

Applying Actuarial Techniques in Operational Risk Modeling

Donald Mango, FCAS, MAAA¹

Presented at
Enterprise Risk Management Symposium
Society of Actuaries

Chicago, IL

April 23–26, 2006

Copyright 2006 by the Society of Actuaries.

All rights reserved by the Society of Actuaries. Permission is granted to make brief excerpts for a published review. Permission is also granted to make limited numbers of copies of items in this monograph for personal, internal, classroom or other instructional use, on condition that the foregoing copyright notice is used so as to give reasonable notice of the Society's copyright. This consent for free limited copying without prior consent of the Society does not extend to making copies for general distribution, for advertising or promotional purposes, for inclusion in new collective works or for resale.

¹ E-mail: donald.f.mango@guycarp.com. Contact: Guy Carpenter & Co., 44 Whippany Rd, Morristown, NJ, 07962, USA. Phone: (973) 285-7941.

Abstract

Motivation. There is a growing need for effective, practical methods of operational risk analysis in all industries. Risk professionals are learning to develop business-unit-level risk distributions, combine those distributions into an aggregate risk model, and use that aggregate risk model to either assign risk charges back to the business units, or to evaluate cost-benefit of mitigation strategies. Operational risk modeling is structurally similar to actuarial risk modeling. The operational risk community will benefit from learning actuarial techniques that can be applied to operational risk modeling.

Method. First, the paper will outline how operational risk management is similar to an internal insurance program. Second, it will discuss an internal risk modeling framework that ties together risk exposure, likelihood, severity, and correlation into an aggregate loss model. Finally, using the model output, it will present several methods to transparently reward risk mitigation efforts and attribute risk costs to their sources.

Conclusions. This paper will demonstrate the potential synergies from applying actuarial techniques to operational risk analysis. It will also demonstrate practical techniques, grounded in actuarial science, to solve operational risk management problems such as risk cost allocation and risk mitigation cost-benefit analysis.

1. Introduction

1.1 Research Context

The Basel Committee defines *operational risk* as “the risk of loss resulting from inadequate or failed internal processes, people, and systems or from external events. This definition includes legal risk, but excludes strategic and reputational risk.”²

Commercial property and liability insurance is therefore essentially the coverage of “insured operational losses.” The same risk analytic techniques—**actuarial** techniques—can be applied to the analysis and management of all operational risks. Actuarial science has more than 100 years of lessons learned and best practices to share. Actuaries are experts in internal risk modeling and probabilistic decision-making, both core components of any sophisticated operational risk management system.

1.2 Objective

First, the paper will outline how operational risk management is similar to an internal insurance program. Second, it will discuss an internal risk modeling framework that ties together risk exposure, likelihood, severity and correlation into an aggregate loss model. Finally, using the model output, it will present several methods to transparently reward risk mitigation efforts and attribute risk costs to their sources.

1.3 Outline

The remainder of the paper proceeds as follows. Section 2 will provide a review of principles of actuarial science, operational risk and draw the parallels with insurance. Section 3 will delve into corporate decision-making under uncertainty. Section 4 discusses probabilistic decision metrics. Section 5 shows methods for assessing the cost of risk and the benefits of mitigation to business units. Section 6 concludes with areas for further research.

² Basel Committee on Banking Supervision, “International Convergence of Capital Measurement and Capital Standards.”

2. Background

2.1 Principles of Actuarial Science

A good overview of actuarial science can be found in the 1999 exposure draft “Principles Underlying Actuarial Science,” published jointly by the Casualty Actuarial Society and Society of Actuaries.³ It succinctly summarizes the field of actuarial science as follows:

Actuarial science is an applied science based on concepts and observations distilled from the experience of practitioners and from other sciences....The primary applications of actuarial science identify and analyze consequences of events involving risk and uncertainty. Actuarial practice involves the analysis and management of these implications and their associated costs. To gain insights about future possibilities, the actuary depends on observation and the wisdom gained through prior experience. Actuaries use these observations and this experience to construct, validate, and apply models. Actuaries continually incorporate additional observations and insights into their models. This feedback cycle systematically addresses discrepancies between these models and observed reality.... Actuarial models are constructed to aid in the assessment of the consequences associated with phenomena that are subject to uncertainty with respect to occurrence, timing, or severity. To construct such models requires:

- understanding the conditions and processes under which past observations were obtained
- anticipating changes in those conditions that will affect future experience
- evaluating the quality of the available data
- bringing judgment to bear on the modeling process
- validating the work as it progresses and revising the model as needed
- estimating the uncertainty inherent in the selection and construction of the model.

