White Paper

Models, Model Risk and Running Effective Model Management Programs

Executive Summary

Financial institutions across the globe are extensively dependent on mathematical models for making business decisions and meeting regulatory requirements. With an explosion in available and captured data, and enhancements in technological capabilities, there has been a marked shift among banks to treat customers individually rather than as a group. This has increased the relevance of analytical models exponentially. For a variety of reasons, this process of building and implementing analytical models is fraught with risk, widely termed as 'model risk'.

These mathematical models attempt to simplify real world phenomena into mathematical equations by making multiples assumptions and approximations while also trying to make the model easy enough to comprehend and implement. A model is an attempt to replicate the realities of a dynamic world into an equation, but almost by definition, no model can fully do so.

This first layer of model risk is inherent and implicit in all models. We call this systematic model risk. The second layer of model risk is the risk associated with a specific model. This risk arises from multiple sources, explained in detail in later sections, which increases model risk for a particular model. These risks can be avoided or mitigated if organizations maintain a rigor and discipline in the model building process. We call this 'unsystematic' model risk.

The third layer of risk arises when consumers of model output ignore the fact that the output is the closest approximation of the truth, but could possibly be wrong. This occurs when model assumptions and approximations are glossed-over, ignored, or forgotten, and the model output is considered sacrosanct. This, strictly speaking, is not model risk, but underscores the risks arising from rampant model usage.

This paper focuses only on the second layer of model risk, and the industry best practices we have observed in mitigating model risk.

Regulators across the globe have issued regular guidance on model risk management. Most notable of these are Office of the Comptroller of the Currency (OCC) Bulletin 2011-12, "Supervisory Guidance on Model Risk Management", and Board SR letter 11-7, "Supervisory Guidance on Model Risk Management". As a part of a bank's Comprehensive Capital Analysis and Review (CCAR) submissions, banks are required to submit documentation regarding their model risk management policy and practices. Also, banks now actively publish their model risk management (MRM) practices in their annual reports.

Even as financial institutions, especially banks, seek to comply with regulatory norms, we believe there is a case to consider MRM as an integrated part of business strategy, rather than from a regulatory compliance perspective alone. In this paper, we outline the challenges that banks face in implementing efficient MRM programs, and our point of view on running such programs effectively.

Models and Model Inventory

What constitutes a model? Interestingly, this fundamental question draws a lot of debate and different banks define models differently. The SR 11-07 Guideline refers to a model as a quantitative method, system, or approach that applies statistical, economic, financial, or mathematical theories, techniques and assumptions to process input data into quantitative estimates.

> A model is a quantitative method, system, or approach that applies statistical, economic, financial, or mathematical theories, techniques and assumptions to process input data into quantitative estimates.
>
> - The SR 11-07 Guideline

Model inventory is a centralized repository of all models being used inside the bank and is a central part of supervisory reviews. To create such an inventory is seen as a challenge by banks because:

- The definition of a model is sometimes unclear in a bank.
- Traditionally, each line of business in a bank has been building and managing their own models with their own definitions, in respective silos and it is difficult to log it on the centralized model repository.

Some banks also have been scoping the model inventory by the risk tiering of the model, i.e. choosing to inventory only high impact models, and not inventorying models used for collections, staffing optimization, marketing models etc. See Footnote 1.

In our experience however, banks must take this opportunity to look beyond mere regulatory compliance. Hence, we recommend banks to inventory all their models, irrespective of their tiering and impact.

Models Risk

Following a brief discussion on models and model inventory, we will now move our focus on the central theme of model risk. Model risk is defined as the risk of loss arising from decisions based on incorrect or misused model outputs and reports.

The following quote from the OCC while releasing the SR 11-07 guidance mentions "...while model validation remains at the core of the new guidance, but the broader scope of model risk management encompasses model development, implementation, and use, as well as governance and controls related to models". What this implies is that model risk can occur anywhere in the lifecycle of the model.

As discussed earlier, systematic risk in models is unavoidable – only the degree of risk differs. A model of an airplane helps us understand the form, features and functions of a plane, but you can never replace a real plane. Similarly, a statistical model only tries to replicate human, price or portfolio behavior, but can never replace it.

