Article from

The Modeling Platform

April 2018
Issue 7

Model Governance: Controls and Culture

By Ryan Krisac, Zohair Motiwalla and Uri Sobel

As actuaries rely on increasingly complex models to influence decisions and calculate financial statements, model validation has become simultaneously increasingly difficult and crucial. In response, many insurers have instituted a Governance or Controls function that oversees these models. That impetus has been accelerated by regulations that mandate calculating reserves with actuarial models based on periodically unlocked assumptions, and reporting requirements supporting those models and assumptions. Regulators and auditors also expect certain reporting requirements and potential calculation demonstrations, which will be facilitated by proper controls.

Numerous studies, surveys and opinions have been published regarding implementing effective controls for maintaining and using complex models. It is important, however, to focus on the goals of an effective governance structure rather than simply craft or follow a set of prescriptive rules. This article identifies several ideas for an effective governance structure that may enhance the capabilities and validity of models, help fulfill regulatory demands, and provide feedback to gauge if the structures are effective.

THE GOALS OF MODEL GOVERNANCE

Why do companies invest significant resources in model governance?

1. **Modeling efficiency.** Strong model governance reduces runtime, since model code will be written to execute efficiently. Model maintenance, updates and upgrades will become streamlined. Human capital will be deployed with the appropriate skill sets and pay grades aligned with the appropriate stages of modeling. New users will become familiar with models faster. Digging up existing models for sensitivity and "what-if" testing will be easier.

2. **Consistency of results across the organization.** Conflicting and/or confusing model results drain company time and resources, and undermine confidence in the models. Solid governance ensures that model results deliver coherent and consistent messages.

3. **Validity, accuracy and applicability of model results.** Model results are only as useful as the validity of the model design and accuracy of its calculations. A robust baseline and ongoing review structure ensures that model design, inputs and calculations are implemented as intended, and that the model is fit for the purpose of its given application(s).

4. **Providing an audit trail.** No matter how thorough the governance structure and model reviews are, questions regarding model results will arise. A strong governance structure will help a company easily identify the source and calculations underlying model results. Many of the questions will have already been anticipated and answered in advance (and if not, companies will have an indication of where to begin to find the answer).

MODEL INVENTORY

Implementing an effective model controls framework starts with a company's basic definitions of what a model is, how many it has, and taking stock of the key features of each.

Different definitions of a model have been presented in various publications. For this article, a model is defined as any tool that involves inputs, calculations and result reports. This is intentionally vague. Actuaries must review the tools they employ for decision-making, reporting or other purposes valuable to the company and decide whether to consider these tools as models.

Consider the following when assessing a tool: Does the tool perform tasks beyond reorganizing inputs provided? Does the tool perform calculations that are either directly or indirectly reported, or are used by management for decision-making? Does the translation from inputs to outputs require review to ensure correctness? If "yes," then the tool may be considered a model.

Once a complete list of models is formed, the company can develop a model inventory. A helpful model inventory provides insight into the uses and operational risks of those models, as well as commentary on other important properties. Suggested characteristics for an effective model inventory are listed here:

- Operational risk rating, based on dimensions such as:
  - Reliability of calculations
  - Audience, particularly internal and external visibility
  - Capability of attribution analysis
  - Efficiency of calculations
  - Software-based user or vendor errors
  - Scope/materiality
  - Documentation rigor
  - Manual adjustments
  - Consideration in company decisions

Table 1
Model Inventory Sample Scorecard

Score	Reliability of Calculations and Data Input	Audience	Capability of Attribution Analysis	Efficiency
1 (low risk)	Highly reliable that model ran correctly; transparent calculations	Internal only/limited audience	Roll-forward process exists, easily built from production run	Little user intervention, timely results
2	Reliable after analysis		Roll-forward process exists, can be pulled from production run	Some user intervention, generally timely
3	Somewhat reliable after analysis; includes black-box calculations	Internal only/wide audience	Roll-forward process exists, but requires some additional work	User intervention, acceptable processing time
4	Somewhat unreliable; multiple simplifications and approximations		Roll-forward process exists, but requires significant work	Significant user intervention, long but predictable processing time
5 (high risk)	Unreliable; highly approximate methods used to gain comfort with results	External	Roll-forward process does not exist or is unreliable	Significant user intervention, unpredictable processing time

- Model confidence rating
- Model owner
- Eligible applications for the model
- Known limitations/exceptions
- Testing, vetting history
- Links to documentation (process for running, change history, financial impacts, etc.)
- Locations for published modeling files and model results

Often, building a sound inventory falls to a company's model steward in consultation with model owners, who will decide on the most appropriate dimensions to apply in determining the operational risk rating. The roles of model steward and model owner are discussed more fully later in this paper. Common definitions are required for the risk dimension conversation to be constructive. A scoring system—for example, on a scale of 1 to 5—in each dimension based on objective, observable information will clarify the relative model risks. An example scoring rubric is provided in Table 1.

