced by the bank over time rather than by ad hoc conservative adjustments.

Conservative overlays and model tuning should be used with caution and are designed to be temporary means for addressing risks in weak models. While model overlays and tuning might appear to introduce conservatism in model use, a weak model can in practice be so misleading that conservative add-ons and model tuning produce overly optimistic predictions.

Bank supervisors should understand the practical necessity of temporary mitigants for model weaknesses but should also be aware of the pitfalls of relying on what might appear to be “excessive conservatism” in model use. If a model is severely inaccurate the bank’s model risk management program should address the situation through remedial measures such as model re-estimation or redevelopment within a reasonable time frame.

Supervisory Standards for Models

Supervisory expectations for bank model risk management should be based on the *principle of proportionality*. Supervisory expectations for the bank resources devoted to a model’s development and governance should be proportional to the risk exposures the model is used to manage, the model’s complexity, the extent to which the bank relies on the model, and the potential costs of model risk associated with the model.

Supervisors should adjust supervisory expectations for bank model risk management based on bank size, complexity and the importance of a model to an individual bank. As such, as a bank increases in size and complexity its supervisor should seek to ensure that the bank’s model risk management program adjusts appropriately to changes in model risk.

This raises the question of how supervisors should set expectations for bank model risk management under the *principle of proportionality*. One important source of information for this question is provided by the bank’s independent validation unit which should rate models in

terms of the model risk they pose for the bank and use these ratings to set expectations for model development and to determine the depth and scope of validations and frequency of re-validations.

Banks can use a variety of criteria for risk-rating models. A model risk matrix can be designed to rate models as low, moderate or high on scales for both inherent model complexity and materiality. A bank should also be able to provide a rationale and support for the assigned model risk ratings. If the supervisor concurs with a model's risk ratings and with a bank's independent model validation procedures based on risk ratings, then supervisory expectations regarding the principle of proportionality for model risk management should be met. The important point here is that supervisors cannot prescribe standards for applying the principle of proportionality without information on the bank and the model under consideration. There are too many factors to consider allowing a uniform approach to this question across banks, so risk-based supervision should be applied in these judgements.

Supervisory expectations for the model development process and its documentation should, in general, be higher for new models than for the re-estimation and redevelopment of existing models. Initial model development should involve extensive consideration of alternative modelling approaches; however, extensive investigation of alternative approaches is unnecessary when the model undergoes periodic re-estimation due to updates to model input data. This is not to say that bank supervisors should not ask banks to reconsider periodically the overall efficacy of the modelling approach. For example, bank supervisors may require banks to conduct *annual model reviews* that look specifically at whether the model is operating as expected and whether existing validation procedures are effective. More generally, banks' model risk management programs need to consider advances in statistical and modelling software, new data, advances in existing modelling approaches, and the development of new modelling approaches, as part of either periodic model redevelopment or annual model reviews.

A common concern among banks is the amount of time and resources needed to align bank model risk management, overall and for individual models, to supervisory expectations. Bank supervisors may use a gradual or phased-in approach in their treatment of models that result from new laws and regulations or are newly applied to a bank. Bank supervisors need time to communicate their expectations regarding models that are new to a bank since model risk management guidance cannot answer many questions that arise when the guidance is applied to specific models. Further, banks need time to develop the capacity to develop and manage new and often complex models that are required to fulfill legal and regulatory requirements.

Bank supervisors typically approach model risk by first considering a model's purpose and materiality. Second, bank supervisors should review individual models on a rotating basis so that all material models are reviewed over the course of a supervisory cycle. Individual model reviews are most effective when linked to the specific area of bank operations being examined. For example, mortgage loan prepayment and default models are typically reviewed during the assessment of banks' mortgage lending policies and procedures. Individual model reviews can also be used for transaction testing during the assessment of a bank's model risk management program, to evaluate whether bank policies have been put into practice.

Supervisory assessments of individual models should consider the model development process, model use and monitoring, and the bank's model risk management. These areas of review can be based on supervisory guidance on model risk management and general principles for sound statistical modelling. Bank supervisors should evaluate the risk of individual models not only by examining the model's appropriateness, accuracy and robustness but also by examining the

integrity of the model development process and the governance of that process. As a consequence, bank supervisors should place considerable emphasis on the documentation of model development and the governance of model risk.

Home-Host Coordination of Supervision

An important issue in model risk management, and bank supervision in general, is the supervision of international banking groups. Internationally active banks are a significant presence in banking markets and can make up a substantial share of banks in developing countries. The headquarters and senior management of international banks are located in the “home” country while affiliated banks and bank branches can reside in a “host” country as well as the home country. IMF (1998) discusses the supervisory challenges for international banking organizations, pointing out that the cross-border expansion of banks can stimulate competition and encourage greater efficiencies in host country banks. “However, cross-border expansion can create a variety of difficulties for supervision. This is particularly true for emerging market countries that are still developing their accounting or legal systems, and where supervisory resources are limited.”