The important sources of model risk that one needs to focus on is the unsystematic model risk, which can evolve primarily from four major sources across the model's life-cycle.

Figure 1 Sources of Model Risk

Model Errors

Model errors primarily arise in the form of models containing mathematical errors or assumptions that are misleading or inappropriate.

To give a better perspective, consider activities such as derivatives trading which depend heavily on complex equations and advanced mathematics. A model is said to be incorrect if there are mistakes in the analytical solution (i.e. in the set of equations or in the solution of a system of equations), or if it is based on wrong assumptions (e.g. about the underlying asset price process) – and this is perhaps both a common and dangerous risk.

Foonote

[1] While definitions can vary, in general a model is classified as "high impact" based on balance sheet influence, complexity and purpose. Generally, such models are significant from a corporate-wide perspective (i.e. corporate Value at Risk (VaR), economic capital, anti-money laundering (AML), etc.).

Provided below are a few examples of the most common errors in model assumptions, especially in the capital markets domain.

In options pricing, the volatility of the underlying asset is considered to be constant and while this was always known to be a risky assumption to make, market practitioners continue to make this assumption without revisiting the assumption even during stressed scenarios. Consider the volatility of the S&P 500 Index, measured by the Chicago Board Options Exchange Market Volatility Index (VIX) which was approximately 15 percent at the beginning of July, 2007, and by the end of July, it was over 30 percent. Later in the financial crisis, in early September 2008, the VIX was around 30 percent. Within two weeks, following the collapse of Lehman Brothers, it had jumped to over 80 percent. Market prices diverged significantly from the theoretical price and any decisions based on a theoretical model induced significant model risk in the balance sheet.

Another example is modelers often assuming that rates of return are normally distributed, i.e., they have a classic bell-shaped distribution. However, empirical evidence points to the existence of "fat tails" in many distributions; in these distributions, unlikely events are in fact much more common than would be the case if the distributions were normally distributed.

In some quantitative risk models, market risk might be measured at a one-day time horizon but may have to be capitalized at a 10-day horizon for regulatory purposes. Often the 10-day value at risk (VaR) is approximated by multiplying the one-day VaR by the square root of 10. The assumption is that the daily profit and loss (P&L) distribution is independent, yet identical across the time period. In reality, this assumption may not hold true.

Yet another example of approximation in market risk modelling is the use of sensitivities or, more generally, Taylor approximation. That is, the necessary re-valuation of certain instruments under changes to market risk factors is not performed by using an elaborate model, but simply by multiplying pre-computed sensitivities of the instruments with the changes (deltas) in the risk factors. The underlying assumption is that the behavior of the change in value of the instrument can be described as a linear function of the change in the relevant risk factors.

Data Errors

The quality of model outputs and hence the decisions we take based on them is dependent on the data that is fed into such models. We can consider the following sources of data issues:

Raw Data:

Raw data refers to the data that may enter the model in states as near to their source. This data may have to be checked for the following:

- Missing or Incomplete Data: e.g. Coverage with external data is only 90%.
- Inaccurate Data: e.g. For some products the entry for the notional value contains the market price.
- Outdated Data: e.g. Use of old data that is not changed or changed infrequently.
- Asynchronous Data: e.g. Data from some geographies may be older while others are current.
- Misinterpreted Data: e.g. Is the assumed notional amount at inception or maturity? Confusion whether the FX rate is EUR/USD or USD/EUR?

Expansion Data:

This is the data created by users; where there are gaps in the existing data which are attempted to be filled by duplication or other simple rules:

- Missing values: e.g. Substituting missing value with the mean or mode.
- Duplication: e.g. If a data point cannot be observed on a day, then use the previous day's value.

Sub-Model Data:

Where models use output of other models (internal or external vendor models), it becomes necessary to check the source models for issues. Consider the example of Expected Loss models which derive their input from Probability of Default (PD), Exposure at Default (EAD) and Loss Given Default (LGD) models. If there are issues with such upstream models, they may impact the downstream models severely.

Introducing the concept of data lineage and why it is important to model risk management

We have been observing regulators keenly focusing on the "data lineage" for models and regulatory reporting. We will elaborate on this aspect in this section.

