Enterprise Risk Management Quantification—An Opportunity

Christopher (Kip) Bohn, ACAS, MAAA
Brian Kemp, FCAS, MAAA

Presented at
Enterprise Risk Management Symposium
Society of Actuaries

Chicago, IL

April 23–26, 2006

Abstract

Enterprise risk management (ERM) has been getting an increasing amount of attention in recent years. While various industries, regions of the world and professional organizations may have coined different names for their general framework, the underlying theme is the same. Companies and organizations are recognizing the value in assessing, prioritizing and quantifying the risk that they face with the ultimate goal of choosing the most effective mitigation or exploitation options available to them.

With the exception of the insurance and banking industries, much of the focus has been on the qualitative aspects of framework with quantification only briefly touched on. Quantification of enterprise risks often requires developing models that are outside the classic casualty actuarial frequency and severity model realm. However, actuaries' experience and understanding of risk presents tremendous opportunity to expand upon our skill set and both assist and steer the future course of operational and financial risk modeling.

1. A Brief Introduction to the Present State of Enterprise Risk Management

Enterprise risk management (ERM) is a hot topic in today's business environment. Demands from analysts, auditors, regulators and stakeholders in response to legislation and business events helped spark initial interest. Business leaders such as the CFO, treasurer, risk manager and chief risk officer (a recent addition to an increasing number of organizational charts) are recognizing the ERM framework as a vehicle to:

- Provide transparency to analysts, auditors and stakeholders;
- Aid in the development of a financial disclosure framework that will support regulatory compliance initiatives; and
- Promote better capital allocation and decision-making.

A number of professional organizations have also taken an interest and are recognizing the value they can add to the advancement of the ERM framework. A brief sample includes the Casualty Actuarial Society, the Society of Actuaries, the Risk and Insurance Management Society and the Professional Risk Managers' International Association.

A number of ERM frameworks are currently being used. While they may vary in name, industry and region, they share a common theme: the identification, prioritization and quantification of risk in order to help corporations effectively manage their exposure. While many of the frameworks focus on mitigation, exploitation of risk should also be considered. Here is a brief description of three popular frameworks:

- Committee of Sponsoring Organizations of the Treadway Commission (COSO): Enterprise Risk Management—Integrated Framework. Perhaps the most popular framework being implemented in the United States. The definition of ERM offered by COSO is purposely broad and is geared to achieving an entity's risk management objectives in four categories: strategic, operational, reporting and compliance. While discussing various techniques for assessing risk, the methods are more qualitative than quantitative in nature from an actuarial point of view.
- Bank for International Settlements, Basel Committee on Banking Supervision (Basel II), International Convergence of Capital Measurement

and Capital Standards; A Revised Framework. Targeted at banks and financial institutions, the standard is based on three "pillars" which include minimum capital requirements, supervisory review processes and market discipline. The standard also separates risks into three broad categories: credit risk, market risk and operational risk. While banks arguably are comfortable quantifying credit and market risk, operational risk is new territory. As pointed out by Rech,[1] the insurance industry is beginning to run in parallel with this three-pillar approach. Of the three examples given, Basel II puts more emphasis on the quantification of risk and suggests a value at risk (VaR) approach for allocation of capital. However, recent attempts to implement a modeling framework produced results that were more widely dispersed and resulted in lower required capital than expected. This, among other factors, has led to pushing back the implementation date of the accord.

- Standards Australia/Standards New Zealand, Australian/New Zealand Standard: Risk Management (AU/NZS 4360). First introduced in 1995, this is currently one of the more popular frameworks being implemented outside of the United States.[2] Like COSO, this standard provides a generic guide for the establishment and implementation of the risk management process and involves the identification, analysis, evaluation, treatment and monitoring of risks. Quantification is addressed but only broadly.

Given its popularity in the United States, for the purposes of this discussion, we will be focusing on the COSO framework. However, as noted above, most frameworks focus more on the qualitative aspects of ERM. For those that do have more discussion around quantification, there is more work to be done.

2. An Opportunity

Studying the COSO framework from an actuarial and quantitative perspective has led us to the following conclusion: There is a clear opportunity for the actuarial and mathematical communities to not only add value to organizations interested in

[1] James E. Rech. "Enterprise Risk Management for Insurers; Actuarial Theory in Practice," *Contingencies* (November/December, 2005).

[2] Most companies in the United Kingdom have been following the Financial Reporting Council's Internal Control: Guidance for Directors on the Combined Code (Turnbull). Recently, The Association of Insurance and Risk Managers—Risk Management Standard has been gaining popularity in the United Kingdom and Europe. As with the other frameworks noted, our opinion is that quantification guidance is lacking.

implementing an ERM framework but to also aid in the development of a more rigorous quantitative framework.