³ “Exposure Draft of Principles Underlying Actuarial Science,” *Joint CAS/SOA Committee on Principles*, 1999. Available online at www.casact.org/research/science.pdf.

Actuarial models are typically used in property and casualty insurance, life insurance, investments, health benefit systems and retirement systems. These are the “traditional practice areas” for actuaries. However, actuaries are essentially **risk analytics professionals** whose skills and techniques have applicability beyond the traditional practice areas. One area where the applications are obvious, and the potential benefits immediate, is operational risk.

2.2 Operational Risk Defined

Operational risk is defined by the Basel Committee on Banking Supervision as “the risk of loss resulting from inadequate or failed internal processes, people, and systems or from external events. This definition includes legal risk, but excludes strategic and reputational risk.”⁴ This definition is deceptively short for such a broad area. To elaborate, the Basel Committee issued a July 2002 consultative paper, “Sound Practices for the Management and Supervision of Operational Risk,” where they defined the following types of operational risk events⁵:

- **Internal fraud.** Acts of a type intended to defraud, misappropriate property or circumvent regulations, the law or company policy, excluding diversity/discrimination events, which involve at least one internal party. Examples include intentional misreporting of positions, employee theft, and insider trading on an employee’s own account.
- **External fraud.** Acts by a third party, of a type intended to defraud, misappropriate property or circumvent the law. Examples include robbery, forgery, cheque kiting, and damage from computer hacking.
- **Employment practices and workplace safety.** Acts inconsistent with employment, health or safety laws or agreements, or which result in payment of personal injury claims, or claims relating to diversity/discrimination issues. Examples include workers compensation claims, violation of employee health and safety rules, organised labour activities, discrimination claims, and general liability (for example, a customer slipping and falling at a branch office).
- **Clients, products and business practices.** Unintentional or negligent failure to meet a professional obligation to specific clients (including fiduciary and suitability requirements), or from the nature or design of a

⁴ Basel Committee on Banking Supervision, “International Convergence of Capital Measurement and Capital Standards—A Revised Framework,” November 2005, Part 2, Section V-A. Available online at <http://www.bis.org/publ/bcbs118.pdf>.

⁵ Available online at <http://www.bis.org/publ/bcbs91.pdf>.

product. Examples include fiduciary breaches, misuse of confidential customer information, improper trading activities on the bank's account, money laundering, and sale of unauthorised products.

- **Damage to physical assets**. Loss or damage to physical assets from natural disaster or other events. Examples include terrorism, vandalism, earthquakes, fires and floods.
- **Business disruption and system failures**. Disruption of business or system failures. Examples include hardware and software failures, telecommunication problems, and utility outages.
- **Execution, delivery and process management**. Failed transaction processing or process management, and relations with trade counterparties and vendors. Examples include data entry errors, collateral management failures, incomplete legal documentation, unapproved access given to client accounts, non-client counterparty misperformance, and vendor disputes. (p. 2-3)

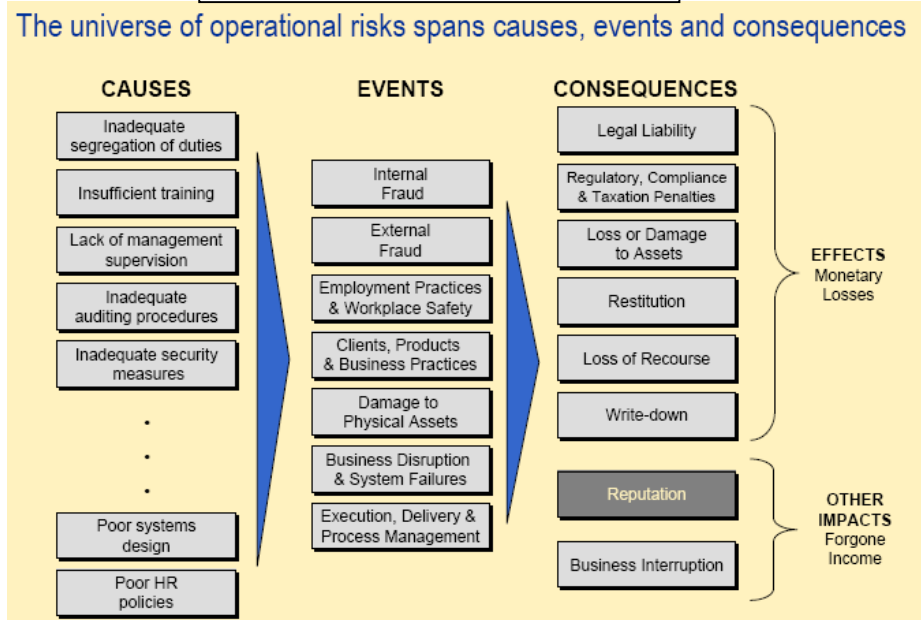
While heavily banking-focused, the Basel committee definition has gained substantial visibility and acceptance as the field of ERM has developed over the last 10 years.