Tracing Data Source is a Challenge

To manage models better, banks need to track and understand data lineage and dependencies in long data transformation chains.

A good data lineage system helps answer some of the following questions:

- Where does the data come from or go to in a specific column, table, view or model/report?
- When was the data loaded, updated or calculated in a specific column, table, view or model /report?
- Which components (models/reports) are impacted when other components are changed?
- Which data, structure or report is used by whom and when?
- What is the cost of making changes?
- What will break when we change something?

From a MRM perspective, regulators want to see whether the model data can be traced back to its source. With advanced visualization techniques available, it becomes easy to maintain and refer to such complex mappings. An example of one such mapping is provided in Figure 2.

As we have seen in the earlier parts of this section, models also feed into other models. It hence becomes necessary to track model dependencies, as depicted in Figure 3.

Implementation Errors

The next stage is implementing the model in production.

Once a model is built by the modelers in an environment (SAS, R, MATLAB, etc.) they are familiar with, these models have to be converted to a suitable programming language in which the model execution can take place in large IT systems.

Since this production code is manually written, it introduces the risk of errors in the code and these may be catastrophic for the financial institution, which uses the outputs of these flawed models for their day to day decision making.

Figure 2 Data Lineage System

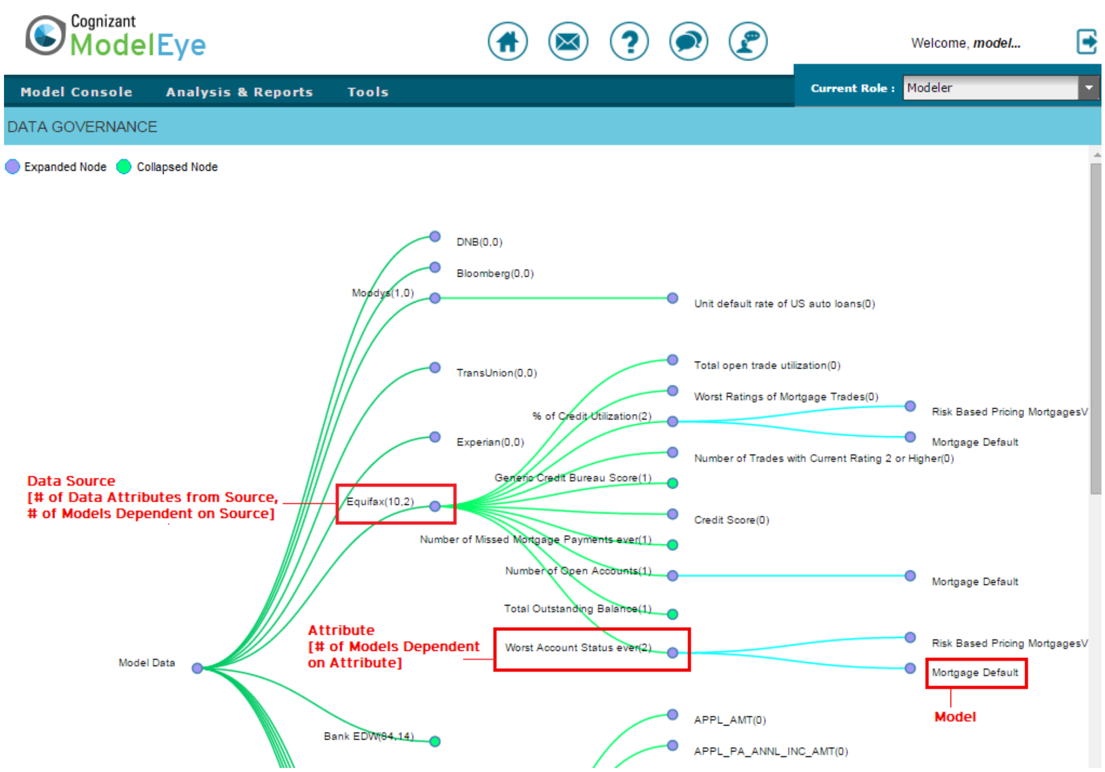

Figure 3 Tracking Model Dependencies

