

Guidance on Advanced Approaches GAA 2014-01: Supervisory Guidance for Data, Modeling, and Model Risk Management Under the Operational Risk Advanced Measurement Approaches

This document provides guidance regarding supervisory expectations for data, modeling, and model risk management under the operational risk advanced measurement approaches (AMA) to calculate a regulated banking organization's operational risk. Staff at the Office of the Comptroller of the Currency (OCC) and the Board of Governors of the Federal Reserve System (Board) worked closely to develop this guidance.

I. Background

Under the advanced approaches risk-based capital rule (rule), a banking organization must use the advanced measurement approaches to calculate its capital requirement for operational risk.^{1, 2} The starting point for this calculation is the banking organization's estimated operational risk exposure, which is the 99.9th percentile of the distribution of the banking organization's potential aggregate operational losses over a one-year horizon (not incorporating eligible operational risk offsets or qualifying operational risk mitigants).³ The AMA requires a banking organization to estimate its operational risk exposure by collecting and using four data elements: internal operational loss event data (internal data), external operational loss event data (external data), scenario analysis, and business environment and internal control factors (BEICF).⁴

The AMA does not prescribe any specific approach for operational risk quantification systems, and provides the flexibility for banking organizations to use the four data elements in the most effective way when estimating operational risk exposure on a forward-looking basis. All four elements are critical, required components of a banking organization's operational risk quantification process. Internal operational loss event data often indicate a banking organization's historical operational risk exposure and can provide a foundation for the forward-

¹ The rule is at 12 CFR part 3 for the OCC and 12 CFR part 217 for the Board.

² For simplicity, and unless otherwise indicated, the rule and this guidance use the term "banking organization" to include national banks, federal savings associations, state member banks, and bank holding companies. Beginning January 1, 2015, this guidance will also apply to any savings and loan holding company subject to 12 CFR 217, subpart E.

³ The definition of operational risk exposure is at 12 CFR 3.101, OCC, and 12 CFR 217.101, Board.

⁴ Under the rule, internal operational loss event data are gross operational loss amounts, dates, recoveries, and relevant causal information for operational loss events occurring at the banking organization. External operational loss event data are gross operational loss amounts, dates, recoveries, and relevant causal information for operational loss events occurring at organizations other than the banking organization. Scenario analysis is a systematic process for obtaining expert opinions from business managers and risk management experts to derive reasoned assessments of the likelihood and loss impact of plausible high-severity operational losses. BEICFs are indicators of a banking organization's operational risk profile that reflect a current and forward-looking assessment of the banking organization's underlying business risk factors and internal control environment.

looking estimation of operational risk exposure. Depending, however, on a banking organization's specific circumstances (e.g., limited internal data or a significant change in a banking organization's business mix), it may be appropriate to increase the weight given to scenario analysis, BEICFs, or external data for a more informed, forward-looking estimate of risk exposure.

Although not required, many banking organizations use the loss distribution approach (LDA) as a core modeling technique in their AMA quantification processes.⁵ Building on "Interagency Guidance on the Advanced Measurement Approaches for Operational Risk," issued in June 2011 (June 2011 guidance), the following sections provide additional guidance on the quantification of operational risk exposure, including frequently encountered issues relating to data, units of measure (UOM), model selection and fitting, diversification, and model risk management.⁶ While many of these concepts are broadly applicable to estimating operational risk exposure, they are particularly relevant when applying the LDA.

II. General Implementation Guidance

A. Data

The credibility of any empirical modeling approach hinges on the relevance, integrity (e.g., accuracy, comprehensiveness, and appropriate classification), and internal consistency of the underlying data. Examiners therefore give close attention to data issues that can affect the credibility of a banking organization's estimate of operational risk exposure.

1. Data Selection

As a general matter, banking organizations should incorporate all relevant data within their frequency, severity, and diversification models. While the rule stipulates a minimum observation period of five years of internal data,⁷ the data window used for estimation purposes (reference period) should, wherever feasible, encompass a longer horizon for which data are available and relevant. A greater number of observations generally enables more accurate estimation of model parameters. Longer reference periods also support estimates that are consistent with a broader range of economic and business

⁵ The LDA is an empirical modeling technique that can be used to estimate value-at-risk measures for annual operational risk losses based on fitted parametric distributions. Using a banking organization's own internal data, at each UOM, the LDA involves estimating probability distributions for the frequency and the severity of operational loss events. The estimated frequency and severity distribution are combined using, for example, Monte Carlo simulation techniques to estimate the probability distribution for annual operational risk losses at each UOM. An estimated probability distribution for overall annual operational risk losses is then calculated by combining the stand-alone distributions for the various UOMs within a diversification model. The LDA-based estimate of a banking organization's overall operational risk exposure is computed as the 99.9th percentile from this estimated distribution.