COSO defines eight key elements to the ERM framework that begins with an understanding of an organization's internal environment, moves on to risk identification and prioritization, touches on the assessment and quantification risk and finally discusses risk response, mitigation and monitoring activities.[3] Of note to those of us with a quantitative background, the 2004 framework doesn't give much guidance on the topic of risk assessment and quantification. In fact, only 8 out of 125 pages are dedicated to the assessment of risk.

In September of 2004, COSO published an application techniques guide. The purpose was to "provide[s] practical illustrations of techniques used at various levels of an organization in applying enterprise risk management principles."[4] This document has more content around the topic of risk assessment (22 out of 112 pages) and discusses both qualitative and quantitative methods. Focusing on the quantitative methods, the guide offers three broad techniques: probabilistic, non-probabilistic and benchmarking techniques. We would argue that from an actuarial point of view, the latter two are really more qualitative in nature while probabilistic techniques are of more interest to our specific skill sets.

The probabilistic techniques discussion touches briefly on "at-risk" models such as VaR, cash flow at risk and earnings at risk. The section says this about modeling risk:

> Certain operational or credit loss distribution estimations use statistical techniques, generally based on non-normal distributions, to calculate maximum losses resulting from operational risks with a given confidence level. These analyses require collection of operational loss data categorized by root cause of the loss, such as criminal activity, human resources, sales practices, unauthorized activity, management process, and technology. Using these loss data and reflecting data on related insurance costs and proceeds, a preliminary loss distribution is developed and then refined to take into account the organization's risk responses.[5]

A clear opportunity is presenting itself to the actuarial community. We can apply our statistical expertise to the determination of non-symmetric probability distributions

---

[3] For additional detail, please refer to The Committee of Sponsoring Organizations of the Treadway Commission (COSO). Enterprise Risk Management—Integrated Framework (2004)

[4] The Committee of Sponsoring Organizations of the Treadway Commission (COSO). Enterprise Risk Management—Application Techniques (2004) 1.

[5] COSO. Application Techniques (2004), 41.

and the creation of stochastic models to determine risk at a given confidence interval. Many of the risks that are being identified through the implementation of an ERM framework do not fall within the classically insurable subset of operational and financial risks, which gives rise to three additional opportunities. First, these risks do not have the luxury of extensive databases of relevant loss and event data. In some cases, the risks identified may have not occurred at all, especially with any significant impact. This presents the opportunity to help develop procedures for the collection and storage of operational and financial loss information that will aid in the quantification of the exposure. Second, understanding your risk will only get you half way. As the ERM acronym implies, you need to *manage* your risk across your organization. There are a number of ways to achieve this: avoidance, mitigation and transfer. Through the quantification of risk, we can aid in the understanding of the cost/benefit tradeoffs of various management strategies. Finally, the ability to quantify risk will also advance the development of new transfer products available in the marketplace.

COSO is evolving under the expectation that organizations such as the Casualty Actuarial Society will step up to the challenge of advancing the overall ERM framework. As stated in Application Techniques, "Over time, we believe that additional guidance will evolve as professional organizations, industry groups, academics, regulators, and others develop material to assist their constituencies." This is the opportunity to add value by applying our actuarial and quantitative expertise in the development of a financial and operational risk modeling framework. This modeling framework should be broad enough to apply not only to COSO but also to all ERM frameworks.

3. A Quantitative Modeling Framework

3.1 Guiding Principles

When developing a model framework, keep in mind the underlying premise of the COSO ERM framework:

> The underlying premise of enterprise risk management is that every entity exists to prove value for its stakeholders. All entities face uncertainty and the challenge for management is to determine how much uncertainty to accept as it strives to grow stakeholder value. Uncertainty presents both risk and opportunity, with the potential to erode or enhance value. Enterprise risk management enables management to effectively deal with uncertainty and associated risk and opportunity, enhancing the capacity to build value.[6]

[6] COSO. Integrated Framework, 1.