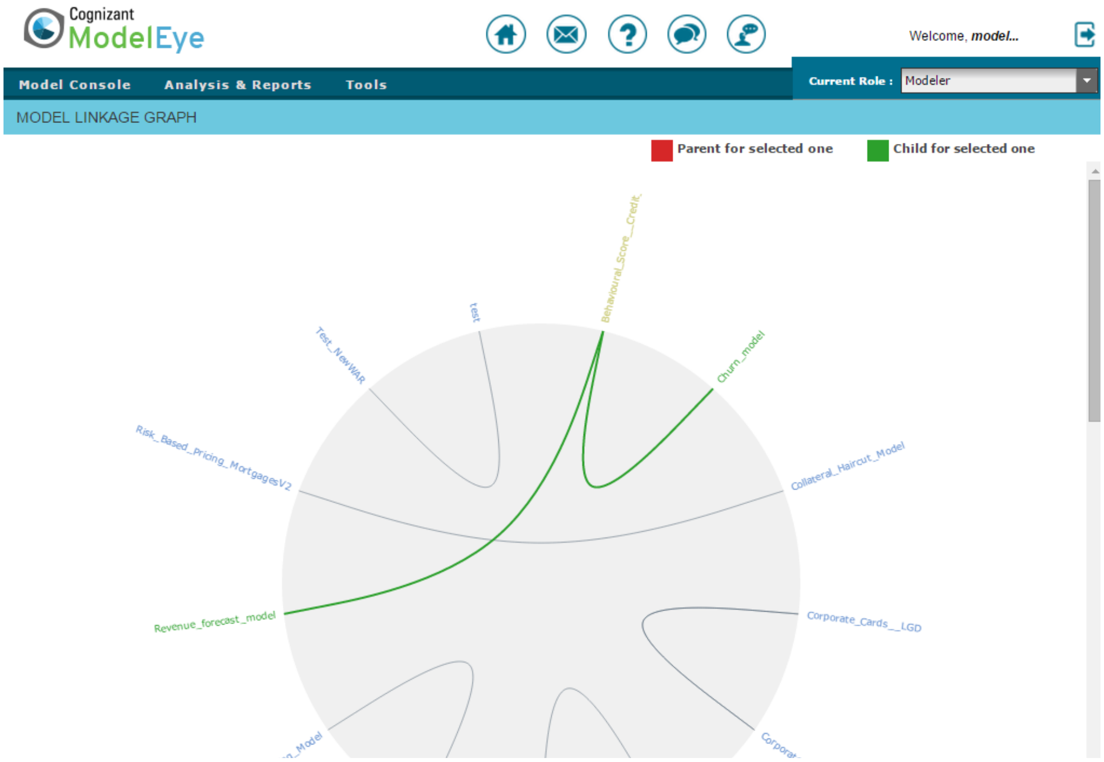

The AXA Rosenberg case, where three AXA Rosenberg entities settled a $242 million enforcement action by the U.S. Securities and Exchange Commission (SEC) relating to a computer programming error, is a classic example.

Here's how this model worked, according to the SEC filing:

The model consisted of three components: the Alpha Model, Risk Model, and Optimizer. The Alpha Model evaluates public companies based on their earnings and valuation. The Risk Model identifies risk on two primary bases – specific stock risk and common factor risks. Common factor risks include, among other things: (i) specific industry risks, which are risks associated with certain industries (such as oil, automobiles, or airlines); (ii) country risks and (iii) stock fundamental risks, which capture price to earnings ratios and similar metrics. The Optimizer takes the output from the Alpha and Risk Models, balances them against each other, and recommends an optimal portfolio for the client based on a benchmark chosen by the client, such as the S&P 500.

Then in 2007, Rosenberg commissioned a new version of the risk model. Computer programmers finished designing the program in 2009, and noticed that it was spitting out weird ("unexpected") results.

Some Risk Model components sent information to the Optimizer in decimals, while other components reported information in percentages; therefore the Optimizer had to convert the decimal information to percentages in order to effectively consider all the information on an equal footing.

Because proper scaling did not occur, certain decimal information was not converted to percentages and the Optimizer did not give the intended weight to common factor risks. This resulted in huge losses for the firm.