Discussing the operational risk ratings in a group setting, such as in committees that span the organization, is recommended in order to eliminate bias and ensure that scales are applied consistently across models. Those responsible for a given model may be inclined to understate risk or overstate confidence so as to maintain their model status quo, rather than venture into a potentially complex, time-consuming model reconstruction project. However, an inventory must contain honest and consistent risk ratings to be credible.

The conversations to establish risk ratings are crucial for legitimizing the inventory as an effective means of comparing models in terms of operational risk. The inventory can then aid in prioritizing which models to review, improve, retire or change in some other way. By fostering understanding of model purpose, the inventory will help mitigate the risk of misuse.

Each model can also be assigned a confidence score. The difference between the two scores can help prioritize model improvements. An example of this model inventory gap analysis exercise is provided in Table 2.

In this case, work on the ULSG GAAP reserve model may be considered higher priority than long-term projections because the gap between overall risk and overall confidence is greater for the ULSG GAAP reserve model. This may be explained by having disparate models or fewer reliable testing tools, as reflected in the capability of attribution analysis and efficiency scores.

"Confidence" is a subjective term that indicates the comfort the owner and steward have in a certain model, counterbalancing

Table 2
Model Inventory Gap Analysis

Model	Reliability of Calculations and Data Input	Audience	Capability of Attribution Analysis	Efficiency	Overall Risk	Overall Confidence
Illustration actuary	4	4	3	1	**3.00**	**4**
Pricing model	4	3	2	4	**3.25**	**4**
Stat Vx	1	5	5	2	**3.25**	**3**
Cash-flow testing	3	5	4	4	**4.00**	**2**
Long-term projections	5	4	5	5	**4.75**	**3**
ULSG GAAP reserve	5	5	4	4	**4.50**	**2**

its inherent risks. Conversations regarding these ratings will help determine which models to review and the prioritization of those reviews.

MODELING ROLES

Another major decision for the governance framework is access—that the necessary people have the ability to change a given model (and that others do not). The underlying principle is as models become increasingly wide-ranging, nuanced and important, auditing those models becomes more complex. Controlling change-level access to models, therefore, limits the number of disparate, unnecessary or unintentional modifications and helps allocate the staff with the appropriate expertise to the appropriate tasks. This, in turn, promotes supportability and consistency.

Depending upon the complexity of the software involved, access may be easily segregated by an IT administrator. If model access can be granted via normal IT account management techniques, a natural audit trail will exist for all model changes.

Regardless of platform, all modelers may be classified in a role relative to the models they run or review. Consider four fundamental roles:

- User
- Developer
- Owner
- Steward

The typical delineation of responsibilities between users and developers lies in access to calculation code. The owner is the actuary designated with responsibility for a given model. The steward (who may or may not be an actuary) acts as a gatekeeper for the production environment of models, tasked with granting access approvals, updating the model inventory, and monitoring the overall effectiveness of the controls in place.

MODEL STEWARD IN THE COMPANY ORGANIZATIONAL CHART

When building out a formal model controls team, a decision must be made regarding where the model steward role fits in the organization. While there is no universal answer, a well-designed steward role maintains independence from the functional areas of the actuarial structure. In other words, the steward does not directly report to any owner, developer or user.

Thereby, the steward can provide unbiased opinions on all actuarial models—including any changes. In this sense, the steward must be able to challenge changes recommended by functional areas. This may range from revealing flaws in code during a focused review, to recommending coding standards that improve legibility. A strong steward role drives company models to achieve the goals of model governance and, in the process, reinforces the roles of actuaries in important company decisions.

PRODUCTION STATE FOR MODELS THAT ARE RECOGNIZED AND ACTUALLY USED

A tenet of model governance is the establishment of certain models as official (sometimes referred to as "published" or "productionized"), and that those models are recorded as such in the company's model inventory. These models generally receive the highest level of scrutiny and review and, therefore, carry expectations of being reliable, accurate, transparent and efficient.

An important corollary to this idea is that model *results* produced and provided throughout the organization are based on those official models directly (running the official model and capturing and summarizing output) or indirectly (with changes from the official models entirely defined and laid out). How should a company ensure that a user does not provide results from an unauthorized model?