Based on this, primary considerations in the development of a modeling framework were not only the quantification of uncertainty surrounding a particular risk or portfolio of risks but also, perhaps more importantly, the ability to assess the cost/benefit tradeoffs of various avoidance, mitigation and transfer options. The result is an iterative six-step process:

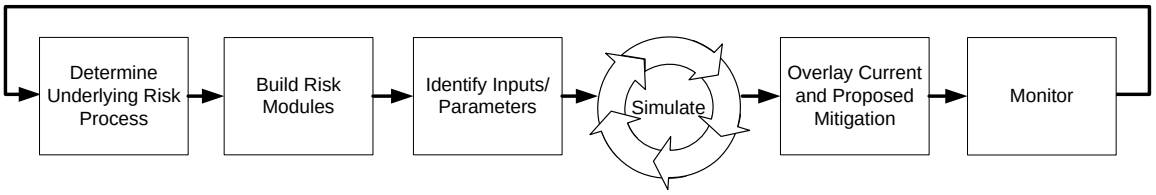

What follows is a discussion of these six steps. For illustration, we have also included excerpts from a case study on a supply chain model we recently developed for a corporate client.

3.2 Determine the Underlying Risk Process

There are a few key considerations that are critical to the first phase of the modeling framework:

Clearly define the risks you wish to model. When defining the risks, it is often helpful to frame the problem into three components: the underlying exposure (which may be as general as the business operations of the company or a specific process or asset); key events that can impact that exposure; and, finally, key consequences that arise from those events. It is important to only focus on significant exposures, events and consequences to the enterprise.

In our case study the underlying exposure was defined as the supply chain for the corporation, from suppliers of raw materials to delivery of the finished products. The supply chain was broken into nine discrete components. Key events were defined as those that could affect an entire location (e.g., natural disaster, fire, etc.) and those that would affect only one process or location. Key consequences were defined as the inability to complete a particular step in the process, which affected each subsequent step. The bottom line consequences were defined as impacts to projected sales (due to lack of product to sell) and potential impacts to future demand (due to loss of market share to competitors).

Determine desired output. Before you begin the design of the model, it is imperative that you have a clear understanding of what outputs or key performance

indicators you wish to track. You should also consider how you wish to measure the risk associated with the key variables. There are a large number of papers on various measures to use: VaR, risk-adjusted return on capital (RAROC), etc., so we will not discuss them here. The key is to understand the risk measures used by the company and design the model accordingly.

For our case study, our output was the deviation to planned sales and cash flow. The company has in place a fairly sophisticated forecasting model that takes into account some of the business risk. As the ERM process has evolved, they have recognized several additional sources of risk, such as the supply chain risk discussed here. The selected model outputs work well with their existing framework.

Keep **in mind potential mitigation strategies that may be implemented.** Risk transfer through insurance is often one of the easiest mitigation strategies to implement in your model (although one should consider the potential for the denial of coverage). The impacts from operational changes, new policies and procedures or perhaps a new manufacturing facility can be more challenging, and it is best to have this in mind when designing the model.

Don't reinvent the wheel. Review any deterministic models that have already been created. In addition to gaining insight into management's current view of the risk process, you will also discover key performance indicators and mitigation strategies that are currently being considered by management.

Map out the risk process. With these considerations in mind, we have found that developing a flowchart of the risk process is beneficial. Historically, when the actuarial community discusses loss modeling, the risk process is based on the determination of a single frequency distribution and a single severity distribution (in many circumstances, a multi-modal severity distribution). While adequate for casualty lines of business where the major mitigation tool is often through insurance products, operational risk modeling often requires a more complex model. It is critical that the design of the risk process be a collaborative effort with those in the organization and the industry that are most familiar with the identified risk.

In our case study example, the risk process essentially followed the manufacturing process. The model captured the dependencies in the process as raw materials were transformed into finished goods. A single frequency and severity model would not be robust enough to adequately model this risk (e.g., outputs from one process are inputs to the next, the physical location of many of the processes are the same and are thus exposed to same loss event). The model accounted for existing risk

mitigation in the form of inventory and excess capacity. Another design consideration was the fact that the model needed to cover a multi-year timeframe which enabled the company to see the change in risk over time. This was important given the lag in implantation of different strategies. During this design phase we worked closely with people responsible for the entire supply chain to ensure our model was a reasonable depiction of their processing and mitigation strategies.

This first stage of the modeling framework does not involve any collection of data or defining of any exposure, event or consequence probability distributions. Rather, it is intended to be the foundational blueprint on which your final quantitative model will be built.

3.3 Build Risk Modules

With the blueprint in hand, you now need to convert the risk process into a stochastic model. The more consideration you gave to the design of your risk process blueprint, the easier the coding of the model will be. There are many software packages available, and your choice will likely depend on your interest in writing actual code as opposed to relying on more familiar spreadsheet applications.

When building or coding your risk model, make sure to take into account the key considerations that were the foundation in the development of the risk process blueprint:

- Identified exposures, events and associated consequences
- Key performance indicators you are interested in tracking
- The ability to overlay various mitigation and transfer strategies.