With regard to model risk management, the supervision of cross-border banks by the host country can be challenging for several reasons. First, international banking groups usually centralize their model risk management through policies that apply to the entire group. Second, model development and independent validation are often centralized, residing in the home country. Third, should the host country need to obtain information on models, or speak with model developers, independent validators and bank management located in the home country, there may need to be bilateral agreements that facilitate information sharing. Finally, the host country might lack the resources and expertise to conduct model reviews of cross-border banks.

IMF (1998) discusses best practices for the cross-border supervision of banks, and states that the host supervisor should obtain any information necessary for the prudential supervision of cross-border banks from the home supervisor. The IMF also encourages host supervisors to allow home supervisors to inspect the cross-border operations of its banks. Basel Committee (2007) discusses principles for information sharing between home and host supervisors in the context of the consolidated supervision of bank capital adequacy under the Basel II Advanced Measurement Approach.

Basel Committee (2012) discusses core principles for effective banking supervision and discusses information sharing between home and host supervisors in Core Principle 13: “Home and host supervisors of cross-border banking groups share information and cooperate for effective supervision of the group and group entities, and effective handling of crisis situations. Supervisors require the local operations of foreign banks to be conducted to the same standards as those required of domestic banks.” The international coordination of cross-border bank supervision can also be assisted by memoranda of understanding and other bilateral agreements.

Overview of Models

Banks analyze information on all aspects of their operations. These analyses inform individual credit decisions, product pricing, asset-liability management, portfolio management, strategic planning, compliance, stress testing and many other activities. These analyses often rely on models to provide metrics that inform business decisions.

The purpose of this section is to guide supervisors through an assessment of a bank's model risk management.

What is a Model?

Models are simplified representations of entities, objects and processes. Models vary in depth and scope; however, it is important to remember that models are not intended to be complete representations of entities, objects and processes. Models focus on the features of what is being described to align the model with its purpose.

Models can be categorized as being either qualitative or quantitative in nature. *Qualitative models* are collections of stylized facts (expert judgements) used to describe an entity, object or process. For example, bank management might have learned through experience that offering uncompetitive interest rates on time deposits can result in a decline in time deposit balances during the course of a year. This is a qualitative model since the process is described in general terms. *Quantitative models* – which are most relevant to this Note – are based on statistical analysis of data, algorithm-based data analyses or a combination of both, and are comprised of the set of assumptions, observations on variables (data) and relationships between explanatory variables and outcomes that describe an entity, object or process.

Banks use models to represent processes - narrow processes such as the determinants of customer default on loans, and broad processes such as the impact of macroeconomic stress on a bank's earnings and capital. The vast majority of bank models can be further classified as predictive models. Predictive model results can take the form of predictions of the likelihood of an event (for example a loan default), a financial outcome (for example the annual loss on a loan portfolio), or both. In banking, predictive models are used in the credit scoring of loan applicants, loan prepayment and default risk estimation, external fraud detection, stress testing, capital adequacy assessments and a host of other activities.

Business tools, such as spreadsheet calculators, that involve simple arithmetic are not considered to be models by many bank supervisors. Business tools that do not have the key elements of a model - assumptions and relationships between input data and outcomes that are determined by an objective function - lack the content for a model review by supervisors. Business tools should, however, be subject to internal controls and audits to ensure their accuracy and security.

Model Accuracy

The objective of all models is to describe entities, objects and processes. How model accuracy is measured depends on the model's purpose. For example, the accuracy of the outcomes of predictive models can be measured in two ways - rank order and value accuracies. If predictive models are used for credit card applicant accept/reject decisions, rank order accuracy is an appropriate criterion for measuring model performance. If, however, predictive models are used to forecast the effect of interest rates on bank deposit growth, the accuracy of forecasted deposit growth (outcome value) is the most appropriate model performance criteria.

Model accuracy also applies to how models describe entities, objects and processes. Specifically, model developers seek explanatory variables that have intuitive relationships with the dependent variable. For example, a loan default prediction model might include credit cardholders' credit bureau scores and a monthly income as explanatory variables; model developers expect credit card default risk to increase as a cardholder's credit score and income decline.

Model Overfitting

The goals of model accuracy and explainable relationships between dependent and independent variables are not only compatible but reinforce one another. To understand why, consider a common problem in model development in which some explanatory variables are found to be useful in outcome prediction in the model development dataset but are not useful for prediction when the model is applied to new data and situations. This can occur when some explanatory variables have *spurious relationships* with the dependent variable in the model development data sample - relationships that are extremely unlikely to extend to other periods, locations and situations.

Spurious relationships are, by their nature, difficult to explain rationally. For example, models that suggest that income is negatively related to credit card defaults in some time periods and positively related to defaults in other periods are described as being inconsistent and spurious relationships. While spurious relationships can improve model accuracy in the model development data, model performance will undoubtedly worsen when the model is applied to new and different data. This phenomenon where model performance in the model development data sample is improved by the inclusion of spurious relationships is known as *model overfitting*. The risk of model overfitting can be mitigated by excluding explanatory variables with counter-intuitive relationships with the dependent variable in the model development sample and by testing models over datasets that were not used to develop the model (out-of-sample testing), before selecting the final model specification.