One solution is to expand the concept of the model inventory to include model lineage. To establish clear lineage, model results

are annotated to indicate which model was the basis for the results, and to describe how (and why!) that model was developed from any of the officially published company models.

The level of review of the changes from the officially published models can also be documented with the model lineage. In many companies, a grade is assigned to a model that determines the level of review it requires. How far removed a model is from the officially published model may be a factor in determining the model's grade. The level of review in the model lineage is ideally commensurate with the model's grade.

An important key is for expectations to permeate the culture of the organization. In some cases, this can be accomplished by consistent messaging—for example that model results will not be considered unless they contain such model lineage. In other cases, technology can be leveraged such that model results can only be obtained from a results warehouse, and results can only be posted to the results warehouse if they contain such model lineage. The good news for modeling actuaries is that such a framework serves to underscore the importance of the work they do.

APPROVALS PROCEDURES

One must not associate model governance with the feeling that someone is looking for mistakes. Model governance is designed to help ensure the validity, accuracy and applicability of model results. A well-designed approvals process for model changes and model results has the potential to improve the overall quality of company models significantly. However, achieving buy-in across the organization may be a harder challenge than designing and implementing the governance structure itself.

First, let's review some items to address the specifics of an approvals process.

Authority

Clearly articulate which persons or groups are authorized to make model changes. Removing any ambiguity regarding who set the assumptions and endorsed the methodologies creates a reference point if and when questions arise.

Library/History of Decisions

A record of when significant model changes were made and the rationale behind them will help answer any questions. If any authorized party decides to revise a previous decision, such written history will help make decisions more efficient and transparent.

Review of Model Change Implementation

Once a model change is authorized, a strong governance process includes a structure to ensure that the model change is implemented correctly. Such a structure features independent roles responsible for

a. making the change,
b. checking the change and
c. reviewing the impact of the change.

The individuals responsible for each of these steps may differ depending on the nature of the model change. Errors can be introduced at either the formula coding or data input level, or both! Some modelers' expertise extends to understanding the formulas in the model coding, while others are assigned to implement or check inputs. Certainly, care needs to be taken that the people checking the model have the appropriate expertise to do so. Too often, the checking role is given to "higher-level" staff who do not understand the inner workings of the model, which can lead to missed problems and a less-than-thorough check. On the other hand, giving the task of input-checking to coding experts is an inefficient allocation of resources.

After a model change has been implemented, it is wise to circle back to the party that authorized the change for a final endorsement that the change matched its intentions (see "User Acceptance Testing" on page 31).

Documentation of Model Change Implementation

It is important to document how model changes were implemented and to include the authority under which they were implemented. Depending on the actuarial software, some of this documentation may exist or even be created on demand within the model itself. Developing separate documentation outside the model provides guidance to new and less familiar users. Evidence of unit testing, no-harm tests and appropriate levels of both technical and high-level peer review are critical to ensure confidence in the model change implementation.

This may be a good time to review the model code itself, in that well-written code can often serve as its own documentation. Modern code is usually easy to read, even for novices to a particular syntax. If an experienced actuary cannot read the code and understand the logic, consider rewriting the code.

Production

After a model change has been authorized, implemented, checked, reviewed and endorsed, it needs to be incorporated into the production version of the model—depending on the structure of a company's models—so that it will be promulgated to all other uses of the model throughout the organization. Such incorporation needs its own level of testing. The model must produce the expected results in the scenarios in which the model changes were developed, and must also produce reasonable results in other potential scenarios (including no-harm tests for certain business or situations where results are not expected to change). Again, a structure of checking and reviewing for this incorporation step will help reduce unexpected, or incorrect, model results.

Exceptions

If an off-production version of a model is needed, prominent documentation in the model inventory and on any results produced by the model must be created. The rationale behind this exception must be included to be considered sufficiently documented. The governance policy may specify who has the authority to grant such exceptions.

Model Lineage

As discussed earlier, having the model lineage attached to all the models and model results will bolster confidence that they are based on officially approved assumptions and methodologies, and will help improve the consistency of results throughout the organization.

Model Results Review

Given all the above, it is still important to review the results of the model before distributing them. The reviewer might consider whether the appropriate margins are reflected, or whether the model has an appropriate level of granularity. Also consider the level of detail provided to an intended audience, consistency with other results within the organization, and reasonability. This level of review ensures that all the prior governance policies were adhered to, including basing results on the latest production model, or understanding the rationale for any exceptions.

Cultural Aspects

Once the appropriate model governance structure has been laid out, model review responsibilities must be defined and carried out in a way that is transparent and educational rather than bureaucratic and prescriptive. Effective review policies also lead to clearer modeling best practices, which developers can readily understand and employ. This reduces key person dependencies and bolsters modeling knowledge across the company.