2.3 Parallels with Insurance

Property and casualty (P&C) insurance is used by all type of firms to transfer the impact of several of the risks from the Basel Committee list. Examples cited include fire, earthquake, theft, workers' compensation, employment practices liability and premises liability for customer safety. P&C actuaries have been analyzing the pricing, valuation and portfolio risk management of the costs of these risk transfers for nearly 100 years. One could even refer to P&C actuaries as "insured operational risk analysts."

Actuaries have developed their own approach to this analysis, but so have the leaders in operational risk analysis and modeling. One such expert, Ali Samad-Khan (2005: "Actuarial Techniques in Banking Operational Risk"), separates the operational risk process into *causes, events and consequences*. See Figure 1:

Figure 1 – from Ali Samad-Khan



A P&C actuary would change terminology as follows:

Samad-Khan	P&C Actuarial
Causes	Perils + Exposures
Events	Occurrences or Claims
Consequences	Losses

2.3.1 Causes vs. Perils and Exposures

Actuaries consider causation as the combination of two factors: perils and exposures. *Perils* are root causes of damage. Examples include fire, flood, wind, earthquake, hail, burglary or automobile collisions. *Exposures* are insured interests (i.e., a property, corporate legal entity or activity) that could potentially be damaged or diminished in value due to impact from a peril. Insurers take on (or *assume*) exposure when they issue (or *write*) insurance policies.

2.3.2 Events vs. Occurrences or Claims

Actuaries distinguish between occurrences and claims. *Occurrences* are events resulting in damage. Examples include a collision involving an automobile, a property fire, or an earthquake. Occurrences are uniquely identifiable with respect to location and duration.

Claims are attempts by policyholders to collect compensation under the terms of their insurance policy for damages. The distinction is crucial to insurers, because not every occurrence results in a valid claim. Some types of occurrence are *covered*, where coverage could be defined by specifying covered perils, covered properties, persons or entities. Uncovered occurrences violate one or more of these policy terms and conditions, and therefore will not be honored if submitted as claims.

2.3.3 Consequences vs. Losses

Once the determination of legitimate coverage has been established, the final damage amount needs to be determined by expert review in a process known as *claim adjustment*. The final amount is subject to negotiation, limitation, court judgment and appeal, and so on. These time lags in resolving claims to final value force actuaries to apply forecasting techniques known as *loss development*, in an attempt to extrapolate these preliminary estimates to their final or *ultimate* amount.

Most operational risk losses will not have this much uncertainty and legal superstructure surrounding their valuation process. The “claimant” and “insurer” are one and the same, namely the entity suffering the loss. However, some operational risks involve lawsuits and disputes (with counterparties, clients, or competitors), or regulatory actions and potential fines. These types of losses will have a time delay similar to some types of insurance claims. These types of losses may need to be “developed” to ultimate using extrapolation techniques similar to actuarial loss development.

2.3.4 Exposure Bases

Exposure bases measure the magnitude of the potential risk assumed by an insurer. In the seminal actuarial reference on the topic, Bouska (1989) highlights three criteria for acceptability of an exposure base:

1. It should be an accurate measure of the exposure to loss.
2. It should be easy for the insurer to determine.
3. It should be difficult for the insured to manipulate.

Exposure bases are proxies for magnitude of exposure to perils. Some exposure bases are obvious—e.g., number of automobiles, or insured value of a home. Others are pragmatic, being the best *available* proxies—e.g., payroll for workers’ compensation or sales for product liability. In these instances, satisfaction of criterion 1 is outweighed by practicalities of criteria 2 and 3.

One should strive to meet these criteria, as well as possible, in an operational risk management system. Failure to do so could damage the internal credibility of the risk management function, leading business units to question the fairness or accuracy of the operational risk measurement.