Rule Sets

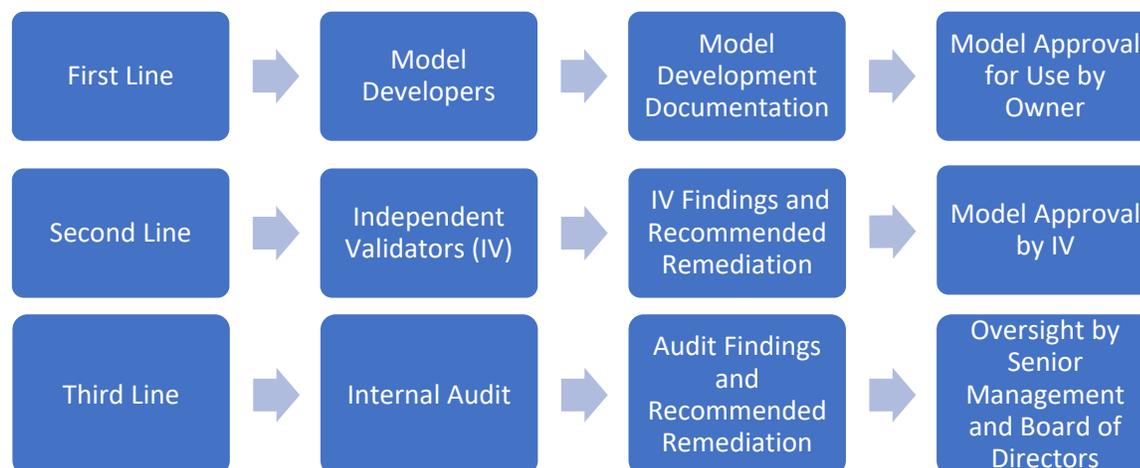
Bank management and supervisors should be aware of the limitations of model-based risk measures. Foremost among these are limits to models' flexibility and scope. For example, an external credit card fraud detection model might flag multiple online computer equipment purchases by one customer in a short time interval as potentially fraudulent and initiate contact with the credit card holder to verify the purchase orders. To limit false suspected fraud alerts, banks create business rules (*rule sets*) that are used to qualify and sometimes over-ride decision rules based on models. For example, multiple online credit card purchase orders for computer equipment by a customer might be accepted by the bank as lower risk during a holiday season when gift-giving might explain such purchases. A bank's rule set for a single activity, such as the approval of credit card purchases, can be comprised of hundreds of simple questions or filters that can be used to amend and/or over-ride model-based decisions. For example, a single rule might ask whether the customer has a very high credit line or whether the customer has earned a high number of loyalty reward points on the credit card. As is the case for business tools, rule sets should be subject to internal review and audit by the bank to assess their effectiveness and security, but this will not be a full model review.

Three Lines of Defence Against Model Risk

In many banks the management of model risk is comprised of three lines of defence (see Figure 1). Supervisors should review the effectiveness of these lines of defence, and the effectiveness of the oversight of model development, validation and usage by a bank's senior management and Board of Directors.

The first line of defence comprises the model developers who build, validate and monitor models and the model owners who accept responsibility for models. The second line of defence is comprised of independent model validators who evaluate the work of model developers and the models themselves. The third line of defence is the bank's internal audit function that reviews whether the first two lines of defence are complying with the bank's policies, practices and procedures for managing model risk, and complying with relevant laws and regulations. All of this should be overseen by the bank's senior management and Board of Directors.

Figure 1. Three Lines of Defence Against Model Risk



Model Development

The model development process is the most important aspect of bank model risk management. As a consequence, supervisory reviews of models focus a great deal of attention on the model development process. A sound model development process helps to ensure that the model is suited to its intended use, that alternative modelling approaches are considered, that model limitations and weaknesses are understood, that the model is appropriately tested, and that management has a plan for continuing model performance monitoring.

Model Purpose

The supervision of model development should begin with the model purpose. Model development documentation should provide a clear statement of the model's purpose(s). The stated model purpose should not be based on circular reasoning, such as "the loan default model's purpose is to predict the likelihood of loan default." The model purpose should address how the bank intends to use the model; for example "the loan default model's purpose is to provide an expected probability of loan default to be used in loan applicant accept/reject decisions".

By stating how the bank intends to use the model, the developers, independent validators, auditors and supervisors can assess whether the modelling approach and accuracy measures are appropriate for the model's purpose. Models can sometimes have more than one purpose. For example, a loan default model might be used for customer credit risk scoring when approving loan applications and also as an input to an expected loan loss model. In the

expected loan loss model, expected losses can be estimated as the product of expected default probability and expected loss given default. In this case, model developers should be able to show that the model is appropriate for all its uses.

Consideration of Alternative Approaches

A common weakness in model development is inadequate consideration of alternative approaches for modelling the process under consideration. For example, for statistical models, there are often various approaches for model estimation that can be applied to a specific situation. The probability of loan default can be estimated using logistic and probit regression, and within each approach, there are various techniques that can be used such as overall versus stepwise model estimation. Model developers might also consider statistical versus machine learning approaches for modelling a process. Consideration of alternative modelling approaches can be aided by a review of industry practices and the academic literature to narrow the number of relevant approaches to be considered. Model developers should be able to provide rationale and support for the chosen approach, whether based on accuracy measures or practical considerations or both, which supervisors are then able to review.