Reviewing results from a regular modeling project inherently requires expertise; therefore, the model owner will likely be responsible for that activity. Clear modeling roles can help steer the owner in their review. For example, if users can only modify model inputs, tools can be created to check these against desired assumptions, population metrics and other sources. If those inputs have been checked and have not identified the cause of the problem, the owner can then consult the inventory for a repository of coding changes and associated testing. Effectively controlled models have traceable change logs that facilitate model investigation.

What happens when a model change is not implemented as expected or does not produce expected results? Do reviewers and implementers become defensive? Are there attempts to convince each other of one's position just to be right? Are decisions based on power struggles? Such situations will not aid in producing the most efficient model. Here are a few ideas to help an organization maintain a strong model governance structure in these situations:

- **Start early.** From the beginning of a person's tenure, preach review as part of the organizational culture, applicable to everyone. Look to hire people open to this type of collaboration.

- **Normalize.** During group meetings or modeling user-group communications, regularly show mistakes (or less efficient implementations) that were found. Discuss why it was normal or subtle to make such a mistake, how it was found, and why the ultimate solution was better. This will help instill a sense throughout the modeling organization that "we're in it together."

- **Rotate roles.** If possible, make sure people serve in both implementation and checking/reviewing roles (for different model changes). This will help ensure nobody is always in the role of being picked on.

MODEL DOCUMENTATION

A full exploration of what to include in model documentation is beyond the scope of this article; however, here are a few questions to consider in this regard:

- Thorough model documentation includes statements of model purposes and limitations. For what applications can the model be used? For what uses is the model inappropriate, or what changes are necessary for it to be made appropriate?

- Are the authorities for model changes discussed earlier recorded and included in the documentation?

- Are the model governance policies and structures themselves documented, with details of model owners and users included?

Typically, commercial actuarial software includes general user guides and documentation, but a company-specific process manual, detailing step-by-step instructions on how to update each element of the model (specific inputs and/or code) and specifying precedent and dependent code is immensely helpful documentation. In general, more detail is better, with descriptions of any internal checks to confirm throughout the process. However, such reliance on step-by-step instructions and internal checks must not lead to a sense of complacency regarding one's own model. Occasionally, it is prudent for a reviewer to observe the process of following the documentation to update the model, in order to provide fresh eyes and question any existing approaches that may no longer be appropriate.

MODEL LIFE CYCLE

Models go through life cycle stages, and a model's life cycle stage informs how the model fits in the governance structure. Broadly, models progress through the following stages:

- Defining purpose
- Initial development
- Vetting and implementation
- Putting into production
- Ongoing review
- Modification
- Retirement

Models in the model inventory are assumed to be in the "ongoing review" or "modification" stages, and will likely alternate between those stages until "retirement." Model reviews are commonly dictated by scores on the model inventory, an audit request, or the desire for a new modeling tool. Given available resources, model reviews can also be preventive as scheduled; focused reviews may reveal problems or opportunities that were previously overlooked.

Choosing who is responsible for performing a model review will be guided, in part, by the balance between expertise required and time available. A centralized steward role is likely to have more time to research but may require significant education on the nuances of the model. On the other hand, the steward may be in a better position to leverage new and outside perspectives or other related model improvement projects.

Regardless of the responsible party, a test plan clarifies the testing and documentation requirements necessary before signing off that the model may be placed back into production. Some common types of testing include:

- **Regression testing.** Run a set of production models that are not directly related to the new code, and confirm zero financial impact.

- **Match testing.** Run the model(s) associated with the coding change. Verify that the change in results matches a company published source.

- **Impact testing.** Run the model(s) associated with the coding change, and gauge reasonability of financial impact, seeking sign-off from model owner.

- **User acceptance testing (UAT).** Confirm the outcomes of the other tests and review any requirements not yet approved. The key is to maintain independence in the final UAT step. Production model changes must not be unilaterally dictated. Instead, they are joint efforts toward commonly understood goals. UAT ensures that owner and steward mutually agree that the model review was performed according to the test plan and that, if necessary, any changes were appropriate.

The steward updates the inventory to reflect the outcomes of the review, with new rankings and details. Findings, especially resulting changes, are communicated with all model owners. Modeling is improved as the findings shed light on better modeling or testing methods to be employed elsewhere.

OTHER BENEFITS OF MODEL GOVERNANCE

The primary advantages of an effective model governance program were described earlier in this paper: efficiency, consistency, validity, applicability and supportability. If too audit-driven—especially to satisfy bare minimum standards—model governance can veer into bureaucracy and superficiality. If designed toward broader model understanding, though, governance provides benefits across the actuarial teams and beyond.