⁶ See OCC Bulletin 2011-21, "Interagency Guidance on the Advanced Measurement Approaches for Operational Risk," and Board SR letter 11-8, "Supervisory Guidance on Implementation Issues Related to the Advanced Measurement Approaches for Operational Risk."

⁷ See 12 CFR 3.122(g)(2)(ii)(A)(I), OCC, and 12 CFR 217.122(g)(2)(ii)(A)(I), Board.

conditions and, most importantly, increase the likelihood of capturing unusually large loss events. Indeed, because the AMA focuses on the potential for experiencing very large annual operational losses, extreme individual loss events typically are the most relevant data for operational risk modeling and have the greatest impact on estimation results, particularly for the LDA.

A banking organization should thoroughly justify and document any exclusion of internal data from its operational risk quantification process. An AMA that excludes large historical losses would call into question the conceptual soundness of a banking organization's approach. In particular, examiners will scrutinize exclusions of large historical loss events that could meaningfully affect the banking organization's estimated overall operational risk exposure. Because large operational risk loss events generally occur infrequently and are, by their nature, unexpected, such observations are often the most informative data from an operational risk modeling perspective. Even historical losses associated with a divested business may be informative of future potential losses, because many elements of a banking organization's control infrastructure and aspects of its loss exposure are endemic to the firm as a whole (e.g., the risk control culture established by senior management, the firm's overall risk appetite, or legal risks associated with certain broadly defined businesses). Thus, past losses associated with a now-discontinued business unit (e.g., representation and warranty losses associated with sales or securitizations of subprime mortgages, which are no longer originated by the banking organization) may indicate potential risks associated with business practices or products more generally.

The exclusion of *any* internal data should be rare and accompanied by strong justification, including BEICF and empirical analysis, especially for large loss events. The exclusion or down-weighting of historical data should be supported by analysis demonstrating that such exclusions or down-weighting results in a more credible, transparent, systematic, and verifiable forward-looking measure of the banking organization's overall operational risk exposure.⁸ In any event, for baseline reference purposes, banking organizations are expected to provide examiners with model-based exposure estimates using all internal data (before exclusions).

2. Loss Amount

A banking organization may estimate its operational risk exposure based on either the gross loss amounts or losses net of recoveries. Because banking organizations must be able to estimate operational risk exposure with and without insurance benefits and other operational risk mitigants, a banking organization must not incorporate insurance benefits or other mitigants when determining the amount of losses net of recoveries.⁹ The process

⁸ For example, if a banking organization attempts to justify the exclusion of historical losses associated with a divested business, supervisors expect the organization's analysis to demonstrate that after the divestiture there is no residual exposure, the excluded loss experience has no relevance to other continuing activities or products, the divestiture has reduced the banking organization's overall operational risk exposure, and the divestiture will not be offset over time by increased risk-taking in other related businesses.

⁹ See 12 CFR 3.161(a)(1), OCC, and 12 CFR 217.161(a)(1), Board.

for selecting the loss amount used in AMA quantification should be internally consistent, well-reasoned, clearly documented, and understood by the banking organization personnel responsible for its implementation. Recoveries can occur over an extended time, and a banking organization should establish credible guidelines for recovery periods and the appropriate discount or other treatment of recoveries that extend beyond the one-year time horizon.

3. Selection of Loss Event Dates

Banking organizations generally record at least three dates associated with each operational loss event: occurrence date, discovery date, and accounting date. For most loss events, these dates are the same, or very similar, but for some types of loss events—particularly legal loss events—there can be substantial timing lags. When data exhibit long lags between the dates of occurrence and discovery, a banking organization should be able to explain the reasons for such lags and, especially for non-legal-related events, should consider whether improvements are needed in the organization's internal data collection systems.

The selection of loss event dates is critical when modeling the frequency of operational loss events and diversification effects among UOMs. Thus, the process for selecting loss event dates used in AMA quantification should be credible, transparent, systematic, and verifiable, and understood by the banking organization personnel responsible for implementing the process.

4. Legal Losses Before Settlement

As noted in the June 2011 guidance, there are cases in which a banking organization, after consulting with appropriate legal counsel, determines that it must record a legal reserve for accounting purposes with respect to a pending or potential claim. In these cases, the banking organization should categorize the legal reserve as an operational risk loss event for regulatory capital purposes and include the legal reserve in its operational risk quantification process. The date that a loss event is recorded for regulatory capital purposes should be consistent with, and no later than, the date a legal reserve is established. The final operational loss amount for that event should be consistent with the amounts to date of the associated legal reserve, as adjusted to reflect any settlement or final judgment, internal legal expenses, external legal fees, and other costs driven by the legal claim.

5. Comparability of Loss Severity Data

The LDA assumes that, for each UOM, the probability distribution of loss severities is unchanging over time. In practice, some factors that may cause this probability distribution to vary over time are challenging to measure (for example, scale of operation). Other factors are easier to incorporate, such as the general price level. Where feasible, banking organizations should adjust historical severity data so that, to a reasonable approximation, the empirical probability distribution of loss severities is

invariant over time. Also, loss amounts denominated in a foreign currency should be converted into U.S. dollars at the exchange rate applicable when the loss occurred.