Data Preparation

Data preparation is the most time-consuming and difficult part of model development and is, therefore, an area on which supervisors should focus.

Data preparation begins with selecting a set of candidate model inputs, namely measures for the dependent variable and explanatory variables. The list of candidate variables should be based on a combination of previous experience at the bank, expert judgement and relevant literature. This stage in model development can be aided by the input of subject matter experts, such as bank loan officers, customer relationship managers and bank officers. Bank subject matter experts also play an important role in later steps in model development by reviewing candidate models for the intuitiveness of the results.

Once a set of candidate model inputs is selected, model developers can collect data on these inputs, typically from several different sources, and combine the data in a format that can be used for further analysis and model estimation. This stage of model development is referred to as data extract, transformation and loading (ETL). Model developers will often need to extract data from the bank's internal sources (for example, relational databases), and external sources (for example, credit bureaus), and combine the data in a format that allows for statistical analysis.

Data ETL can be a complicated process and depends on the integrity of the bank's information systems, as well as that of external data sources. Model developers should, therefore, investigate the accuracy of extracted data. This can be done by replicating samples of the data from primary sources and third-party sources where possible. Additional data cleaning tasks include identifying outlier values for model inputs that can distort model estimations. Outliers can be dealt with in several ways, such as deleting outlier observations when they are few in number, capping/flooring outliers (winsorization) and normalization (standardization) of variables.³

³ For example, a common approach to dealing with outliers in the data set is to normalize all inputs by subtracting their mean values from each observed value and divide that difference by the standard deviation of the variable, creating what are also known as Z-scores.

After assembling the model development data set, model developers should conduct a univariate analysis of the relationships between candidate explanatory variables and the dependent variable to identify those candidate model variables with a statistically significant relationship with the dependent variable. Model developers should check for consistency in these relationships across subgroups and over time, and evaluate whether candidate explanatory variables have an intuitively expected and explainable relationship with the dependent variable. Further, model developers can use graphical and statistical analysis of relationships between dependent and explanatory variables to identify potentially useful data transformations that might improve model fit.

After candidate model inputs are selected, model developers should test alternative model specifications in terms of input variables. Since some inputs might be highly correlated with other inputs, model developers should identify redundant variables that can distort subsequent model estimations. Further, model developers should identify segments of the model development data sample that behave similarly. Similarly behaving sample segments, if modelled separately, can improve overall output accuracy. An important part of multivariate model testing is reviewing the model agreement with statistical assumptions used by the chosen estimation technique (for example linear regression), by analysis of appropriate statistical diagnostics.

Model Selection

Once the candidate variables and modelling segments have been identified, model developers should compare the results of model accuracy tests and the intuitiveness of results across alternative model specifications. Use of a variety of model accuracy measures is recommended, as each measure has different strengths and weaknesses (for example, the adjusted R-squared statistic and the Akaike Information Criterion). Model developers should use clearly stated *model selection criteria* when choosing the final model, which may consist of one overall relationship between dependent and explanatory variables or be comprised of a number of modelling segments that similarly relate model inputs. Model selection criteria should align with the model purpose(s). For example, a model's ability to rank-order the riskiness of loan applicants is important for credit scoring models, while a stress test model's ability to find meaningful relationships between economic conditions and the financial position of a bank is important.

Models typically perform best for the model development sample but are often less accurate when applied to out-of-sample data not used in model development. Model developers should test model accuracy for out-of-sample data, particularly data from time periods subsequent to those used in model development (robustness testing).

Out-of-sample tests can only consider the accuracy of the selected modelling approach and of the estimated model. Another test of model accuracy is comparison of the developed model results with those of models that use different input data and/or modelling approaches (benchmark tests). The benchmark models can be developed by the bank (for example earlier models), as well as by outside groups. Finally, model developers should investigate the sensitivity of model results to small and large changes in model input data (sensitivity analysis) to identify possible areas of model risk due to fluctuations in underlying market conditions.

Similarly, a supervisory review of model development should focus on these key aspects of model selection.

Model Monitoring Plan

Model developers should include a model monitoring plan as part of model development. The plan should identify the model performance measures that will be tracked; thresholds for accuracy measures or triggers for remedial actions; types of remedial actions; escalations and steps to be taken should model performance deteriorate significantly; reports and frequency of reports; and the roles and responsibilities for model monitoring.

For example, the performance of a credit applicant risk scoring model might be tracked by measuring the default rates of accepted applicants over different periods of time, such as the first 3 months, second 3 months, and so on after obtaining a loan, with default rates computed by risk score deciles. The bank should expect default rates to increase with risk scores and also with the time since the loan was issued. If the bank finds default rates do not increase with risk scores and time since issued this might trigger remedial actions such as a review of the model and possible updating of the model. Similarly, should the default rates exceed thresholds set by the bank for each risk score decile, a model review might be required. Should the model review and update be insufficient to address model weaknesses, the response might escalate to a model redevelopment or replacement.