2.3.5 Classification Plans

Fairness and accuracy of risk assessment are core actuarial principles, part of what is known as *actuarial equity*. Actuarial equity refers not to an investment in the insurer, but to the goal that every policyholder pay a premium proportional to his or her particular loss potential, as measured by various rating characteristics. These rating characteristics are assembled into what are known as *classification plans* (or class plans for short), for purposes of determining the premium to be charged. These plans categorize potential customers based on the values of the rating variables. Examples from automobile insurance include driver gender, driver age, type of vehicle, territory (where the car is garaged and driven) and type of usage (business or pleasure).

In a competitive insurance market, insurers are constantly refining their class plans. If Insurer X were the only one failing to differentiate in its rating (i.e., by charging uniform premiums), they would be exposed to what is known as *adverse selection*. The result: high-risk individuals would find X's uniform premiums lower than those quoted by other insurers, and low-risk individuals would find the opposite. The end result for X would be an imbalanced portfolio of high-risk individuals written at medium-risk prices.

In an operational risk management setting, what is important is identifying (i) the key risk indicators to use as rating variables, and (ii) the discounts or surcharges based on their value.

2.4 Operational Risk Management as Internal Insurance

The structural parallels between operational risk and P&C insurance are compelling. The necessary elements of a sound operational risk management system are akin to components of an internal insurance program. This suggests that a centralized operational risk management group is a de facto *portfolio risk* manager. The portfolio consists of all the retained operational risk exposures. Drawing on the key elements of insurance portfolio risk management, the necessary steps for operational risk portfolio management might include:

1. Identify exposure bases for each key operational risk. Examples include payroll, headcount, transaction count, revenue, number of counterparties and notional outstanding.
2. Measure the exposure level for each business unit (BU) for each of those operational risks.
3. Estimate the loss potential (frequency and severity) per unit of exposure for each operational risk, reflecting the existing level of internal controls and process effectiveness.
4. Combine #2 and #3 to produce modeled BU loss frequency and severity distributions.
5. Estimate the impact (reduction) of mitigation or improvement efforts on the BU loss frequency and/or severity distributions.
6. Establish a risk transfer program between each BU and the corporate risk function. This involves setting a BU *retention level*, an expected amount of loss the BU can bear on its own. Amounts in excess of the retention (i.e., unexpected losses) are transferred to the corporate risk function.
7. Determine and assess the cost of this risk transfer.

Actuarial techniques could prove beneficial in the following areas:

- Steps 1-4 will benefit from the actuarial risk analytic *framework* and techniques, more than from specific insurance-specific information (e.g., risk distribution parameters). Identifying the best exposure bases, frequency and severity models for each operational risk is unfortunately outside the scope of this paper. Actuaries can provide tools and techniques, but much research remains to be done.
- Step 5 will likely be an area where significant expert opinion is required. It is unlikely that organizations will have significant amounts of operational loss data both *before and after* every possible mitigation effort. Actuaries are often in a comparable position. A good example is the mandated premium discounts for tort reform. Actuaries do not have the luxury of “re-running the experiment” before and after tort reform—that is, what would the claim experience have been had these reformed tort laws been in effect over the past several years. Significant detailed analysis of claim and policyholder information is required, along with informed estimation and judgment.