6. Aggregation of Losses Associated With a Common Trigger or Causal Factor

An important objective in quantifying operational risk exposure is capturing the full extent to which individual operational losses are correlated, especially losses resulting from the same fundamental source. Quantification processes that do not accurately reflect these correlations may misrepresent a banking organization's operational risk exposure.

A fundamental assumption of the LDA is that loss events within each UOM are independent and identically distributed. Sometimes, however, individual losses have a common underlying trigger or instigating factor or a clear relationship to each other. In such situations, a generally acceptable approach is to aggregate losses having a common trigger or instigating factor, or a clear relationship to each other, and treat these related losses as a single event. This aggregation principle would apply regardless of whether such losses are spread over time or across business lines or transactions. Alternatively, if losses that are closely related or positively correlated are not aggregated in this manner, supervisors expect a banking organization to explicitly account for the dependence among individual loss events within each UOM. The rule requires that a banking organization must demonstrate to the satisfaction of its primary supervisor that its process for estimating dependence is sound, robust to a variety of scenarios, and implemented with integrity and that it allows for uncertainty surrounding the estimates.¹⁰

A banking organization that uses the LDA should have a clear, well-documented policy for addressing losses that are closely related or positively correlated, including procedures for applying the aggregation principle and criteria for determining when multiple losses should be aggregated and treated as a single event. This policy should establish clear guidelines for deciding the circumstances, types of data, and methodology for aggregating losses as appropriate for its business, risk management, and operational risk exposure modeling. In addition, processes should be in place to ensure that there

- is a firm-wide understanding of the data aggregation policy.
- is appropriate sharing of loss event data and information across businesses to implement the policy effectively.
- are adequate controls (including independent review) to assess ongoing compliance with the policies.

The aggregation of operational losses may also be relevant for approaches to estimating operational risk exposure other than the LDA. As such, a banking organization should implement these policies and processes when relevant and appropriate.

¹⁰ See 12 CFR 3.122(g)(3)(i)(D), OCC, and 12 CFR 217.122(g)(3)(i)(D), Board.

7. Treatment of Internal Loss Data in Mergers or Acquisitions

Under the rule, *all* expenses associated with the same operational risk loss event must be aggregated and treated as a single loss for internal data and modeling purposes, except for opportunity costs, forgone revenue, and costs related to risk management and control enhancements to prevent future operational losses.¹¹ This requirement is applicable even when such expenses are spread over time or over multiple business units. A frequently encountered practical issue is how to apply this requirement in the context of mergers and acquisitions, which often involve combining firms with different corporate cultures, internal control environments, and risk management practices.

The combination and use of internal loss data following a merger or acquisition poses special challenges, particularly when internal loss data is weighted heavily in the banking organization's estimate of its operational risk exposure. In most circumstances, a reasonable approach involves combining the operational loss histories of the underlying firms and treating the resultant loss history as if it had occurred at a single entity. This approach involves linking and aggregating losses into a single loss event for modeling and risk management purposes when there is a common trigger or instigating factor (including losses occurring before the merger or acquisition).¹²

For LDA modeling purposes, examiners generally find the aggregation of historical data from the combined banking organizations to be reasonable until the consolidated firm's actual internal loss experience demonstrates otherwise. This treatment is a practical alternative to explicitly modeling dependencies among loss events across the consolidated firm. A banking organization may, however, establish a reasonable threshold below which loss data would not be aggregated. The threshold should be credible, well documented, and supported by sensitivity analysis.

In some circumstances, aggregating the historical data for the underlying entities could misrepresent a banking organization's forward-looking risk exposure. For example, a banking organization may acquire a firm with a manifestly different internal control environment as evidenced by a materially greater loss experience, such as the acquisition of a failed or distressed banking organization. In such instances, a case for excluding or down-weighting historical loss data of the distressed acquired firm within the LDA modeling process may be supportable.

The decision to exclude or down-weight data from an acquired firm should be accompanied by strong justification, including BEICF and empirical analysis. The analysis should demonstrate that the exclusion or down-weighting of loss experience from the acquired firm results in a credible, transparent, systematic, and verifiable

¹¹ See definition of operational loss at 12 CFR 3.101, OCC, and 12 CFR 217.101, Board.

¹² Consistent with the concept of a threshold for data collection as discussed in the June 2011 guidance, a banking organization may set a reasonable threshold for aggregating data following a merger or acquisition. The threshold should not, however, exclude important loss event data, and should permit the banking organization to capture a substantial dollar value of the combined operational losses.

forward-looking measure of the consolidated firm's overall operational risk exposure. Examiners will review closely the assumptions and application of expert judgment or forward-looking considerations, such as BEICFs, when assessing the overall credibility of the decision to exclude or down-weight loss data from the acquired firm. Examiners will also review model-based exposure estimates that do not exclude or down-weight data to understand the impact of these decisions on AMA capital estimation.

