



Achieving Near-Real-Time

Risk Monitoring

Risk Appetite Revisited

Achieving Near-Real-Time Risk Monitoring

Risk Appetite Revisited



Contents

Introduction	1
Risk Measurement Models	3
Risk Monitoring With First-Generation Risk Models	4
Enterprise Risk Monitoring Case Study	6
Risk Measurement Models	15
For More Information	17

Achieving Near-Real-Time Risk Monitoring is the second paper in a three-part series that revisits the concept of risk appetite and how insurers can extract full value from its implementation. The first paper in the series, *Another Bite at the Apple*, looks at enhancing risk appetite articulation with clearer linkages to the business.

Visit [towerswatson.com/riskappetite](https://www.towerswatson.com/riskappetite).

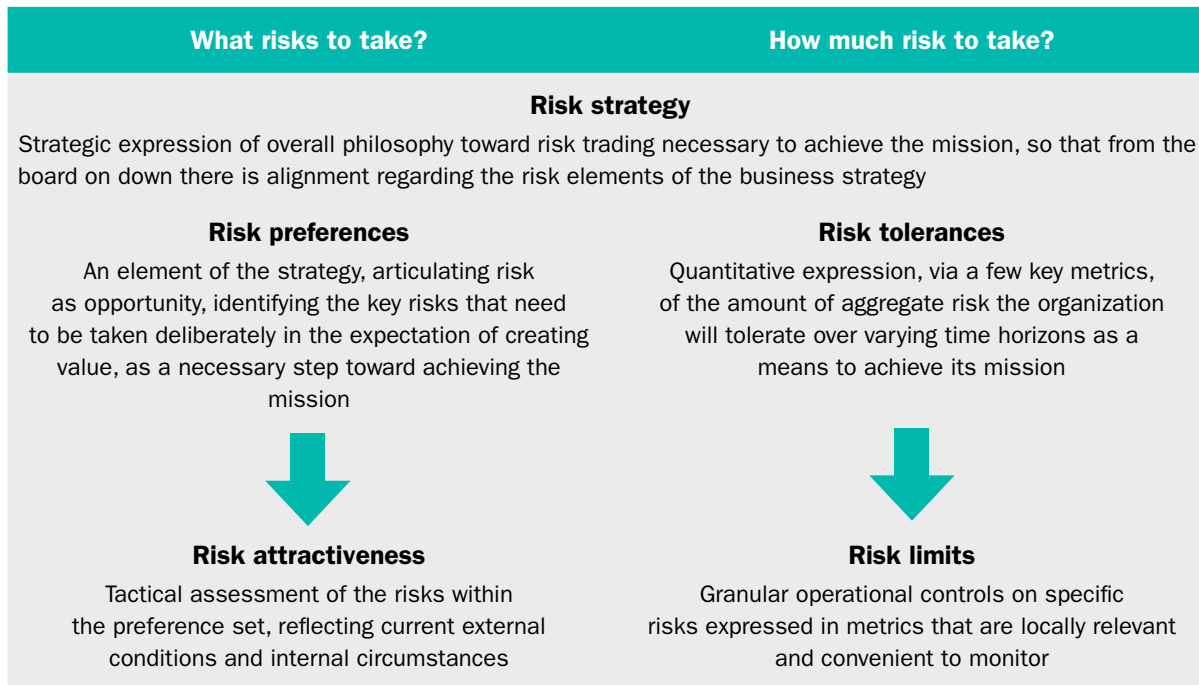
Introduction

While many insurers have by now developed risk appetite statements, they have also expressed dissatisfaction with the exercise — risk appetites must be more clearly linked to business mission and strategy, and levels of risk must be continuously monitored against risk appetite. This requires risk models capable of quick turnaround, so that management can make appropriate decisions based on near-real-time risk information.

In this paper, we will discuss the importance, practicality and utility of a near-real-time enterprise risk measurement model to monitor risk appetite. Internally, the risk appetite framework shown in *Figure 1* provides a basis for consistent risk-trading decisions; externally, it gives stakeholders comfort regarding the company's exposure to risk. We will focus on the right column of the risk appetite framework and look more specifically at risk tolerances, which are a quantitative extension of the risk strategy.

Risk tolerances express quantitatively the amount of aggregate risk that the company is willing to accept, expressed in probabilistic terms, time horizons and unacceptable mission impairment impacts. Risk tolerances are set at the overall enterprise level, across the full spectrum of risks contemplated by the business strategy. Most importantly, risk tolerances must be monitorable to assure that the tolerances aren't breached, and actionable to assure that if tolerances are breached, they can be managed back within tolerances. Otherwise, the risk appetite process becomes largely academic. To be monitorable, the actual levels of accumulated risk relative to the tolerance must be measurable on an ongoing basis (not just annually), with sufficiently timely results to allow management to act as risk tolerances are approached or breached. (For a more complete discussion of the risk appetite framework, see our companion paper, *Another Bite at the Apple: Risk Appetite Revisited.*)

Figure 1. Elements of the risk appetite framework



Risk limits are more granular than risk tolerances. They are set for specific risk sources, business units or products, and are used to help implement the risk tolerances. Risk limits are often expressed using more practical metrics that are measurable and relevant to managers at the local level. Risk limits must be set so that they are aligned and linked with risk tolerances, so that they are effective in controlling to risk tolerances. (Establishing this linkage will be the focus of the third paper in this series.)

Risk tolerances can be linked to *adaptive buffers*, organized around four *risk quadrants*. These concepts are illustrated in *Figure 2*. Adaptive buffers are resources that allow the organization to manage through any bumps in the road that may occur. For example, holding buffer capital, above

the minimum level required to avoid regulatory or rating agency actions, affords management the ability to adjust business plans after a capital loss. Bank liquidity facilities and reinsurance can act as adaptive buffers. Adaptive buffers can be financial or nonfinancial in nature, spanning all of the four quadrants shown and covering all mission-critical aspects of the business.

Adaptive buffers are linked to risk tolerances, because these tolerances often express the enterprise's willingness to potentially exhaust buffers through risk-trading activities. For example, many companies express a capital risk tolerance in terms of their unwillingness to take risk above levels that expose their buffer capital to a modeled probability of exhaustion above a predetermined percentage.