When building your model, we suggest taking a modular approach that will allow you to easily add or remove exposure, event and consequence modules. You will gain the ability to add consistency to your overall loss model. A simple example of this could be having a single property catastrophe event module that impacts several risk exposure modules.

Finally, consider the incorporation of correlation and causation into your loss modules. While an in-depth discussion is beyond the scope of this paper, it should be noted that many outside of the mathematical community often confuse correlation with causation. At a minimum, you should understand the relationship or potential

relationships between your identified exposures, loss events and consequences. As mentioned earlier, there were certain loss events in our case study that would affect several processes given their physical location. In addition, given that some raw materials are used in multiple products, a loss event for a particular supplier could impact sales and cash flow for multiple products. By creating a single model (built from multiple modules) for the entire supply chain process, we were able to directly establish these relationships rather than having to rely on a correlation matrix.

3.4 Identify Inputs and Parameters

This stage of the framework involves determining the probability distributions and their associated parameters. In many typical actuarial applications, there is a wealth of organization-specific loss and event data. In lieu of this, we would prefer to fall back on industry data. However, as stated earlier, many of the operational risks that organizations are interested in do not fall within those historically underwritten by insurance companies. This would not be cause for concern if organizations had been tracking losses and events associated with these risks. Unfortunately, this is often not the case. A prime example is the banking industry. Before Basel II, banks were not tracking many of the operational risks that they now are accountable for from a capital adequacy standpoint. Many banks have begun to collect this data and a few consulting groups have recognized the opportunity to compile and supply industry event data.

If no organizational or industry event data is available, one alternative is to rely on our experience regarding the general shape of risk distributions. For example, we can be reasonably sure that the distribution around the size of a court decision is not likely normally distributed.

Once we have determined reasonable shapes for the distributions and have incorporated them into our loss modules, we need to determine their parameters. Again, the lack of data may make some in the actuarial community squeamish. However, even in the insurance world, actuaries will sometimes have to rely upon the opinion of experts and judgment in pricing new coverages for example. We argue that this expertise should be used as a starting point when determining parameters in the absence of hard historical data. Querying risk experts within the company about what the specific parameters are for particular probability distributions will not likely get you very far. Rather, consider collecting expert opinions on qualitative statistics such as the average, minimum, maximum foreseeable, most likely events that might be pulled from the distribution in question. This information can be used to select reasonable parameters and validate the distributions selected in your model. This will likely be an

iterative process. What may seem like reasonable assumptions may produce unreasonable results.

In our case study, we concentrated on events that would impact a production location or process for certain specified time periods. It proved easier for our experts to envision scenarios that could lead to three- or six-month shutdowns than to assess the probability of a shutdown of any length. This project also demonstrated the value of the iterative process. After the first set of probabilities was incorporated into the model, we produced some benchmark results and reviewed them with our project sponsors. The magnitude of the losses was higher than expected. After reviewing the model and assumptions, we discovered the issue: although the probabilities felt reasonable in isolation, they did not make sense when aggregated in the model. The second iteration of the model used an aggregate probability of loss to scale the probabilities of losses at individual locations and processes to what was felt to be a more reasonable level.

3.5 Simulate

At this point, you are ready to actually run your model, which is likely made up of a number of modules, and evaluate if the results are reasonable. If you have relied on professional judgment to determine the shape and parameters of various distributions, it is wise to sensitivity test those assumptions. Also, pay close attention to the number of iterations that your model cycles through. Increasing the number of distributions and modules you incorporate into your overall model will increase the number of iterations required. Many simulation packages offer a feature to determine if the aggregate distributions of your key performance indicators are converging. Whether or not your software package includes this feature, it may be worthwhile to sensitivity test the number of iterations you cycle through your model.

Additional value can be derived from your model through scenario analyses. Completely turning off all variability is one option. The result that the model creates should match the expected plan of the organization. Another option could be to manually select events that will impact your modeled exposures and let the consequence distributions vary. This will give you a feel for the potential risk associated with specific events. Finally, scenario testing is another way to test the reasonableness of your model. If you can reach agreement that the results from a specific scenario or set of scenarios make sense, your model gains credibility.

This has been a key step in the validation process in our case study. We were able to walk our expert contacts through specific loss event scenarios showing how the

event in one step impacted subsequent steps and the resulting outputs. By showing them how the model works in a deterministic sense, we increased their comfort level that the aggregate loss distribution produced was reasonable.