8. Role of BEICFs When Estimating Operational Risk Exposure

BEICFs are indicators designed to provide a forward-looking assessment of a banking organization's business risk factors and internal control environment. In principle, BEICF-based analyses can be used to reflect forward-looking changes in the risk profile—such as the impact of discontinuing a line of business, a change in the internal control environment, or developments in the external environment—that may not be directly observable from models based solely on historical data. BEICF-based analysis might be used *ex post* to adjust the estimated operational risk exposure upward or downward, or within the modeling process itself, such as by weighting certain historical data more or less heavily. The approach for applying BEICFs should be credible, transparent, systematic, and verifiable, both conceptually and empirically.

A banking organization should have a clear, well-documented policy and process for using BEICFs. The policy should identify the BEICFs to be considered and their specific applications within the AMA quantification process, including applicable quantitative limits or other restrictions on their use. The policy should also establish rigorous validation mechanisms around the use of BEICFs, including governance processes over, robust challenge to, and independent review of related methodologies, processes, and controls.

As required by the rule, the banking organization must periodically compare its actual operational loss experience against prior BEICF assessments.¹³ A banking organization is required to have credible, transparent, systematic, verifiable BEICFs that are appropriately weighted and used in the capital estimation process.

B. Units of Measure

Under the rule, a UOM must not combine business activities or operational loss events with demonstrably different risk profiles.¹⁴ Granular UOM segmentation is encouraged when available data are sufficient to allow for accurate and stable estimates of operational risk exposure. When evaluating homogeneity within a UOM, a banking organization should employ sound statistical procedures and other appropriate considerations, such as risk management and business line requirements. In addition, the basis for choosing a particular UOM segmentation should be well justified and documented, particularly when the approach departs from the loss event types described in the rule or industry standard business line

¹³ See 12 CFR 3.122(g)(2)(ii)(D), OCC, and 12 CFR 217.122(g)(2)(ii)(D), Board.

¹⁴ See 12 CFR 3.122(g)(3)(i)(B), OCC, and 12 CFR 217.122(g)(3)(i)(B), Board.

classifications. Use of an unorthodox segmentation should be demonstrably credible, and a banking organization should quantify and document the impact of this decision on its operational risk exposure estimate.

When internal data are insufficient to credibly support segmentation by the loss event types described in the rule and by industry standard business line classifications, the banking organization should document this condition. To compensate for sparse data, it may be appropriate for the banking organization to consolidate UOMs, provided that a credible business case can be made for the appropriateness of the consolidation. When UOMs have been consolidated because of sparse data, the banking organization should, periodically on an ongoing basis, evaluate whether sufficient data have been accumulated to achieve a desirable level of granular UOM segmentation.

A banking organization may explore mixing external data with internal data at the UOM level. The June 2011 guidance indicates that, when using both external and internal data, external data typically are modeled separately and then combined with the results of an internal data model; the direct mixing of external and internal data in the same model, however, may be acceptable in cases where insufficient data preclude the reliable estimation and validation of separate models.

Before mixing external and internal data in the same model, a banking organization should investigate, support with statistical analysis, and document that these data sources are reasonably consistent with one another. In some cases, the goal of homogeneity within the UOM may motivate a banking organization to explore scaling of external losses so that the losses are consistent with the banking organization's risk profile. To ensure that the approach meets the rule's requirement that the data elements be credibly weighted, direct (weighted) mixing of external and internal data or (weighted) averaging of external-data-driven and internal-data-driven UOM loss models that lead to reductions of operational risk exposure estimates vis-à-vis internal-data-only results will be subject to heightened supervisory scrutiny.¹⁵

C. Model Selection and Estimation

The modeling of operational risk loss frequency and loss severity pose distinct challenges. This section discusses selected modeling issues that commonly arise in each area.

1. Frequency Modeling

Model selection

Banking organizations using the LDA generally employ either the Poisson or negative binomial distribution to model frequency, with the majority using the Poisson distribution. The Poisson distribution is a simple, one-parameter probability distribution, with equal mean and variance. The negative binomial is a two-parameter distribution, where the mean and the variance are allowed to differ. Where there is sufficient

¹⁵ See 12 CFR 3.122(g)(3)(i)(C), OCC, and 12 CFR 217.122(g)(3)(i)(C), Board.

operational loss data, a banking organization should compare the sample mean and sample variance when choosing between these two distributions. If frequency data exhibit trends or discontinuities, and the use of a shorter reference data period does not solve these issues or leads to poor statistical accuracy, banking organizations should explore other modeling methodologies, such as factor models.¹⁶

Accounting for Losses Below the Collection and Modeling Thresholds

As noted earlier, a banking organization may refrain from collecting internal data for individual operational losses below internal threshold amounts if the banking organization can demonstrate that the thresholds are reasonable, do not exclude important internal data, and permit the banking organization to capture substantially all the dollar value of its operational losses. When the data collection threshold is above zero, the observed raw loss frequency for a UOM generally does not fully account for all loss events experiences, and frequency estimates based on such data could be biased downward. In addition, a banking organization may model its operational risk exposure using a threshold that is higher than the data collection threshold; however, using the raw loss frequency above, this modeling threshold as the total frequency estimate for the UOM would also lead to a downwardly biased estimate.