Figure 2. Four quadrants for adaptive buffers and risk tolerances



Risk Measurement Models

We define a *risk measurement model* generically, as a system that measures the financial impact of one or more *risk drivers on a business portfolio*. Examples of risk measurement models include credit risk measurement models for fixed-income portfolios, catastrophe risk measurement models for property insurance portfolios and life insurance policy projection models, as well as non-life claim simulation models. Companies may have any number of risk measurement models. These can be stochastic, involving thousands of trials, or deterministic, focusing on a defined set of stress scenarios. In a risk measurement model, the financial impacts can be measured in a variety of ways relevant to one or more of the four risk quadrants, such as loss of statutory capital, diminution of IFRS earnings or strain on liquidity. Risk drivers are external and internal factors that can cause these impacts. Business portfolios are collections of relatively homogeneous assets or liabilities that are susceptible to the risk drivers. *Figure 3* illustrates these three key dimensions of risk measurement models.

An *enterprise* risk measurement model is a special case, in which the business portfolio is the entire organization, and risk is aggregated across all of its businesses and calculated using enterprise-level financial impacts. Enterprise risk measurement models would include internal capital models under Solvency II, as well as the groupwide capital adequacy models used for Own Risk Solvency Assessments (ORSAs) and the comprehensive capital assessment models now required of some insurers by the U.S. Federal Reserve. They would also include company-wide models that assess the potential for subpar earnings or other financial performance metrics, or the potential for liquidity problems. Enterprise risk models are essential to the monitoring of risk against the company's chosen risk tolerances.

“Enterprise risk models are essential to the monitoring of risk against the company’s chosen risk tolerances.”

Figure 3. Examples of elements of a risk measurement model



Risk Monitoring With First-Generation Risk Measurement Models

Once an organization has set its risk tolerances and limits, the focus naturally shifts to a monitoring process that assures that actual risk levels stay within them. However, many companies have found that, having set tolerances and limits, they don't actually have the risk measurement models in place to achieve effective and timely monitoring. Typically, while some monitoring of specific risks against limits is possible, collecting all risks to get a comprehensive enterprise view within a reasonable time frame is not. And the lack of an enterprise view of risk substantially inhibits the ability to meaningfully manage overall risk against enterprise risk tolerances.

The risk monitoring shortcomings of existing models at many insurers are rooted in the models' development. The first-generation risk measurement models were often designed to meet a variety of applications rather than specifically developed for enterprise-wide risk monitoring. These other applications were generally more tolerant of extended run times and required a high level of detail to accurately capture the optionality of the underlying insurance contracts and assets. In fact, since use of these systems was embryonic and risk management still under development, near-real-time monitoring wasn't yet a priority at most companies. The need for real-time information did not become fully apparent until the financial crisis.

The issues with first-generation risk measurement models are manifold, but center on the following:

- Enterprise risk measurement models are often large and complex, requiring substantial resources to maintain them and substantial time to run them. This is particularly true for models at large, diverse organizations. Many of these models require point-in-time snapshots of detailed risk portfolios (at the individual policy or security level), extracted from transactional systems, that take substantial effort and resources to obtain, load into the model and run through even a modest number of scenarios. These limitations reduce the frequency of model runs and output, which is only available annually or perhaps quarterly at best.
- Turnaround can be measured in months because of model size and complexity. So risk information is only available significantly in arrears, making the information substantially less useful for management decision making.
- More specialized models can sometimes produce some information more frequently or on a timelier basis, but cannot produce the metrics that are needed or that are inconsistent with the core risk measurement models. For example, at some companies, credit risk models within the investment department are well developed and produce near-real-time measurements, but they are typically not integrated with the policy liabilities and are unable to capture the net risk position necessary for an enterprise view.

“The risk monitoring shortcomings of existing models at many insurers are rooted in the models' development.”



Clearly, without the ability to effectively monitor enterprise risk, risk tolerances are largely an academic exercise. Companies that can demonstrate that a process is in place to set them may satisfy regulatory and rating agency requirements, but the only way to extract real value from the investment in building risk measurement models is to have them drive decisions within the business.

Some have questioned the need for frequent measurement of actual risk levels, given that risk portfolios evolve slowly over time. This is true for many risk portfolios, such as an established block of auto insurance policies, where customer turnover is relatively slow, and growth is modest. In these cases, the amount of insurance risk won't be expected to change very much from month to month, and monthly movements in the amount of aggregate insurance risk won't be very interesting. However, we believe this argument misses the point: In normal times, most monitoring will not be exciting. It is in times of crisis, when the situation is changing dramatically, that management will be looking for real-time assessments to use in making critical decisions.

The recent credit crisis drives this point home. Prior to the crisis, mortgage-backed securities were perceived as a relatively safe, stable asset class. As our auto insurance example illustrates, monthly measurement of the risks associated with the mortgage-backed security (MBS) portfolio wouldn't change very much. However, during the crisis, the MBS market collapsed — along with the stock market, some banks and one large multinational insurer. We know firsthand that as the crisis unfolded, management at several companies was looking to the CRO for answers on a daily basis about risk levels relative to tolerances, and the impact on capital, earnings and liquidity — and the CRO was inadequately prepared to respond.

Finally, we would point out that even if the level of risk doesn't change much over time, the level of available buffers may be subject to substantial fluctuations. Risk tolerances and limits are often a function of both. So, for example, the capital tolerance could be breached by a decline in actual capital, even if the modeled requirement for capital remained relatively constant.



Enterprise Risk Monitoring Case Study

To illustrate one potential solution to these issues, we present a case study. Our case study company is an amalgamation of several multinational, multiline insurers, writing both individual and commercial insurance, which were all caught in the crosshairs during the credit crisis. The results are combined and stylized so that none of the companies are discernible. Finally, while our case study focuses on an enterprise economic capital model, the concepts apply equally to monitoring of tolerances relating to other risk quadrants such as liquidity and financial performance. These are touched on at the end of the paper.