Usage Errors

A model can be mathematically correct, with the right assumptions and technically robust, and still be used inappropriately, thereby introducing model risk in the last leg of the model lifecycle.

For example, some term structure models that are widely used to value fixed-income instruments depend upon the assumption that forward rates are "log normal" — that is, their rates of change are normally distributed. This model seems to perform relatively well when applied to most of the world's markets most of the time — with the exception of Japan for the last 15 years and the United States and Europe in the immediate post-crisis years because Central Banks implemented quantitative easing in monetary policy, when such assumptions did not hold good. In these conditions, different statistical methods (e.g. Gaussian and square root models) for interest rates work much better.

In the same way, models that are safe to use for certain kinds of products might not perform well when applied to subtly different instruments. Many over-the-counter (OTC) products have options embedded within them that are ignored in the standard option pricing model. For example, using a model to value warrants may yield biased results if the warrant is also extendable.

Our Observations

The fundamental question that arises then is that why do modelers not address these issues upfront before models are put into production and used in decision making.

The answer lies in the fact that modelers are in a continuous struggle to find the best trade-off between speed of build, complexity (to better represent reality), simplicity (to improve the tractability of their modeling) and efficiency (demanding minimum run time for the model).

It is this trade-off which needs to be managed and any combination of the above parameters tend to bring different types of model risk into play. These aspects of model risk can be controlled by a robust model approval process and independent validation.

A robust model approval process and independent validation process check the models for its conceptual soundness. It evaluates all the assumptions made in the process of building the model, it evaluates the theories being used in those models, it questions the relevance of these theories and assumptions, and also checks the production code thereby helping to reduce model risk significantly.

In the end, a lot depends on the modeler himself and his ability to take the right trade-off in the entire process. In this light, we found the best guidance to modelers in the article "The Financial Modeler's Manifesto" by Emanuel Derman and Paul Willmott, where they formulate the Modelers Hippocratic Oath and urge modelers to treat models as they are to be treated, "as models, period".

The Modeler's Hippocratic Oath:

- I will remember that I didn't make the world, and it doesn't satisfy my equations.
- Though I will use models boldly to estimate value, I will not be overly impressed by mathematics.
- I will never sacrifice reality for elegance without explaining why I have done so.
- Nor will I give the people who use my model false comfort about its accuracy. Instead, I will make explicit its assumptions and oversights.
- I understand that my work may have enormous effects on society and the economy, many of them beyond my comprehension.

Running Effective Model Risk Management Programs

Running effective model risk management (MRM) programs requires planning and effective leadership, whose view goes beyond meeting regulatory expectations. Mentioned below are some points on what constitutes a good MRM program.

Formation of Model Risk Management Organization

The first step towards running an effective model risk management (MRM) program is instituting a qualified team of senior, experienced and empowered individuals to run the program. Banks have been forming model governance teams following the release of the guidance in 2011.

In some banks, we find the model governance group as an overarching layer above the modeling and model validation teams. The modelling teams and model owners largely stay in the business unit, and the validation teams are usually centralized with the MRM team directly under the command of the chief risk officer.

Another structure that we prefer and recommend is the "hybrid" structure, wherein some members of the MRM organization are drawn from the modeling teams or business units, some are drawn from the model validation team and the rest are from the model production teams.

We find this structure very effective as it focuses on bringing a "left-shift" in the process of ensuring model quality, which would mean that the best practices and institutionalized modeling knowledge (e.g. good documentation standards) start becoming a part of the model development cycle itself. However, this does not mean any compromise to independence in the validation process.

Review of Model Governance Policies

It is important for organizations to do a complete analysis of their existing model governance policies, compare it with the target state and with the expectations of regulators. All gaps should be noted and remedied for in the next release of the model governance policy. Care should be taken that the policies are written clearly and are as unambiguous as possible. Also, an action plan should be laid for addressing any of the gaps.

An important aspect of model governance is the standards of documentation mandated for each model. While modelers do documentation, it is often sub-optimal and quality is compromised in as teams move on to other activities drawing comfort from the fact that the modeler knows the model well. This prevents the institutionalization of knowledge, which is the biggest risk banks face since the knowledge of the model moves with the modeler.