3.6 Overlay Current and Proposed Mitigation

At this step, the organization can utilize the model results to make strategic decisions about the amount of risk it wishes to retain, transfer or avoid entirely. If not already built into the model, the current mitigation and transfer strategy should be incorporated to set a baseline. Running various alternative scenarios through the model and comparing the reduction or increase in risk will be of interest to the organization. Considering the additional cost or savings of the alternative mitigation strategies will enable the organization to compare the risk/return tradeoffs of various risk mitigation strategies.

Possibilities for risk exploitation should also be investigated. Through your quantification of risk, you may determine that certain risks do not carry the exposure to loss or variability that was originally perceived. You might also discover that natural hedges exist within the organization. In these instances, relaxing the mitigation strategies and shifting risk management capital to other risks should be considered.

Another benefit of comparing various financing and mitigation strategies is that it may prompt an organization to think more diligently about its appetite for risk. In many cases, the mitigation strategy for a particular risk is based on a qualitative perception of individual risk characteristics. The quantification of risk may alter these perceptions and cause an organization to rethink its appetite for a particular risk. Also, by combining multiple risks, the organization is able to recognize the portfolio effect and may determine that a more aggressive mitigation strategy is warranted.

Finally, additional reasonability and consistency checks can be made at this time. For example, if the move to a less aggressive risk retention strategy does not result in a reasonable reduction in risk, the modeler should revisit the process, distribution and parameter assumptions.

3.7 Monitor

As with the overall ERM framework, risk modeling should not be considered a one-time analysis but rather a continual process implemented within the organization. As time passes, a number of elements of your quantitative model will likely change:

- You will likely refine the risk process. Complexity or additional modules may be added to your model to make it more robust. Conversely, you may determine that certain modules do not add precision or reflect your current view of the risk process and thus a simplification is in order.

- As time passes, the organization will evolve and new risks may be identified as candidates for quantification. Conversely, some risk may diminish and no longer warrant a modeling exercise.

- Probability distributions, parameters and key performance indicators can change over time.

- New mitigation options and insurance products may become available.

- Changes in the business model, competitive landscape or regulatory environment.

It is also important to recognize innovations in technology, computing power and modeling techniques that are sure to present themselves in the future. Indeed, it is our expectation that this basic framework will also evolve and be built upon over time.

3.8 A Word of Caution

George E.P. Box is attributed with the following quote, "All models are wrong, some models are useful." We need to be wary of process, parameter and model risk. It is important to fully disclose assumptions, and simplifying assumptions, that we build into our models. Care should be taken to understand the sensitivity of the model to these items, and to make sure that decision-makers understand them too. What must be kept in mind is that we are not building a model that will take over the decision-making process. We are building a tool to assist in the process. In this case, it helps the decision maker understand what can happen if things go wrong and help value options for reducing, eliminating or transferring that risk.

4. Conclusion

The goal of ERM is to identify and manage risk, align risk appetite with the organization's strategy, enhance risk response decisions and improve the deployment of capital. Most, if not all, ERM frameworks recognize the added value that the quantification of risk brings to the overall process. The need for the development of a quantification framework is a clear opportunity for the actuarial and mathematical communities. Through the development of this framework we will be presented with additional opportunities. Aiding organization in the understanding of risk, the quantification of the cost/benefit tradeoffs of various mitigation strategies, the development of new databases of loss and event data and the development of new transfer products are just a few.

Currently the focus of the actuaries has been in the insurance and banking world with respect to ERM. ERM initiatives are much broader than that. There is an opportunity for actuaries to apply a skill set already developed to this broader world as well.

Risk modeling cannot evolve without venturing into new territories. In the 19th and mid-20th century, weather modeling accuracy was a fraction of what it is today. But, without the development of that initial foundation, we would not have the sophisticated models that we rely on today. Finally, we leave you with this food for thought:

> The world is moving into a new age of numbers. Partnerships between mathematicians and computer scientists are bulling into whole new domains of business and imposing the efficiencies of math. This has happened before. In past decades, the marriage of higher math and computer modeling transformed science and engineering. Quants turned finance upside down a generation ago. And data miners plucked useful nuggets from vast consumer and business databases. But just look at where the mathematicians are now. They're helping to map out advertising campaigns, they're changing the nature of research in newsrooms and in biology labs, and they're enabling marketers to forge new one-on-one relationships with customers. As this occurs, more of the economy falls into the realm of numbers. Says James R. Schatz, chief of the mathematics research group at the National Security Agency: "There has never been a better time to be a mathematician."[7]

[7] Stephen Baker with Bremen Leak. "Math Will Rock Your World," in BusinessWeek Online, http://www.businessweek.com/print/magazine/content/06_04/b3968001.htm?chan=gl (23 January 2006)