We also want to make clear that the case study describes one solution for near-real-time monitoring. The reader should not infer that the approach taken by the case study companies is the only solution, as other approaches may be suitable.

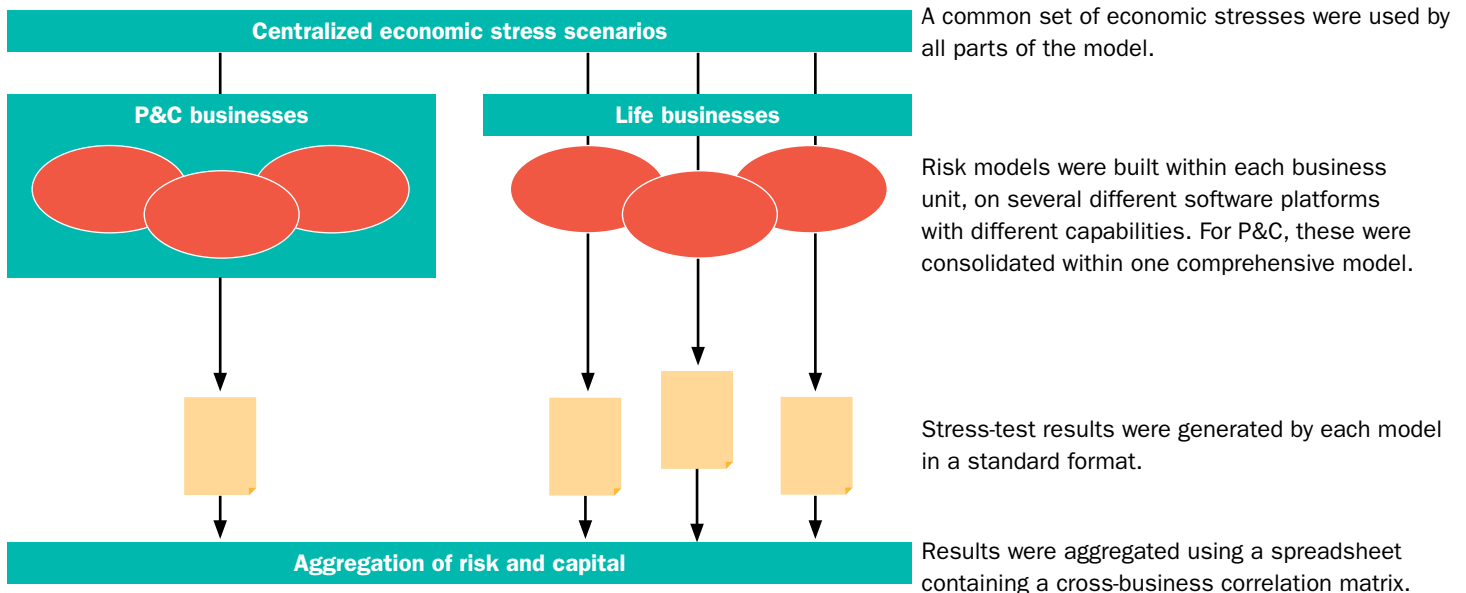
As the credit crisis unfolded, managers were asking for information about risk and capital that CROs were simply unequipped to provide. This included questions about the erosion of the actual capital buffer relative to targets and the extent to which key risk tolerances had been breached. There were also more granular questions about the level of risk in various parts of the business in light of the dramatic shifts in the markets and environment that were taking place.

After the crisis had stabilized, some very frank internal discussions took place focusing on what critical information management needed and why the CROs were unable to respond in a timely way. These discussions led to a fundamental rethinking of the enterprise economic capital model that had been developed.

The first-generation aggregation model

With the benefit of hindsight, the company recognized that, up front, before even building the enterprise economic capital model, it had done an inadequate job of identifying its business requirements. As is the case at many companies, the initial requirements focused principally on “building a tool that would accurately measure and aggregate the risks of the business to provide a view of capital adequacy.” Because its business was diverse, the resulting model was very large and complex. The developers also sought to leverage existing submodels that were built for narrower risk measurement needs within different parts of the business. This led to a decentralized model with components built and maintained by different business units. The complexity made the enterprise model both opaque and inflexible, and the decentralized approach made for long turnaround times. *Figure 4* is a conceptual schematic of the original enterprise risk aggregation model.

Figure 4. Conceptual schematic of the original enterprise economic capital model



As shown, the core of the enterprise model is comprised of a series of risk measurement submodels built within each business unit, reflecting the business unit's own view of risk and any necessary local requirements (e.g., cash-flow testing for life insurers). These submodels were built using different approaches and different software platforms. With different inputs, the submodels produced different outputs. The enterprise model was essentially a "wrapper" around these submodels to standardize inputs and generate overall output at the lowest common denominator. To produce an aggregate view, it was necessary to refresh and run all submodels.


The property & casualty (P&C) business chose to use a single comprehensive stochastic submodel, with modules for each product within each business. The level of detail in the submodel was quite granular and varied by product to reflect specific characteristics. Only attritional claims were modeled inside the P&C submodel; a file of outputs from vendor catastrophe models was imported to include catastrophe claims. The P&C submodel also incorporated the various outward reinsurance covers that were part of the current program, necessitating the modeling of individual large claims for some product lines. The P&C submodel was stochastic, capable of producing a full distribution of outcomes across many trials. However, a substantial computational platform was required to achieve reasonable run times across 10,000 trials.

In contrast, the life business chose to use distinct individual risk measurement submodels for each major product in each country. Each of these was highly detailed, taking policy-level and asset-level data files as inputs. Because of the complexity of the required liability calculations, only a set of deterministic stress scenarios could be run through them within the available timelines. For each stress scenario, the output consisted of the economic impact of that stress versus the base case scenario.

To obtain an enterprise view of required capital, the company had to specify a set of stress scenarios to be used by all submodels. The stress scenarios were reasonably comprehensive across risk categories, but to accommodate the run-time limitations of the life submodels, they were limited to events at a single probability level. While this approach made aggregation of a single risk additive, aggregation across risks was very approximate, using a simple correlation matrix. So the stochastic capabilities of the P&C model were largely unused in the exercise, as the stress capabilities were necessarily set at the lowest common denominator.