Robust Model Inventorying and Model Approval Process

As discussed earlier, building a model inventory is the first step in carrying out effective MRM programs.

In order to ensure the inventorying process happens well, banks require a robust model approval process which would ensure no model can be put into production unless all the necessary policies have been followed and the model is approved by relevant stakeholders. We have also seen some banks following an attestation process where each business unit attests on a regular basis that they have been inventorying and reviewing all models. However, the effectiveness of such self-attestations cannot be solely relied upon and hence the "hybrid" structure helps in these cases.

An inventory should capture complete metadata about a model (including but not limited to its owner, risk-tiering of the model, dates of model build and reviews, implementation date, and most importantly information on all the assumptions, documents, linked models, assets and/or products it is linked to).

The bank should have a clearly defined model approval process that each model has to follow before a model is finally put into production.

Model Implementation Guidelines

MRM often stops at model validation and tends to downplay the importance in model implementation. Models take several weeks to implement and that takes away the ability of model users to leverage the model when it is most effective. Delays in model implementation are a risk in itself, apart from risks arising from errors in code. Banks should adopt processes to accelerate implementation, and completely review the production code before the code is rolled-out into production. The usual process of testing the code is to compare the output of the model in platform in which it was built to the platform in which it is getting implemented. It is important to note that such tests should also be performed under severe / extreme conditions as some models usually break in production in these extreme conditions.

The inventory should contain the latest version of the approved code and this should be the single source of code for the technology teams. As banks adopt meta-data management systems, it becomes easy for the technology teams to track data lineages before implementation. It should be ensured that the implementation team refers to this data lineage before implementing the model.

Model Validation and Ongoing Monitoring

Initial model validation prior to implementation and the monitoring or full-scope validation post implementation, are essential to keep the MRM process efficient. It is important to have a model monitoring plan decided and approved even before the model is put into production.

Model validation is a larger topic for discussion and hence we will not dwell upon that in this paper. However, we present a model validation framework in Figure 4 that showcases the major processes to be followed in such activity.

The frequency of such monitoring of the banks models is usually dependent on the risk-tiering of the models. It is important to archive and store the results of such monitoring across time so that it could be used for reporting and analysis.

Some model risk management systems provide for alerts to be sent to relevant model stakeholders on the performance of models based on pre-configured performance thresholds. This should be not limited to stakeholders of a particular model, but also extended to stakeholders of all linked models.

Approval processes should be in place for maintaining degraded models in production for unavoidable reasons with appropriate comments and justifications, and those should be signed-off by the relevant authorities.

Reporting and Dashboarding of Model Risk

The advantage of having a structured database on all model inventory, approval workflows, task management, documentation and monitoring is the ability to get a holistic enterprise wide view of the MRM status in an institution. This goes well beyond regulatory requirements but allows senior management to understand where model risk is building up, areas of policy violation, linkages of models with data sources and balance sheet impact of different analytics models. Banks have been in the process of creating such dashboards for their senior management, but today these dashboards are largely manually prepared, and some advanced banks are seeking automation of production and circulation of these reports.

Conclusion

As highlighted in this paper, we believe financial institutions must not approach model risk management purely as a compliance initiative. It needs to be driven, in equal measure, as a business and risk program that seeks to understand and then mitigate the different sources of model risk. This understanding is critical to ensure that all MRM policies address these sources of risk, so that model risk is effectively reduced. It also ensures that model risk management moves beyond a model inventory database and factual analysis of model performance to being a source for understanding potential risks associated with the asset side of the balance sheet.