Thus, one of the great benefits of governance is model literacy, with far-reaching effects on the ability of actuaries to understand what their models are calculating, how they can be modified and how to validate results. With a shared model inventory, all the modeling participants become more familiar with the risks, purposes, capabilities, limitations and vulnerabilities of the company's models. With effective governance, subject-matter experts rightly leverage their expertise on products and processes, rather than focusing on manipulating modeling software. As increasingly complex calculations are becoming standard, more actuarial models will rely on complicated coding modules in order to maintain reasonable runtime. The actuarial group must keep pace with understanding and, when necessary, debugging those modules.

Occasionally (or for some companies, frequently), a model improvement project may be too disruptive, time-consuming or complex for a specific functional area to tackle. In these situations, a dedicated team can perform the coding work and initial testing, consulting with the functional area for its expertise and UAT. When sharing the end result with other model owners, this collaboration further contributes to the common knowledge of a given modeling improvement and how it might be leveraged elsewhere.

Another benefit of effective governance is increasing management's understanding of the inherent risks of the company's actuarial models, relative to other enterprise risks. In a broader sense, model governance can be incorporated with company-wide controls programs, Own Risk and Solvency Assessment (ORSA) modeling and change management procedures. Merging with these other programs helps communicate model risk to nontechnical audiences and executives. A sound model inventory

demonstrates how models are performing without requiring extensive knowledge of the calculations. Also, incorporating standard IT change management procedures can strengthen the connection between actuarial and IT teams, potentially revealing many opportunities for automation. Links outside actuarial teams underscore the value that model governance provides, while also clarifying model risk for upper management.

A well-functioning model governance program enables the company to steer conversations with auditors better. Actuaries can provide evidence of controls and tests from readily available materials. Necessary sign-offs, documentation and validations can be provided for model updates in consistent formats, while change management procedures and documentation confirm that no unintended alterations were made to production models. These materials can be presented as self-audit findings and demonstrations of comfort developed in the normal course of business, instead of being instructed to prove, after the fact, that calculations were correct.

GOVERNANCE OF GOVERNANCE

We conclude with some suggestions for confirming that the company's controls and governance structures are working, and improvement ideas if they are not—essentially, governance of the governance process (in these authors' opinion, governance of governance just becomes silly). Consider the following:

- Incorporate "incidence reporting" into the model inventory, tracking errors and responding to questions such as:
  - How were the errors discovered? By whom?
  - What steps were taken previously to attempt to prevent such an error? Why were they not effective? What will we now do differently, and why will it be more effective?
  - In what other models might this error reside?
  - What impact does discovering this error have on the risk score of the model under consideration, or other models?

  An "incident" does not need to be limited to the discovery of a mistake. It could also be when conflicting (or even misunderstood) results are provided, or when validation issues were overlooked until a later point. Crucially, incidence reporting should not be practiced as a blame game, but as an opportunity to improve the governance, modeling and efficiency of the company.
- Incorporate runtime statistics as part of the model inventory. Runtime is influenced by many factors other than governance, but maintaining such statistics may help determine general trends in governance efficacy.
- Maintaining metrics on cost and hours required for model upgrades and enhancements can indicate if the appropriate personnel are allocated to the appropriate tasks. Narrative descriptions of the pitfalls and efficiencies realized during upgrade and enhancement efforts can indicate the efficacy of the governance process and structure. Further, sharing these narratives among company modelers promotes teamwork, standardization and education.
- Are they the right people running company models?
- Regularly review notes from modeling group meetings. Are people struggling to understand the controls and governance? Are there patterns to modeling challenges that point to holes in the governance structure?
- Consider occasionally shifting staff to test the thoroughness and reliability of existing documentation. Similarly, having subject-matter experts (who typically serve in a reviewing role) run through the whole modeling process themselves may uncover issues that lower-level users may not notice.

Einstein had it right: $E = MC^2$, or, Efficiency = Model Control times Culture. ■

Ryan Krisac, FSA, MAAA, is director of model controls at Penn Mutual Life Insurance Company in Horsham, Pennsylvania. He can be reached at *Krisac.Ryan@pennmutual.com*.

Zohair Motiwalla, FSA, MAAA, is a principal and consulting actuary at Milliman in Seattle. He can be reached at *zohair.motiwalla@milliman.com*.

Uri Sobel, FSA, MAAA, is a principal and consulting actuary at Milliman in Little Falls, New Jersey. He can be reached at *uri.sobel@milliman.com*.