“The level of detail in the submodel was quite granular and varied by product to reflect specific characteristics.”

Prior to the crisis, the overall model was run annually based on data at year-end, with results available in late April. Information at the corporate level was limited to a comparison of required economic capital to the actual available financial resources at year-end. There was no mechanism to refresh the results during the year. Very little sensitivity testing had been done, given the resource requirements to produce alternative runs. It is fair to say that while the submodels worked reasonably well for local risk management purposes, senior management had very limited confidence in the aggregate capital results the aggregate model produced because of the limited sensitivity testing.



The second-generation aggregation model

In its redevelopment effort, the company started by setting a clear set of business requirements for the new enterprise capital model. These focused on the model's use as a decision-making and support tool for the business, particularly in the context of a dynamic risk environment. The CRO stated at the time that "effective use in decision making is the only way to realize value from our investment in the model."

The company set the following high-level business requirements for the new enterprise economic capital model:

- The results from the model must be available in near real time, so that they are usable and relevant to the decisions management is making. This will include routine, timely production quickly after the close of each month, plus urgent ad hoc production when circumstances require.
- The results must also be arrayed and distributed in a form that is usable in decision making by business unit leaders throughout the organization. This means the results must be presented in a format that is simple, consistent, understandable and actionable, and circulated with commentary that highlights emerging issues on a regular basis.
- The model must be substantially more transparent, with clarity around risk distribution and dependency assumptions, so that management is more confident in making decisions based on its output. The CRO should be able to explain the results produced by the model, and it should be relatively easy to drill down to obtain additional insight. An important path to transparency is sensitivity testing. The model must be capable of sufficient sensitivity testing to achieve understanding and acceptance by management. Sensitivity tests around key assumptions must be sufficient to assure that the limitations of the model are understood. Testing will focus on the robustness of results as they relate to decisions, as this is the litmus test for any decision support model and the correct basis for assessing model risk.



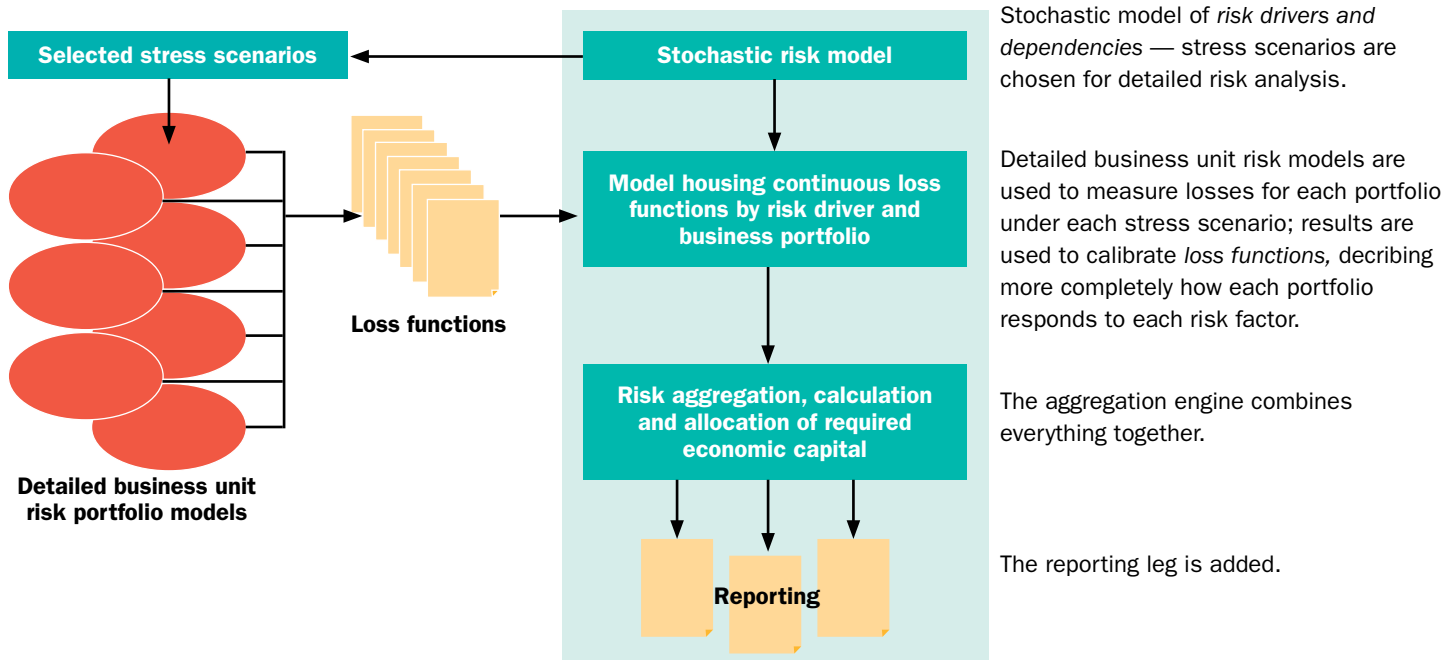
- The model must be sufficiently flexible to accommodate changes in the business and environment in which it operates, so that usable output will be available during these times. The model must facilitate changes to assumptions in response to changes in the environment without massive redevelopment and retesting. It must similarly accommodate acquisitions, divestitures and restructuring.
- Existing business unit submodels must still play a role, to the extent that they have the necessary detail needed to produce highly accurate measurements of the impact of key risks on specific business portfolios. Measurement of the same risk in the submodels and the enterprise model must be consistent.

"An important path to transparency is sensitivity testing."

With these business requirements delineated, the company set out to design and build a new enterprise model. Considerable time was spent on thoughtful design, process and output, all necessary to achieve the requirements.

Figure 5 on the next page displays the conceptual schematic of the new enterprise economic capital aggregation model. While the new model contains many of the elements of the original model, when the two are compared, one can see that a significant restructuring has occurred.

Figure 5. Conceptual schematic of new enterprise economic capital model



In viewing Figure 5, it is important to understand that the new enterprise capital model is only the shaded part shown in the center. The detailed business unit submodels are still used for risk measurement, but in the context of the new enterprise capital model, they are only used to derive inputs, described more fully below. Once these inputs are set, the enterprise capital model can be run without them. This restructuring goes a long way toward improving run times and throughput.