Model Documentation

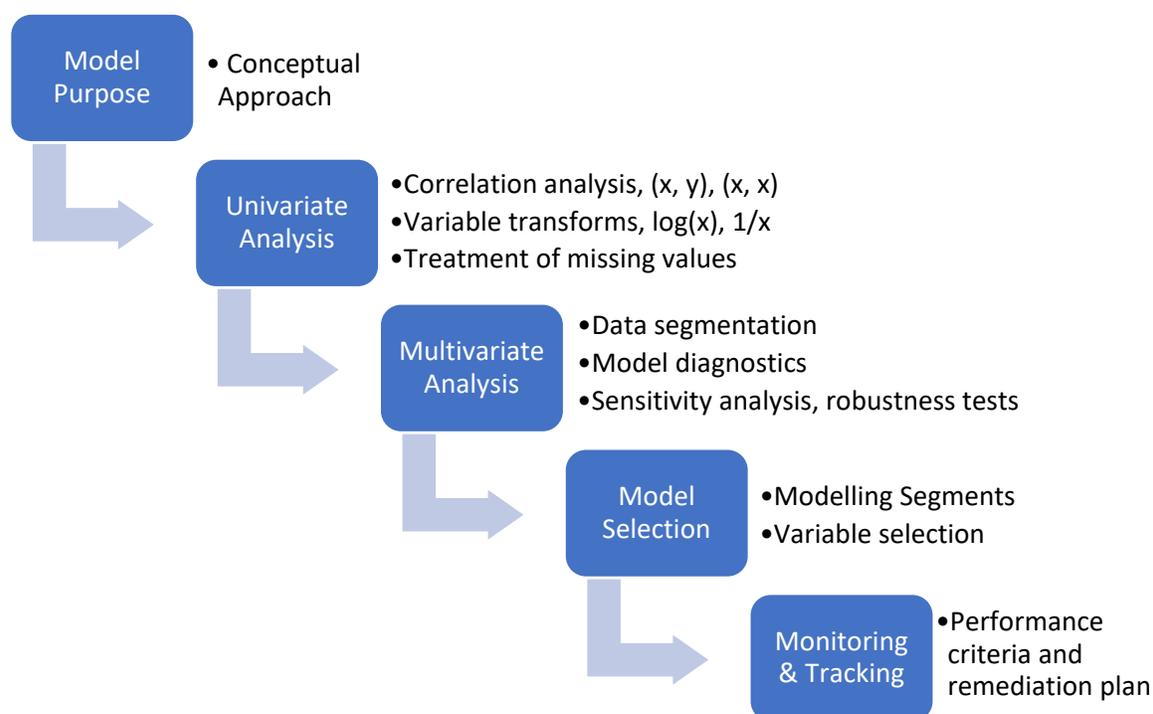
Model development documentation is a critical component of model development and necessary for independent and supervisory model reviews. Importantly, model documentation helps mitigate key person risk by ensuring the storing of information and institutional knowledge that might otherwise be lost with the normal turnover of bank employees. Model documentation should include descriptions of all the model development steps described above, and should be clear and thorough enough such that individuals with appropriate technical skills and knowledge of bank operations could repeat the process.

Model documentation affords model developers and supervisors the opportunity to consider the model development process as a whole, as well as model performance, and to discuss model weakness and risks.

Model Development Process

Figure 2 provides an illustration of the model development process. Should model developers believe certain recommended steps and tests do not apply and/or are not possible or practical, they should include those reasons in model documentation and discussion of their model validation.

Figure 2. Model Development Process



Box 1 illustrates a hypothetical model development process in a bank, and possible findings and recommendations that a supervisor may make regarding the process.

Box 1: Hypothetical Model Supervisory Findings and Recommendations

Hypothetical Model

- The bank recently developed a model to provide credit risk scores for a new line of business - loans to fisheries.
- The loans are for financing the construction of fishing vessels, repair of vessels and construction of aquacultural facilities.
- The bank wishes to estimate the default risk for borrowers and use the estimates as inputs in subsequent models of loan losses and loss reserves.
- The bank has made approximately 200 fisheries loans since lending began three years ago, and nine loan defaults have occurred.
- The bank asked loan officers to identify borrowers and loan attributes that are likely to influence fisheries loan default risk and to assign expected default probabilities to the existing loans.
- The bank then estimated a logit regression model that related assigned default probabilities to borrower and loan attributes that were thought to determine default rates.
- The bank applies the regression model to new fisheries loans to predict their expected default probabilities.

Supervisory Findings

- The fisheries loan default model uses the judgement of subject matter experts in selecting candidate model explanatory variables.
- The model development data are severely flawed because the default model dependent variable is not based on actual defaults.
- The candidate explanatory variables cannot be accurately tested for statistical significance and their impact on loan defaults because the dependent variable is an expert opinion and not a fact.