To address this potential bias, banking organizations generally have adopted one of the following approaches:

- If aggregate losses below the data collection threshold are known, a banking organization may perform a top-side adjustment to the UOM's estimated operational risk exposure. This is calculated by first estimating the operational risk exposure based on the observed historical losses above the threshold, then adjusting this initial estimate upward by the amount of annual average aggregate operational losses under the threshold. Under this approach, the frequency above the threshold would be used with no adjustment.
- If aggregate losses under the threshold are not known, standard statistical techniques are available for adjusting estimated model parameters to deal with truncated data. In the present context, for a given UOM, such a technique could be used to estimate the probability of incurring a loss below the modeling threshold.¹⁷ This probability could then be used to appropriately scale up the frequency parameter(s) for the UOM.

¹⁶ Some banking organizations are experimenting with factor models for frequency. The factors typically consist of firm-specific variables (e.g., revenue, noninterest income, total assets, or number of employees), macroeconomic variables, or some combination thereof. When assessing the appropriateness of a factor model, a banking organization should pay close attention to basic regression diagnostics, such as the statistical significance of the explanatory variables, the overall model fit statistics, and the forecast accuracy. In addition, a banking organization should address the possibility of model misspecification stemming from spurious correlations and should consider the conceptual soundness of the approach, including the economic rationale for the explanatory factors used in the model.

¹⁷ See, for example, Dempster, A.P.; Laird, N.M.; and Rubin, D.B., "Maximum Likelihood From Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society, Series B (Methodological)* 39 (1977): 1-38.

2. Severity Modeling

Model Selection

Compared with frequency modeling, severity modeling is more complex. Severity model selection involves more difficult goodness-of-fit and over-fitting considerations, and banking organizations are expected to establish sound and well-defined modeling criteria to account for these considerations. Furthermore, banking organizations are expected to understand and document the sensitivity of loss exposure estimates to different modeling choices.

A critical aspect of severity modeling is the assumed family of probability distributions characterizing loss severity within each UOM. Industry data support modeling the probability distribution of loss severity as a fat-tailed distribution, which includes members of the sub-exponential family of distributions (e.g., log-normal, generalized Pareto, Burr, and log-gamma).¹⁸ In addition, supervisory experience indicates that the statistical characteristics of loss severity typically differ dramatically across UOMs, necessitating different distributional assumptions across UOMs to ensure accurate fits. Thus, banking organizations should test a menu of alternative severity models for each UOM and base their model selection choices on sound and well-defined goodness-of-fit and over-fitting criteria.

Goodness-of-fit criteria: Measures of goodness-of-fit quantify the degree to which sample data are consistent with an assumed or estimated probability distribution. Under standard LDA modeling assumptions (e.g., independence and identical distribution of loss events within each UOM), such measures can be used to construct statistical tests of whether the estimated probability distribution for loss severity within a given UOM comports with the banking organization's historical loss severity data. Because one purpose of the AMA is to quantify the likelihood of incurring infrequent, large unexpected losses that may affect capital adequacy, the goodness-of-fit criteria used by a banking organization should emphasize the tail of the distribution. In addition, the criteria should be theoretically consistent with the underlying LDA modeling assumptions and promote reasonable stability in parameter estimates over time. At a minimum, when assessing the developmental evidence for a given model, examiners expect the banking organization's documentation for each UOM to include a variety of goodness-of-fit measures, including the Kolmogorov-Smirnov statistic, the Anderson-Darling statistic, and quantile-quantile plots, as well as summary measures of historical loss severities, such as a histogram of severities, the first several sample moments, and a tabulation of the largest loss events.

¹⁸ See, for example: 1) de Fontnouvelle, P.; Jordan, J.S.; and Rosengren, E.S., "Implications of Alternative Operational Risk Modeling Techniques," NBER Working Paper W11103 (2004); 2) Moscadelli, M., "The Modeling of Operational Risk: Experience With the Analysis of the Data Collected by the Basel Committee," Discussion Paper 517, Banca D'Italia (2004); and 3) Dutta, K.K., and Perry, J., "A Tale of the Tails: An Empirical Analysis of Loss Distribution Models for Estimating Operational Risk Capital," Working Paper 06-13, Federal Reserve Bank of Boston (2006).