Of course, the business unit submodels continue to be used for risk analysis and other risk measurement applications. The new enterprise model doesn't replace them, but is really just a different wrapper around the core models, designed to meet the near-real-time monitoring needs and other business requirements.

At the top of the new enterprise capital model is a new stochastic risk model. The company identified roughly 100 *risk drivers* (interest rates, credit spreads, longevity, pandemic, claim cost misestimation, catastrophes, inflation and foreign exchange rates, among other factors) that it believed could have a material impact on the business. These were built into a distinct model, with probability distributions for each risk and the dependency structure between risks represented by copulas. The stochastic risk model can generate many trials quickly, as there is no attempt (at this stage) to link the risk drivers to their financial impact.

Note that the stochastic risk model need only be rerun when a risk distribution or copula parameter is changed by new research, or a change in conditions; an intermediate output file containing random trial values for each risk driver can be saved and reused in the downstream part of the model. The separation of this process greatly facilitates speed and sensitivity testing.

In the new model, the detailed risk measurement is organized into relatively homogeneous *business portfolios*, reflecting the various classes of financial assets and liabilities that comprise the company's business. As stated earlier, a business portfolio can be a business unit or a product within a business unit; the business portfolio can include the assets that support the product or those assets can be placed in separate business portfolios. Ultimately, the choice of business portfolio structure within the model reflects how the business and its risks are managed in practice, so that model output can be linked to clear accountabilities for the management of risk taking for each portfolio. The only requirements are that (a) every part of the business is incorporated into one and only one business portfolio, and (b) a submodel exists to measure the financial impact of risk on every business portfolio.

The case study company has selected a business portfolio structure that is reasonably granular, with approximately 75 distinct portfolios. These are defined in a manner that allows the business to be rolled up to either a legal entity or a business segment view of risk that would, respectively, be useful for external regulatory purposes and internal management purposes.

A key breakthrough in the design of the new enterprise model was recognizing that the enterprise capital calculations required only a small subset of the output available from the detailed submodels. Consequently, the calculations could be abstracted up to a higher level, so that all of the detailed submodels did not need to be rerun every time new results were requested. This was accomplished via the creation of loss functions, which are analytic representations of the financial impacts of the risk drivers on each given business portfolio.

In the current implementation, the loss functions measure the financial impacts in terms of economic gains and losses from the baseline, using the company's internal framework for economic valuation. However, the model itself is agnostic regarding the choice of measurement yardstick. The company could have just as easily chosen to measure financial impacts in terms of changes in statutory surplus or IFRS net assets.

Essentially, each loss function is an equation that calculates the impact directly from the values of the risk drivers without the need to go back to the detailed business portfolio model. To illustrate, a simple loss function might describe the financial impact of a shift in the risk-free yield curve on an existing portfolio of intermediate-term U.S. Treasury bonds. (For example, given an upward shift of 100 basis points, what is the loss in the market value of the portfolio?) While one could go back to the detailed portfolio and get an accurate measure of the impact by revaluing each security, the loss function provides a shortcut to the answer.

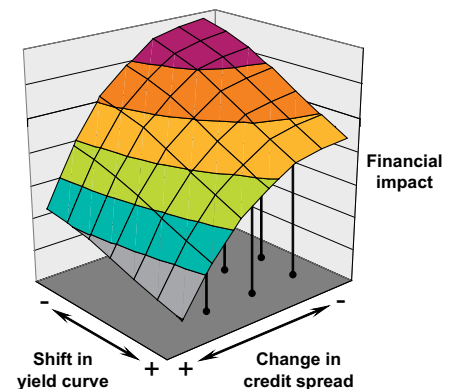
- Loss functions are fitted to the financial impacts produced by the detailed risk measurement submodels. Typically, a defined set of stress-test scenarios are run through the detailed models to produce a set of discrete impacts at a fixed set of sample points. The loss function is then fitted to these points to provide a continuous set of impacts (Figure 6).

- The fit of the loss functions to the sample points is not exact, but it is very close. Ultimately, the user must decide on the level of precision; additional terms can be added to the loss function to achieve whatever degree of precision is required, with the trade-off being increased complexity to the loss function and greater effort required in the fitting process.
- Because business portfolios can be influenced by multiple risk drivers, the loss functions are multidimensional surfaces — in practice, often functions of 10 or more risk drivers. In Figure 6, we show an example comprising two risk factors, with the loss function shown using the third dimension.
- Calculating the financial impact of a single risk scenario using an analytic loss function is almost infinitely faster than calculating them using the underlying business portfolio models. For example, the difference might be a nanosecond versus multiple hours of run time (or even several days for a large life insurer).
- Developing the loss functions is not an insubstantial effort. Each submodel for each business portfolio must be run through the defined stresses for the relevant risk drivers, and a loss function must then be fit to the results. However, the exercise is done on a periodic basis, depending on the sensitivity of the loss portfolio to the risk factors, and the work can be scheduled to occur throughout the year. This approach effectively separates the ongoing research into risk drivers and loss functions from the production of model results.

“Loss functions are fitted to the financial impacts produced by the detailed risk measurement submodels.”

Figure 6. Illustrative loss function for investment-grade bond corporate portfolio

- Financial impacts are measured for a sample of risk factor scenarios, using the detailed risk measurement models to obtain the most accurate measures possible.
 - Sample scenarios are illustrated by the dots with the lines extending upward from the base.
 - Sample scenarios include linear movements (up or down) in the risk-free Treasury yield curve and changes in the corporate bond market conditions.
- Loss function is a fitted surface to the resulting sample of financial impacts.
 - Fitted impacts are tolerably close to measured sample impacts.
 - Continuous surface provides estimated impacts at all intermediate points.



Once the loss functions are derived from each of the 75 detailed business portfolio models, they are combined with the stochastic risk model to produce an aggregate distribution of financial impacts at the enterprise level, reflecting the specified probabilities and dependency structure among the risk drivers. In turn, the aggregate distribution can be used to assess required economic capital and other risk metrics.