Supervisory Recommendations

- The bank lacks the data necessary to estimate a statistical model of fisheries loan default risk.
- The bank's assignments of loan default probabilities, however, use an accepted modelling approach - qualitative modelling.
- The bank should document how it developed the qualitative estimates of fisheries loan default risk.
- Independent model validation should then review the qualitative model for approval of use.
- The bank should explore the possibility of using data from peer banks on loan defaults and losses from their fisheries loan portfolios for statistical model development.
- The bank should collect information on defaults and losses on its fisheries loan portfolio for possible use in statistical loan default models, once sufficient data are available.

Model Validation

The model validation performed by model developers is the first line of defence against model risk. Model developers should conduct their own validation of the model to ensure that appropriate steps were taken in model development, as determined by the bank's policies and supervisory guidance. The validation should tie model development steps to recommended procedures and supervisory guidance.

In addition to evaluating the work of model developers, supervisors should review independent model validations, as documented in model validation reports and through meetings with independent model validators.

Supervisors should keep in mind the distinction between laws, regulations, and guidance. Model risk management guidance is comprised of a set of recommended procedures with supporting rationale and does not have the force of law. Supervisors should weigh whether departures from guidance increase model risk and threaten bank safety and soundness. Compliance models, however, should follow the applicable laws and regulations.

Independent Model Validation

While model developers work to ensure models are suited for their intended use, limitations in model development data and modelling techniques, and a failure to follow sound model development practices, can increase model risk. The role of independent model validation is to provide a second line of defence against model risk. Specifically, the role of independent model validation is to provide *effective challenges* to model developers.

Independent model validators should not be part of the model development team and should not be part of the management team that owns and uses models. Independence from model

development and use, combined with sufficient stature within the organization, are necessary conditions to allow independent validators to challenge modellers and management effectively. Independent model developers typically work in a unit that is separate from other bank operations and report to senior management, such as the chief risk officer, to avoid conflicts of interest with other bank operations.

Independent model validations should include thorough reviews of each of the model development steps discussed above, as well as model development documentation to ensure compliance with bank policies and supervisory guidance. Independent model validations should not rely exclusively on model documentation, however. Independent replication of select model input data, statistical tests and model estimations can serve as a check against errors in the first line of defence. Model validators should also consider alternative model inputs and modelling approaches as a test of model developers' choices.

The frequency, depth and scope of independent model validations should be determined by model complexity, the risk exposures being modelled, the degree to which a bank relies on the model, and the potential costs of associated model risk. In many countries, independent validations of newly developed models are expected to be the most thorough review of the model. If models are periodically re-estimated to incorporate new data or re-developed to consider new variables and approaches (*model change*) the depth and scope of the independent validation should adjust to the level of change being considered.

Independent model validators typically assign a model risk rating to models that are used when assigning the frequency of periodic model re-validations in the absence of model change. Model risk rankings are assigned based on the factors a bank's independent model validation unit determines to be most appropriate for models. Commonly used factors are a model's complexity and the potential costs of model risk. Model complexity is considered to be important because of its association with model uncertainty. The potential for unseen, adverse consequences increases with the number of assumptions used in model development, the number of approximations used in data preparation, and the degree of complexity in model estimation. Model risk also increases with the extent to which a bank relies on a model, the risk exposures managed by the model, and the resulting potential costs of model risk.

Table 1 illustrates how a bank might assign a model risk rating. The risk weighting of a model's potential financial impact and complexity should be determined by the independent model validation unit based on their experience at the bank.

Table 1. Model Risk Ratings: Illustrative Example

Model Complexity	Financial Impact		
	High	Moderate	Low
High	10	5	2
Moderate	8	4	2
Low	6	3	1

The end result of model validation is the independent validation report which documents the validation process and support for the independent validators' findings, and the independent validation unit's approval or disapproval of model use. In most instances, independent model validators recommend a model for use, but request model developers to address concerns and questions raised by the validators. In some cases, a model's continued use is dependent on the

bank successfully addressing independent model validation findings within a specified time frame.

Vendor Models

Banks can employ models developed externally by vendors who have expertise and access to data and techniques not readily available to the bank. *Vendor models* can be an efficient and effective modelling solution for banks. For example, models developed by vendors for the detection of money laundering and terrorist financing can make use of data shared by many banks with the vendor to augment model development data and model capabilities.

In these situations, the ability of independent model validators at the bank to conduct validation activities is limited. Vendors typically do not share their model development data, modelling techniques and model development processes with banks and consider this information to be proprietary. Supervisors should still expect banks to apply some model validation to vendor models, including the use of vendor reports on model validations, sensitivity analysis of vendor models to check that variables affect results as expected, and internal tracking of model performance.

Independence

Independent model validators may become involved in the model development process. Such involvement can stem from many aspects of independent model validation. To begin, banks might ask the independent model validation unit (IVU) to review model development work on a periodic basis so that model developers can avoid wasting time and effort by pursuing incorrect avenues of model development. In this role, the IVU can and should voice concerns about potential pitfalls in the model development process.