Standard goodness-of-fit measures sometimes have little power to distinguish among competing model specifications. These situations can raise supervisory concerns to the extent that modeling choices are largely subjective and have a substantial impact on the banking organization's estimated operational risk exposure. For example, the modeling of extreme value theory-based distributions requires a choice of modeling threshold. Standard goodness-of-fit measures may provide little objective guidance in this area when the modeling threshold primarily affects the estimated probability distribution at very high severity levels, where there may be few, if any, historical observations. Hence, the choice of threshold may rely heavily on subjective visual inspection of data (e.g., hill plots and mean excess plots).

In cases where goodness-of-fit criteria alone are insufficient to differentiate among alternative model specifications, a banking organization should conduct ample sensitivity analysis to ensure that the impact of modeling choices on loss exposure estimates is well understood and that model selection reflects appropriate conservatism. More broadly, model selection should be governed by well-defined policies overseen by appropriate levels of senior management. Specific model choices should be well documented, and there should be effective control and challenge processes in place to ensure consistent and sound application of policies.

Over-fitting criteria: When modeling loss severity, the relatively small sample sizes found in operational risk data coupled with the high quantiles being estimated make over-fitting a significant concern. As the number of parameters or underlying complexity of the model increases, the estimated probability distribution for severity will often tend to converge around the empirical distribution of the historical data, meaning that the model will tend *not* to extrapolate much beyond the largest historical loss event. Also, over-fitting can produce estimated severity model parameters that are highly uncertain and unstable.

To address such concerns and to appropriately capture its operational risk exposure, a banking organization should be able to demonstrate that its model selection process does not result in over-fitting or excessive parameter uncertainty or instability. Model parsimony is desirable because it prevents the modeler from intended or unintended over-fitting of the data.¹⁹ When comparing alternative model specifications, a banking organization should employ well-established information criteria that penalize specifications with more parameters, such as the Akaike information criterion, the Bayesian information criterion, or the deviance information criterion. A banking organization should not increase the number of parameters unless there is a theoretical justification for the more complex distribution, the information criterion improves, and the benefit in terms of goodness-of-fit is substantial.

¹⁹ Gabaix, X., and Laibson, D., "The Seven Properties of Good Models," *The Foundations of Positive and Normative Economics: A Handbook*, Oxford University Press (2008).

Concerns about over-fitting often arise when using mixtures of probability distributions to model severity distributions.²⁰ Such techniques can greatly increase the number of severity model parameters that must be estimated. To limit over-fitting concerns in this context, banking organizations are discouraged from using mixtures of more than two distributions, and should use mixture distributions only when the historical severity data is heterogeneous. Unfortunately, spurious heterogeneity is quite common, especially in small datasets, and so a banking organization should also provide a strong justification for heterogeneity—for example, external and internal data are directly mixed, heterogeneous risk types have been combined within a single UOM due to data limitations, or, in the case of fraud, losses tend to be either low-severity events caused by individuals or infrequent high-severity events caused by organized crime. A specific technique for assessing the appropriateness of mixtures involves examining the relative weights assigned to each component of the mixture. When a single component of a mixture has a very low weight, that component is likely focusing on only a few data points. In these cases, a strong rationale or supporting evidence for inclusion of that component should be documented.

Estimation

Maximum likelihood estimation (MLE) has become the predominant methodology for estimating the parameters of severity models and diversification models (discussed later). Parameter estimates, however, can be quite sensitive to the choice of estimation technique. Thus, banking organizations are encouraged to explore the use of techniques other than MLE, such as Markov chain Monte Carlo methods, the method of moments, and quantile matching methods. If a non-MLE technique is selected as the basis for the model's final calibration, the banking organization should document its reasons for this decision and should carry out a sensitivity analysis relative to MLE. Regardless of the final fitting methodology, for each UOM a banking organization should document the confidence intervals around the severity model's estimated parameters and the implied estimate of operational risk exposure.

As with frequency models, parameter estimates for severity models should account for any modeling thresholds. A variety of modeling techniques for addressing truncated data are available for this purpose;²¹ however, a practice commonly referred to as shifting the distribution is appropriate only in limited circumstances. Subtracting the modeling threshold from losses before fitting distributions produces a “shifted distribution.”²² When the modeling threshold is above zero, the use of a shifted distribution for severity modeling may introduce bias into capital estimates by not capturing the probability

²⁰ Rousseau, K., and Mengerson, K., “Asymptotic Behavior of the Posterior Distribution in Over-Fitted Mixture Models,” *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* 73 (2011): 689-710.

²¹ Dempster, A.P.; Laird, N.M.; and Rubin, D.B., “Maximum Likelihood From Incomplete Data via the EM Algorithm,” *Journal of the Royal Statistical Society, Series B (Methodological)* 39 (1977): 1-38.