Model output for management

The main output from the enterprise capital model is the company's required economic capital. However, in contrast to the predecessor model, the new model does not produce just a number. Rather, it can produce the required economic capital at any chosen risk statistic, for any chosen collection of business portfolios. For example, the chosen risk statistic could be the 99.5th percentile financial loss or the 99% conditional tail expectation. Users can easily interrogate the model to see the results at different percentiles. Similarly, the chosen portfolios could be all those belonging to a particular legal entity, or all those belonging to a business unit or profit center, or a single product line across all legal entities. The first is used for regulatory reporting purposes, while the second and third are used for management purposes.

In addition, for the chosen statistic, the model allocates the required economic capital based on the aggregate risk distribution back to individual business portfolios and to risk drivers to give a complete picture of the sources of risk that are driving results. *Figure 7* presents an illustrative example.

This is a summary view for senior management, with the 75 business portfolios rolled up into four major business segments and the 100 risk drivers rolled up into 10 major categories. The model has the capability to display this grid at virtually any level of detail desired, so that more detailed grids are available (e.g., relating to each major business unit of the company).

The rows of the grid (the risk drivers) reflect all the identified key risks within the model that could impair the company's mission through capital depletion. The columns of the grid divide the overall business into distinct business portfolios, showing how they contribute.

The row totals in the left block of figures show the total amount of capital associated with each risk driver across the whole business, with the balance of the row indicating the relative contribution to capital for that risk from each business unit.

Conversely, the totals at the bottom of the table show the total amount of capital associated with each business unit, with the column above showing the capital contribution from each source of risk for that business unit.

Of course, allocating the required capital to the cells of the grid requires the adoption of a capital allocation method. Methods for capital allocation, and the various pitfalls of different methods, could be the subject of an entire paper.

“The main output from the enterprise capital model is the company's required economic capital.”

Figure 7. Required economic capital by business unit and risk driver (illustrative)

Risk category	Business segment			Life segment		Business segment	
	Life	P&C	Total	Individual	Corporate	Individual	Corporate
Credit spread	1,244	1,087	2,331	871	373	217	870
Asset default	852	381	1,233	596	256	76	305
Counterparty default	5	159	164	-	5	95	64
Credit risks subtotal	2,101	1,627	3,728	1,467	634	389	1,238
Interest rates	1,658	18	1,676	1,326	332	4	15
Equity markets	6,744	212	6,956	6,070	674	42	170
Foreign exchange rates	758	98	857	531	228	20	79
Market risks subtotal	9,160	329	9,489	7,927	1,233	66	263
Life insurance	1,864	-	1,864	1,118	746	-	-
P&C insurance: Reserves	-	2,152	2,152	-	-	323	1,829
P&C insurance: Pricing	-	1,576	1,576	-	-	315	1,261
P&C insurance: Catastrophe	-	5,728	5,728	-	-	4,582	1,146
Insurance risks subtotal	1,864	9,456	11,320	1,118	746	5,220	4,236
Total	13,125	11,412	24,537	10,512	2,613	5,675	5,737

Returning to Figure 7, one can see that the grid itself is a powerful source of management information, as it shows how the various risks and business units are contributing to the overall need for capital by the enterprise. In essence, economic capital becomes the standard unit of measure for all types of risk, across all types of business unit.

The information displayed in Figure 7 is the aggregate required capital created by the risks taken by each business unit, with the output showing the sources of the requirement by business unit and risk driver. One can invert the process, establishing risk budgets — representing the maximum allowable allocation to a business unit or a risk driver, based on risk appetite considerations — as an input and reworking the business plan to assure that the company operates within the budgets. The information in Figure 7 would then allow one to discern the emerging risk hot spots, as well as areas of potential underutilized risk capacity, versus the risk budget.

Making the model near real time

So far, we have described how the new enterprise model is structured to achieve quicker run times and throughput, but we haven't addressed how the model facilitates near-real-time risk monitoring and reporting. This key functionality is introduced via *scale* and *shift parameters* applied to both the risk-driver distributions and the loss functions. Scale parameters magnify or diminish without shifting the origin; shift parameters move the origin without a change in size. Examples of the application of scale and shift parameters are:

- As a particular business portfolio shrinks or grows over time, the model provides a convenient scale adjustment that magnifies the loss function proportionately. For example, if the in-force business grows by 10%, then one can approximate the impact on the loss function by scaling it up by 10%.
- If the risk driver environment should change, then the model provides a convenient shift adjustment to account for the change. For example, if credit spreads widen, the risk distribution can be shifted to a base at the higher spread level.

There are a great many additional details to the scale and shift parameter capabilities that are beyond the scope of this paper. We merely want to give the reader a sense of the general approach.

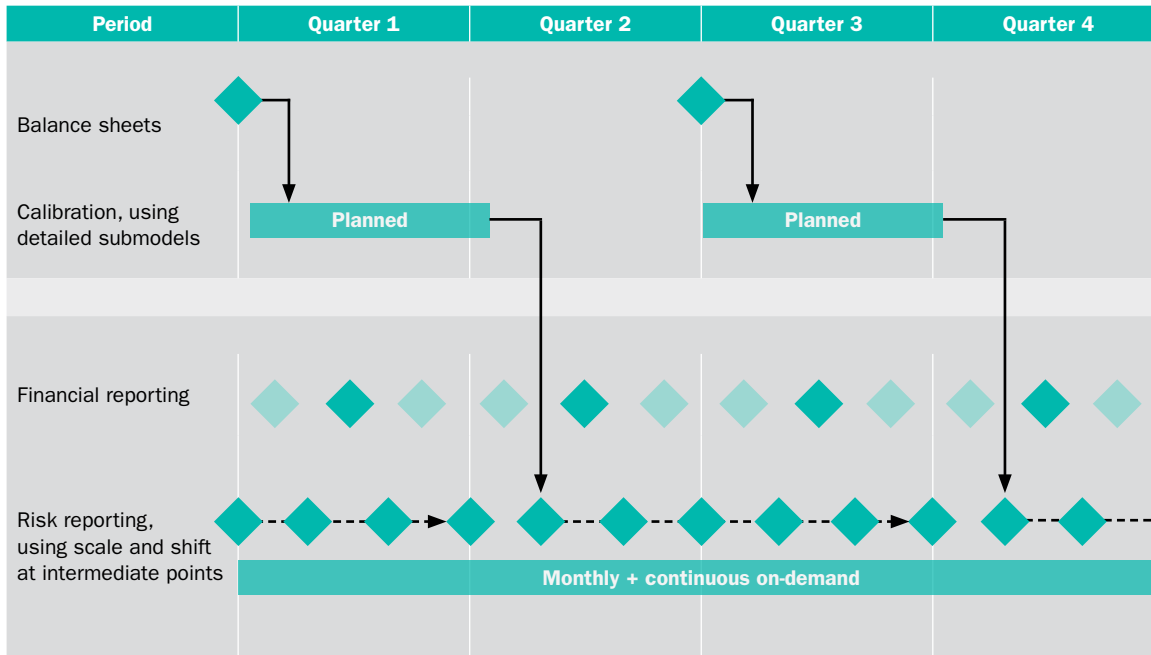
To implement the near-real-time capabilities, the company needed to establish ongoing processes