However, the IVU should be careful not to become prescriptive in its recommendations so it can maintain its independence from model development. For example, if the IVU is asked to review proposed models, in terms of both inputs and modelling approach, IVU recommendations that modelers consider additional inputs and/or modelling approaches does not threaten IVU independence. But if the IVU makes specific recommendations that prescribe how the IVU wishes model development work to proceed, then the IVU has become part of the model development team and its independence is compromised. Once the IVU becomes prescriptive, it is effectively reviewing its own judgments and it becomes part of the model development team, with the result that there is no second line of defence against model risk.

This same concern applies to supervisors. If supervisors are too prescriptive in their recommendations regarding a model, they too become part of model development and lose independence from model development.

A similar concern about the independence of the IVU arises when the IVU is pressured by senior bank management to weigh regulatory deadlines for new models, model development costs and other factors that, while important to the bank, lie outside the scope of independent model validation. In these situations, the appropriate response to competing pressures for model approvals is for the bank to discuss its competing demands and limited resources with its supervisor to work out a timetable for meeting supervisory demands.

Internal Audit

Internal audit has an important role to play in mitigating model risk and is the third line of defence against model risk. Internal audits are conducted to check that a bank's model governance policies, practices and procedures are followed by model developers, model owners, and independent model validators. Model audits can also evaluate whether banks are following supervisory guidance on model risk management and comply with relevant banking laws and regulations.

Audit findings should be reported to senior management and the bank's Board of Directors. Model audit reports should collectively discuss the entirety of the model life cycle, from initial development, independent validation and periodic re-validation, continuing model use and monitoring, compliance with banking laws and regulations, and the eventual retirement of models. As is the case with IVU model validation findings, audit findings should be reported to model developers and bank management, with recommended remediations and time frames for remediation.

Model Governance

In addition to examining the model development process, **supervisors should examine banks' governance of model risk.** Model governance reviews should cover how a bank's policies, practices and procedures for model development, documentation, independent validation, use, and monitoring are established and overseen by the bank's senior management and Board of Directors.

A bank's senior management and Board of Directors are responsible for developing a bank-wide model risk management program. The model risk management program should be comprehensive and fit within the bank's overall risk management program. The model risk management program should establish standards for model development, independent validation, model use and monitoring; and it should incorporate supervisory expectations for model risk management and specify the roles and responsibilities of the three lines of defence against model risk. While the Board of Directors is ultimately responsible for the program, the Board typically delegates to senior management the responsibility to execute and maintain the model risk management program.

Machine Learning Models

Machine learning (ML) algorithms, such as artificial neural networks, gradient boosting, support vector machines and random forests, are important modelling approaches for banks. ML algorithms have been available for many years. The first artificial neural network was developed in 1943, followed by support vector machines (1963), random forests (1995) and gradient boosting (1999). These algorithms have been improved over time, such as the introduction of error correcting backpropagation to artificial neural networks in 1986. As a result, there are variations of each category of ML algorithms that are useful for different types of data and objectives.

The Advent of ML

Over the past 25 years, there have been significant advances in computer capabilities, computer languages and model development software. These advances have combined with increases in the availability of data that in some instances only ML models can fully utilize,

thereby making ML the preferred modelling approach in many situations. ML models offer increased flexibility over traditional statistical models in several ways:

- ML algorithms can model linear and non-linear relationships between dependent and explanatory variables, while statistical models can only model linear relationships.
- ML models can use non-traditional data as model inputs. For example, convolutional neural networks can be trained to recognize handwriting.
- The model development process can be faster and less costly for ML models than statistical models for a number of reasons:
 - ML models do not make assumptions about the probability distributions of model input data; hence, modellers do not need to check ML models' conformance with the assumptions made by statistical models, for example, assumptions of best linear unbiased estimator regression models.
 - Determining optimal model data segments (similarly behaving groups of entities) can improve statistical model performance but is unnecessary in many types of ML models since data segmentation is subsumed by the model (for example, random forest decision trees).
 - Controls for model overfitting can be incorporated in ML algorithm objective functions through regularization terms that penalize model complexity but reward model accuracy. Regression models do not include overfitting controls and rely on post-estimation tests for overfitting.

Supervision of ML Models

There has been some discussion about the challenges of applying supervisory guidance on model risk management to the independent validation of ML models, as well as the challenges for bank supervisors.⁴ For example, US bank supervisory guidance on model risk management focuses on statistical models because ML was not common practice at banks when the guidance was published in 2011.⁵ US bank supervisory guidance on model risk management has shown, however, that it is broad enough in scope to be applied to ML models, and US bank supervisors have been examining ML models at banks for a number of years and as yet have not changed model risk management guidance.⁶ That may be about to change, however.

In March 2021 the US bank regulatory agencies issued a request for comment on bank use of artificial intelligence⁷ (AI), including ML, for business purposes.⁸ The agencies list a number of uses of AI by banks, and point out the potential benefits of AI (for example, greater efficiency and accuracy, expansion of credit access) and potential costs (for example, reduced transparency of the model, increased risk of violation of consumer protection laws, challenges resulting from dynamic model updating). The request for comments focuses on the challenges in developing, using and managing AI.