²² The term shifted distribution comes from the insurance industry, where deductibles are subtracted from losses before applying the loss distribution approach to modeling losses. The logic and reasoning used in insurance modeling do not apply to operational risk modeling when deductibles do not exist.

distribution for loss severities below the modeling threshold. Except in specific circumstances justified by underlying modeling or distribution assumptions (e.g., when using a generalized Pareto distribution to model the tail of a spliced or extreme value theory distribution),²³ a banking organization should not use a shifted distribution approach.

D. Diversification Modeling

Examiners recognize that diversification modeling is a highly challenging aspect of operational risk modeling. Under the rule, to use internal estimates of dependence a banking organization must demonstrate to the satisfaction of its primary federal regulator that its estimation process is sound, robust to a variety of scenarios, and implemented with integrity, and that it allows for the uncertainty surrounding the estimates; otherwise, a banking organization is required to sum exposure estimates across the UOMs.²⁴ Moreover, internal consistency within the overall AMA implies that, when quantifying these relationships, the banking organization's diversification model should be consistent with its UOM-level frequency and severity models. Thus, credible frequency and severity models at the level of each UOM are a precondition for a sound diversification model. The remainder of this section discusses additional modeling issues that frequently arise in the context of diversification modeling.

1. Copula Models

Copula models have emerged as the industry standard approach for quantifying operational loss dependencies among UOMs, reflecting the ability of copula models to capture a wide range of potential correlation structures and their analytical and empirical tractability.²⁵ The flexibility of the copula approach, however, often comes at a cost in terms of significant uncertainty around model parameters. Thus, the specification of the diversification model should reflect sound and well-defined internal goodness-of-fit and over-fitting criteria, and should incorporate adequate conservatism to address parameter uncertainty and ensure robustness to alternative loss scenarios.

When reviewing a banking organization's copula modeling approach, examiners consider whether the model is conceptually sound; whether the model reasonably captures the likelihood that when the loss in one UOM is relatively high, losses in other UOMs are likely to be relatively high as well (upper tail dependence); and whether the model incorporates sufficient conservatism in light of underlying uncertainties.

²³ Estimates of the shape and scale parameters for the generalized Pareto distribution are insensitive to shifting.

²⁴ See 12 CFR 3.122(g)(3)(D), OCC, and 12 CFR 217.122(g)(3)(D), Board.

²⁵ Trivedi, P.K., and Zimmer, D.M., "Copula Modeling: An Introduction for Practitioners," *Foundations and Trends in Econometrics* 1 (2005): 1-111.

2. Conceptual Soundness

Consistent with the one-year loss horizon of the AMA framework, the diversification model should assume at least one year for correlation effects to materialize. Because using annual data to estimate correlations will likely result in a lack of data, banking organizations should conduct sensitivity analyses around different time horizon assumptions.

The diversification model should also be consistent with other components of the LDA framework. For example, dependence of losses between UOMs may be caused by frequency dependence, severity dependence, or both. Thus, when modeling diversification, banking organizations should account for both potential sources of dependence. Modeling dependence between UOM aggregate losses has become the industry norm to address this challenge.

3. Upper Tail Dependence

Because the rule establishes a 99.9th percentile standard for quantifying a banking organization's overall operational risk exposure, the ultimate goal of diversification modeling is to quantify the degree to which annual operational risk losses among UOMs are interdependent at high loss levels. This can be a challenging statistical exercise. Even when internal loss data exhibit fairly low loss correlations among UOMs, such estimates can be misleading because of the sparseness of data at high loss levels. Moreover, standard measures of correlation often quantify dependence as a linear, constant relationship independent of loss size, and, therefore, are generally ill-suited for assessing dependencies at high loss levels. While copula-based dependence models can be an effective approach for dealing with these challenges, it is important that the model specification allows for the possibility of positive dependence when UOM losses fall within the upper tails of their distributions (upper tail dependence). In particular, copula specifications that effectively imply independence between UOM losses in the upper tails would not be appropriate, including Gaussian copulas (which assume zero asymptotic tail dependence) and t-copulas with many degrees of freedom (which have properties comparable to Gaussian copulas).

4. Conservatism in Dependence Modeling

Given the challenges in quantifying complex dependence relationships with limited data, the confidence intervals around the estimated parameters of a copula model can be quite wide, and standard goodness-of-fit diagnostics may have limited power to discriminate among alternative model specifications. In light of such modeling uncertainties, the agencies noted in the rule preamble that banking organizations are expected to adopt conservative assumptions when modeling dependence.²⁶ The final section of this guidance discusses the role of conservatism in dealing with AMA model risk generally. In the specific context of diversification modeling, at a minimum, examiners expect banking organizations to adopt the following prudent principles:

²⁶ 72 Federal Register 69, 317 (December 7, 2007).