Figure 4 Model Validation Framework

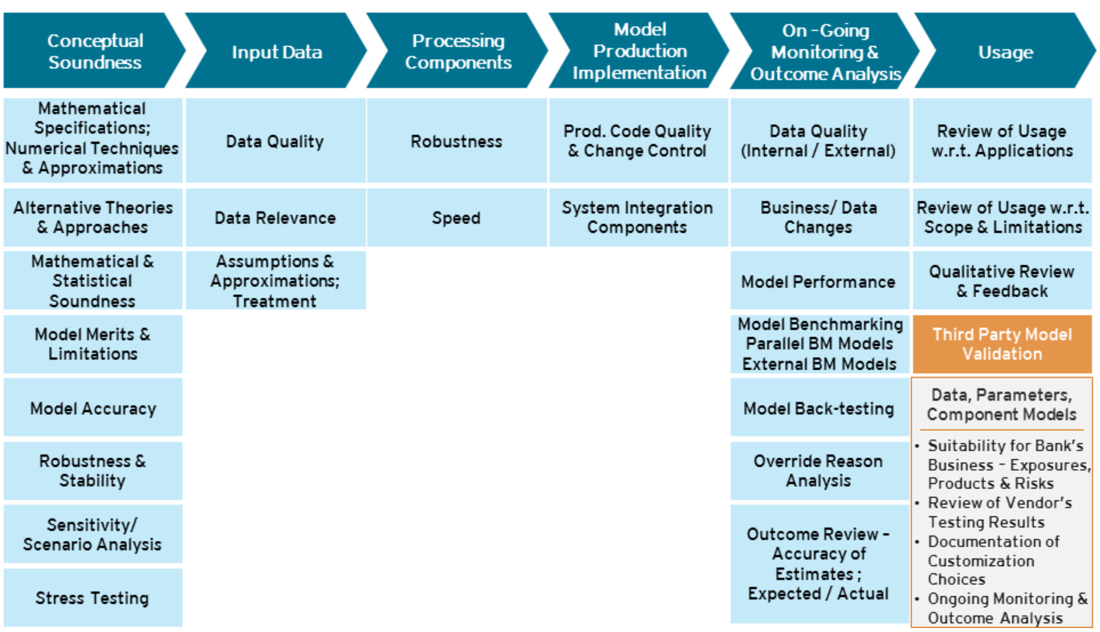

References

1. SUPERVISORY GUIDANCE ON MODEL RISK MANAGEMENT, SR Letter 11-7, Federal Reserve.
2. Risk Model Validation, A Practical Guide to Address the Key Questions, By Peter Quell and Christian Meyer.
3. Models Behaving Badly: Why Confusing Illusion with Reality Can Lead to Disaster, on Wall Street and in Life, By Emanuel Derman.
4. Cognizant's ModelEye.

About the Authors

Vasant Rao is a Senior Director and leads the Banking & Financial Services Analytics Practice within Cognizant. He has 18 years of experience in the financial services domain across risk, analytics, corporate finance and capital markets in multiple geographies. He can be reached at Vasant.Rao@cognizant.com.

Swaminathan Aiyer is a Senior Manager and leads the Products & Alliances group within the Banking & Financial Services Analytics Practice, with a specific focus on risk management solutions. He has extensive experience in analytics and credit bureau business having set up operations for numerous of credit information companies. He can be reached at Swaminathan.Aiyer@cognizant.com.

About Cognizant

Cognizant (NASDAQ: CTSH) is a leading provider of information technology, consulting, and business process outsourcing services, dedicated to helping the world's leading companies build stronger businesses. Headquartered in Teaneck, New Jersey (U.S.), Cognizant combines a passion for client satisfaction, technology innovation, deep industry and business process expertise, and a global, collaborative workforce that embodies the future of work. With over 50 delivery centers worldwide and approximately 137,700 employees as of December 31, 2011, Cognizant is a member of the NASDAQ-100, the S&P 500, the Forbes Global 2000, and the Fortune 500 and is ranked among the top performing and fastest growing companies in the world.

Visit us online at www.cognizant.com for more information or follow us on Twitter: Cognizant.

Cognizant

World Headquarters
500 Frank W. Burr Blvd.
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243
Toll Free: +1 888 937 3277
Email: inquiry@cognizant.com

European Headquarters
1 Kingdom Street
Paddington Central
London W2 6BD
Phone: +44 (0) 207 297 7600
Fax: +44 (0) 207 121 0102
Email: infouk@cognizant.com

India Operations Headquarters
#5/535, Old Mahabalipuram Road
Okkiyam Pettai, Thoraipakkam
Chennai, 600 096 India
Phone: +91 (0) 44 4209 6000
Fax: +91 (0) 44 4209 6060
Email: inquiryindia@cognizant.com