The US agencies will use the responses “to assist in determining whether any clarifications from the agencies would be helpful for financial institutions’ use of AI in a safe and sound manner and in compliance with applicable laws and regulations, including those related to consumer

⁴ See, for example, Protiviti (2019) and Model Risk Managers’ International Association (2021).

⁵ Board of Governors et al. (2011).

⁶ There are no published studies on bank preferences for statistical versus ML models of which the author is aware. Hence, in this section of the paper the author is relying on his experience in over 60 examinations of bank models at 25 large and complex US commercial banks between 2012 and 2020.

⁷ AI refers to the broad category of models designed to simulate human thinking and behavior, while ML is a subset of AI in which the model learns from data without being programmed.

⁸ Board of Governors et al. (2021).

protection.” It appears, therefore, that the agencies believe the existing supervisory guidance on model risk management applies to AI (and ML) and that a wholesale rewriting of supervisory guidance on model risk management does not appear to be necessary at this time.

A likely reason for the continued usefulness of US supervisory guidance on model risk management is that the US bank supervisors’ approach to examining bank business models is risk-based and not prescriptive. Specifically, the focus of the US examinations of bank business models are on the model development process, its appropriateness for the model’s purpose, transparency, conformance with low-risk model development practices, and the bank’s governance of model risk. However, there is a strong element of prescriptive supervision of models used to meet compliance requirements for areas such as capital adequacy and stress testing.

Another likely reason for the lack of new US supervisory guidance on model risk management is that the approaches used by ML models are too complex and diverse to be addressed in the same technical manner as is used for statistical models in existing guidance. Moreover, advances in ML result in new approaches and expand existing approaches at a pace that is likely to make model-specific supervisory guidance quickly outdated.

New algorithms continue to be developed at a rapid pace, expanding the field of artificial intelligence into the area of dynamic model re-development as new data become available.⁹

These advances in models pose serious challenges to the supervision of bank models.

Traditional supervisory model reviews are based on information about the model as it is being implemented at the time of the review. This raises the question of how to examine a model-based decision process that is continually adapting to new information. There does not currently exist a framework for the supervision of dynamic modelling applications.

Bank supervisors are likely to permit bank use of algorithms that change the underlying model in real-time and will adapt the model review process accordingly. For example, models that can incorporate new explanatory variables through automated processes can pose the risk of possible violations of fair lending and consumer protection laws by adding variables (for example, gender or race) that result in disparate impacts on some customer segments. This risk can be mitigated by pre-screening candidate variables available to the model, controls on the weight (importance) the model can place on individual variables, and periodic monitoring of model output for compliance with consumer protection laws.

Advances in modelling approaches have the potential to expand access to credit to underserved market segments by utilizing alternative information on creditworthiness to risk rank credit applicants more accurately. Some market segments, for example, first-time borrowers, lack information used by traditional credit scoring models, such as credit bureau scores, employment and credit histories. ML models have the capacity to risk score such credit applicants by using non-traditional model inputs as alternative information, for example, school grade-point averages, educational degrees, and occupation. Supervision of model risk can similarly advance by using alternative approaches to reviewing models.

⁹ Dynamic model updating is an automated process in which the software can update/redevelop a model as new data and information become available and almost immediately put the new model into production.

Conclusion

This Toronto Centre Note has discussed the supervision of model risk management. It has stressed the importance of supervisors:

- Tailoring their supervision to the purpose of the models used by a bank (or indeed by any supervised financial institution).
- Taking a risk-based and proportional approach, based on the materiality of a bank's models.
- Reviewing all the stages of the model life cycle, namely model development, independent validation, implementation, monitoring, change and eventual retirement.
- Reviewing the governance framework that determines how bank decisions that concern models are made, and the model risk management program that implements the governance framework for day-to-day model risk management activities.
- Where applicable, taking account of a bank's use of machine learning models.

References

Basel Committee on Banking Supervision. [Core Principles for Effective Banking Supervision](#). September 2012.

Basel Committee on Banking Supervision. [Consultative Document: Principles for home-host supervisory cooperation and allocation mechanisms in the context of Advanced Measurement Approaches](#). February 2007.

Board of Governors of the Federal Reserve System, Bureau of Consumer Financial Protection, Department of the Treasury, Federal Deposit Insurance Corporation, National Credit Union Administration, and Office of the Comptroller of the Currency. [Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, Including Machine Learning](#), Federal Register Vol. 86, No. 60, March 31, 2021, pp. 16837-16842.

Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency. SR Letter 11-7, [Supervisory Guidance on Model Risk Management](#). April 4, 2011.

European Central Bank. [ECB Guide to Internal Models](#). October 2019.

International Monetary Fund (IMF). [Toward a Framework for Financial Stability. World Economic and Financial Surveys](#). January 1998.

Model Risk Managers' International Association, [Machine Learning and Model Risk Management](#), Technical Report 2021-01, Version 1.0. March 8, 2021.

Protiviti. [Validation of Machine Learning Models: Challenges and Alternatives](#). 2019.