- First, absent compelling conceptual or empirical support, a banking organization’s diversification model should not imply negative dependence between losses in different UOMs (i.e., a large loss in one UOM implying a small expected loss in another UOM). That is, pair-wise correlations should be floored at zero.
- Second, a banking organization should not presume that all pair-wise correlations have the same value. A single correlation value that appears “conservative” when compared to the average correlation across all UOMs may not capture markedly higher correlations among those UOMs having the greatest impact on the banking organization’s overall AMA capital charge. If specific pair-wise correlations cannot be measured with any confidence, then a banking organization should perform an analysis of dependence for clusters of UOMs taken together.

In addition, to ensure that estimates of dependence are robust to a variety of scenarios, a banking organization should conduct a comprehensive sensitivity analysis of its diversification model, including (but not limited to) evaluating the impacts on estimated operational risk exposures of different specifications for the underlying correlation structures and different families of copulas. Among the families of copulas considered, banking organizations are encouraged to explore using asymmetric copulas, which allow the degree of dependence among UOM losses to be greater at higher loss levels compared with lower loss levels.

E. Model Risk Management

Model risk is a critical consideration whenever banking organizations use models for risk management or regulatory capital calculation purposes. As outlined in the *Supervisory Guidance on Model Risk Management* issued in April 2011 (Model Risk Guidance), banking organizations should use a variety of methods to evaluate their models and ensure that their modeling processes are sound.²⁷ A key requirement of the rule is that banking organizations have sound processes for validating AMA quantification systems.²⁸ Another critical component of a sound modeling process is the application of appropriate conservatism in models’ specifications and calibrations. When dealing with inherent model risks, sound practice requires that banking organizations apply appropriate conservatism when specifying and calibrating their operational risk models. Consistent with the existing supervisory guidance related to credit risk quantification,²⁹ banking organizations should have in place processes for documenting how such conservatism has been achieved and, along with other elements of a banking organization’s advanced systems, this process should be subject to ongoing validation.

²⁷ See OCC Bulletin 2011-12, “Supervisory Guidance on Model Risk Management,” and Board SR letter 11-7, “Supervisory Guidance on Model Risk Management.”

²⁸ See 12 CFR 3.122(i)(4), OCC, and 12 CFR 217.122(i)(4), Board.

²⁹ See Basel Coordination Committee Bulletin 2013-5, “Applying the Requirement for Conservatism to the Parameters in the Advanced Approaches.”

1. Validation

Validation of AMA models should be consistent with the Model Risk Guidance, including a robust challenge process and monitoring of outcomes in relation to actual experience through back-testing and benchmarking at both the UOM and aggregate level. At the UOM level, a banking organization should undertake sufficient validation to ensure that the severity model, when used in tandem with the frequency model, produces estimated probability distributions for operational risk losses that are reasonable in light of the banking organization's historical loss experience. The banking organization should identify where the largest individual losses fit into the estimated severity distributions. In addition, the banking organization should benchmark various quantiles (e.g., 50th, 75th, and 90th) of the historical loss distributions against the quantiles implied by the model.

Validation should also be undertaken to assess the quality of the annual operational risk exposure estimate for each UOM against the AMA's 99.9th percentile standard. At a minimum, for each UOM the banking organization should compare the estimated operational risk exposure to the largest historical loss event and to the largest sum of historical losses over any four-quarter interval. Capital breaches (instances where the estimated operational risk exposure is less than these benchmarks) are potential indicators that the model does not adequately measure risk exposure.³⁰

2. Benchmarking

Validation should also include benchmarking of the banking organization's overall operational risk exposure estimate. An aggregate capital breach (e.g., an instance where the estimated overall operational risk exposure is less than the banking organization's largest cumulative loss over a four-quarter interval) may indicate that the operational risk model is fundamentally flawed. Another analysis examiners recommend compares the banking organization's estimated overall operational risk exposure against one or more nonparametric benchmarks. In one variant, the nonparametric benchmark is calculated as the 99.9th percentile of annual losses implied by a "bootstrapping" approach, which uses the same frequency and diversification estimates as the LDA, but replaces the estimated UOM severity distributions with empirical loss distributions. Examiners would review closely circumstances in which a banking organization's estimated operational risk exposure falls below such benchmarks at the UOM level or in the aggregate.

A banking organization management's internal benchmarking processes—including the use of peer analysis and comparison—are important inputs when assessing the credibility of operational risk models and their outputs and whether the banking organization has incorporated appropriate conservatism. While benchmarking and peer analysis may reflect differences in banking organizations' management of their risk exposures, these tools provide an important perspective for the model risk management process.

³⁰ Under a 99.9th percentile loss standard, capital breaches should only have a 0.1 percent likelihood of occurring. If the UOM breach is large or if it causes a capital breach at the consolidated bank or bank holding company level, it should be aggressively investigated to ensure model validity.

Examiners will review a banking organization's benchmarking approaches and the use of benchmark information in assessing the adequacy of operational risk capital estimates. In addition, examiners will use peer and supervisory information to assess whether a banking organization's estimated operational risk exposure is in line with the loss experience and estimated exposures for other banking organizations.