to generate the information needed for scaling and shifting. Management chose to establish a monthly update cycle, consistent with internal financial management reporting. Each month, information on the growth and evolution of each business portfolio is produced, along with statistics on changes in each risk driver. Procedures have been developed to determine how to translate that information into scale and shift parameters. The scale and shift parameters are input into the model to modify the original risk distributions and loss functions. This allows updates to the output, such as that presented in Figure 7, to be produced on a monthly basis, shortly after the monthly management reports are available. Because the model runs quickly, results are available as soon as the scale and shift parameters have been updated. To illustrate the point further, the company has developed a schedule like the one illustrated in Figure 8 on page 13 that defines the planned frequency and timing of recalibrations. This facilitates resource management over the course of the year.





Figure 8. Enterprise capital model reporting and recalibration schedule



In addition to monthly production of *required* economic capital, the company also uses the enterprise capital model to measure *available* capital, permitting capital adequacy and related risk tolerances to be assessed on an ongoing basis.

Of course, the scale and shift parameters are approximate, with increasing errors over time as the business portfolios evolve. The extent of approximation errors dictates the frequency with which the loss functions need to be recalibrated by reference back to the detailed risk models. However, the experience to date suggests that approximation errors are not much of an issue for normal evolution, at least for the majority of the business portfolios. As a safeguard, the company does some spot testing by running a few detailed models on current

portfolios, and comparing the measured impact from the detailed models to the approximated impact via the scale and shift parameters. When portfolios are undergoing more fundamental change, more frequent updates will be required.

Finally, we need to emphasize that the implementation of the near-real-time enterprise capital model, with its abstracted loss functions, does not negate the need for the detailed portfolio risk measurement models. The latter are still necessary to support ongoing recalibration of the loss functions, as well as to meet other, more local, risk management needs within the business units.

Meeting the business requirements

The company believes that the new model has successfully met the business requirements that were established for it. In this section, we offer a few comments on each major requirement.

Near real time. This business objective has largely been met, as the model can produce updated results on a monthly basis, shortly after the management reports are available. In addition, in special circumstances, the model can be run more frequently — as quickly as the appropriate scale and shift parameters can be determined.

Useful for decision making. The information from the model is significantly better than the single number produced by the predecessor model. The ability to distinguish capital deployment by business unit and risk driver provides real insight. As comfort with the results of the model has grown, results are increasingly being used within the business.

Transparency. This is perhaps the least tangible but most important success of the new enterprise capital model. The risk driver/business portfolio/loss function paradigm has facilitated a much better understanding of the model. In turn, the paradigm has channeled the internal debate about risk in a useful and productive manner. The parameterization of the risk drivers and dependency structure generated considerable internal debate, going to fundamental beliefs about risk. This debate occurred without reference to the financial impacts — a healthy separation of thought processes. The ability to perform sensitivity tests on all key assumptions also contributed to model understanding and a much higher degree of comfort with the model output. Since transparency is an important source of model risk, the company believes it has substantially reduced its model risk with the new approach.

Flexibility. Because of the model's business portfolio structure, it has been easy for the company to consider strategic alternatives such as exit strategies for certain product lines or sale of certain business units. In addition, the scale and shift capabilities have allowed the company to look at alternative growth strategies.



Planned future extensions

As the use of the model increases, the company is considering several extensions that it believes will make it even more valuable as a management tool.

- Currently, the model's financial impacts are all measured using an internal economic valuation framework. The company may create a second instance of the model measuring the financial impacts according to U.S. GAAP (its financial reporting standard). This would give the company the ability to look at risk issues through a second, important lens. Creation of a GAAP instance of the model would necessitate the construction of a second set of loss functions, reflecting the income effects of the risk drivers under GAAP. While this would entail some extra work, it is far less than a doubling of the loss function effort. Many of the detailed models can already measure impacts on both an economic and GAAP basis. In some cases, the GAAP impacts are virtually a by-product of the current economic measurement process.
- The company is also considering adding a risk dashboard to the back end of the model. This would change the format and the delivery of existing information, and perhaps add relevant information from other sources. Delivery of this type of information via dashboard would represent a key step forward in the use and acceptance of the model, and would be an important signal in the development of the company's risk culture.
- The financial impacts are currently measured over a one-year time horizon, consistent with Solvency II requirements. However, the company is considering whether it would like to have an additional instance of the model, which would look at cumulative capital losses over a three-year time horizon — a useful capital-planning tool.

“The ability to distinguish capital deployment by business unit and risk driver provides real insight.”

Risk Measurement Models

Beyond the Case Study

The case study presents one solution for achieving near-real-time monitoring of enterprise risk tolerances relating specifically to capital. As we have said, while we believe it represents a good solution, there may be equally good alternative solutions. For example, some companies may be able to address the run-time issues of more detailed models through technology such as grid computing, which facilitates massively accelerated computational speed. Also, certain elements of the case study solution, such as the use of risk drivers and the two-dimensional capital allocation plan, can be implemented within the context of a first-generation-style model.

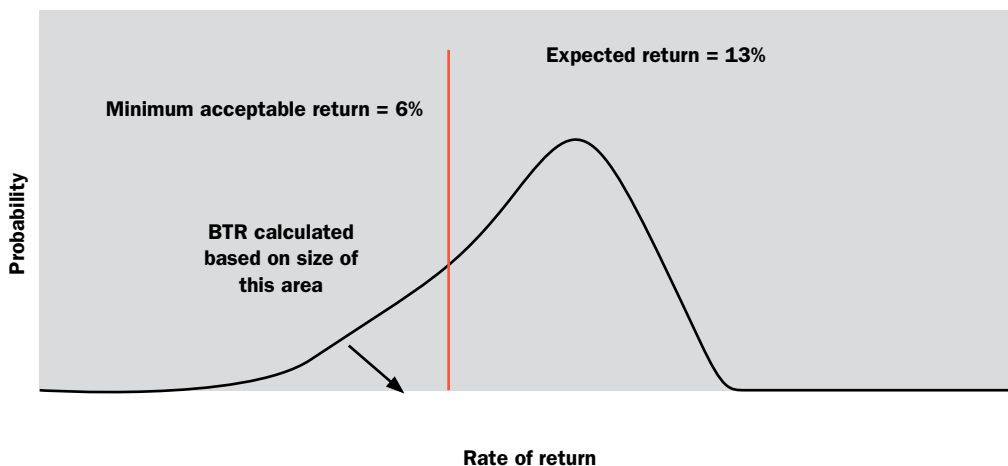
It is also important to point out that the case study approach can be extended beyond economic capital modeling to consider other key risk tolerances — in any of the four risk quadrants — that can be monitored using an enterprise risk measurement model. Such models could be focused on earnings (measured either on a public reporting basis or an internal economic basis) or liquidity.

Our concept of risk budgeting is not unique to the risk of capital loss — for example, one could set a risk tolerance relating to the risk of financial nonperformance, and use an earnings-based risk measurement model to allocate the risk of nonperformance to business units and risk drivers in a manner analogous to what was described in the case study. To illustrate further, we are aware that some companies measure the risk of financial nonperformance using below-target risk (BTR) as their metric. They employ a risk measurement model that generates a probability distribution around their annual GAAP return on equity, then measure BTR as the downside standard deviation of the returns, considering only the adverse results, measured as the distance below a minimally acceptable target return (*Figure 9*). These companies could set risk tolerances around acceptable levels of BTR, and allocate the enterprise BTR to business units and risk drivers as part of their risk budgeting process. The resulting allocation would then show each business portfolio's contribution to the risk of enterprise nonperformance. Several years ago, we proposed a similar risk-based asset allocation investment management concept.

Figure 9. BTR measure

BTR focuses only on downside volatility, and is particularly useful when returns from different products or business units exhibit varying skewness.

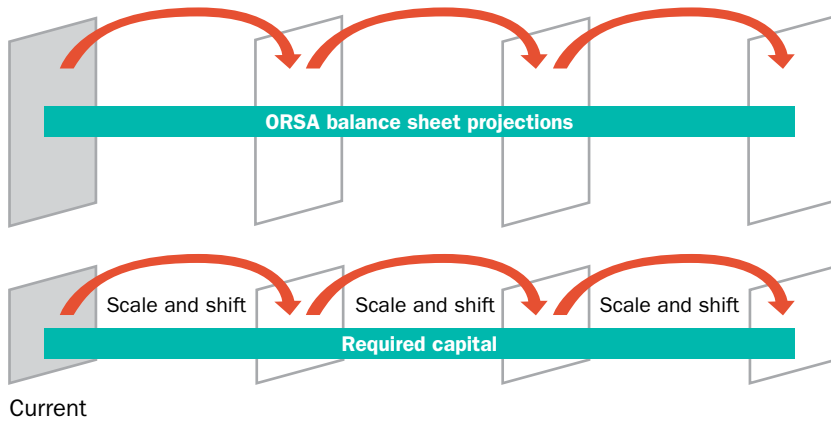
Like standard deviation, BTR is calculated as the square root of the sum of the squares of the distance from the target (minimum acceptable, in this example), considering only outcomes below the target.



Finally, an enterprise risk measurement model like the one described in the case study can also be useful in the context of capital projections, such as those required by the new regulatory ORSA. Once the future balance sheets have been projected according to the business plan, the scale and shift parameters can be used to obtain the required capital at each subsequent point in time (Figure 10). This is a much simpler solution than building a full-blown stochastic projection model at each future balance sheet date.

In this paper, we have attempted to describe the importance, practicality and utility of a near-real-time enterprise risk measurement model. We believe, and the case study demonstrates, that when such models are properly implemented, they can add real value to the enterprise by providing risk-based support to decision makers. Their implementation is also critical to the success of risk appetite frameworks, which we believe will remain hollow without them.

Figure 10. Using scale and shift to estimate future required capital



For More Information

towerswatson.com/riskappetite

Americas

Stephen Lowe
+1 860 843 7057
stephen.lowe@towerswatson.com

Mark Scanlon
+1 212 309 3974
mark.scanlon@towerswatson.com

Martha Winslow
+1 952 842 6527
martha.winslow@towerswatson.com

Asia Pacific

Penny Fosker
+852 2593 4539
penny.fosker@towerswatson.com


Europe

Ian Farr
+44 20 7170 2395
ian.farr@towerswatson.com

Graham Fulcher
+44 1737 284869
graham.fulcher@towerswatson.com

Michael Klueggens
+ 49 221 8000 3433
michael.klueggens@towerswatson.com

John Rowland
+44 20 7170 3853
john.rowland@towerswatson.com



About Towers Watson

Towers Watson is a leading global professional services company that helps organizations improve performance through effective people, risk and financial management. With 14,000 associates around the world, we offer solutions in the areas of benefits, talent management, rewards, and risk and capital management.

Copyright © 2013 Towers Watson. All rights reserved.
TW-NA-2013-34967

towerswatson.com

TOWERS WATSON 