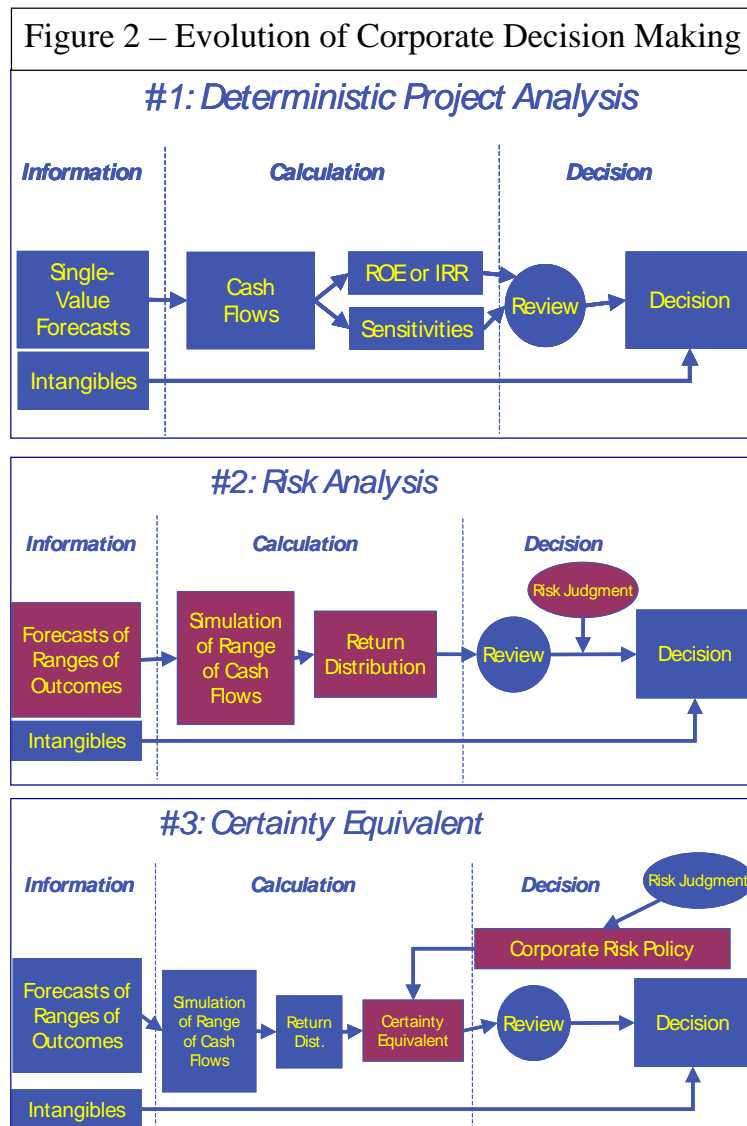
- Steps 6 and 7 are most directly comparable to what insurers face. All firms must assess the cost of risk within a portfolio. The set of operational risks facing an organization can be thought of, and managed, as a portfolio of risks.

The remainder of the paper will focus on elaborating on Steps 6 and 7. That elaboration begins by grounding this exercise in the overarching framework of corporate decision-making under uncertainty.

3. Corporate Decision-Making under Uncertainty

3.1 Evolution of Corporate Decision-Making under Uncertainty

Internal risk modeling is becoming common practice in risk management of all types of firms, but particularly financial institutions. It represents a special case of simulation (or *Monte Carlo*) modeling, a widely used methodology in all forms of corporate decision-making under uncertainty. Spetzler (1968) outlines the evolution of this decision-making (see Figure 2):



1. Deterministic Project Analysis—using a single deterministic forecast for project cash flows, a decision variable such as present value or internal rate of return is produced. Sensitivities to critical variables may be shown. Uncertainty (along with other intangibles) is handled judgmentally (i.e., intuitively) by decision makers.
2. Risk Analysis—forecasts of distributions of critical variables are fed into a Monte Carlo simulation engine to produce a distribution of present value of cash flows. Risk judgment is still applied intuitively.
3. Certainty Equivalent—an extension of Risk Analysis, quantifies the intuitive risk judgment by means of a corporate risk preference or utility function. The utility function does not replace judgment, but simply formalizes the judgment so it can be consistently applied.

Best practice in actuarial risk modeling has evolved to Step 2, the Risk Analysis stage, in what the actuaries call dynamic financial analysis (DFA). However, there is not yet consensus that the jump from Step 2 to Step 3 is proper or even meaningful. The debate centers on the role of corporate risk preference⁶ in an efficient market/modern portfolio theory world.⁷ An attempt to summarize the debate might go:

- Diversified investors, with many small holdings of all available securities (the *market portfolio*), are only concerned with non-diversifiable (a.k.a. *systematic*) risk.
- Diversifiable (a.k.a. *firm-specific*) risks do not command any risk premium in the market, since they can be diversified away by holding the market portfolio.
- Investors require a risk premium (additional return above risk-free) as compensation for bearing systematic risk.
- Firm managers should focus on maximizing shareholder value.
- Since their shareholders can diversify away firm-specific risk, they are indifferent (risk-neutral) towards it.
- Therefore, firm managers ought to be indifferent to firm-specific risk as well.

⁶ See Walls (1995) for discussions of corporate risk attitudes and preferences.

⁷ There are numerous references on efficient markets and modern portfolio theory—for example, Bodie and Merton (2002).

Within such a theoretical framework, there is no apparent place for the Step 3 certainty equivalent approach, with its mathematical formulation of corporate risk preferences. Managers should only care about the wealth, risk preferences, and other investment opportunities of the firm's owners. Absent a real-time survey system, the best proxy for this information is the record of stock prices—both their own stock and other comparable stocks. Firm managers and equity analysts perform these (admittedly complex) analyses in an attempt to discern the probable impact of major firm decisions on the stock price.

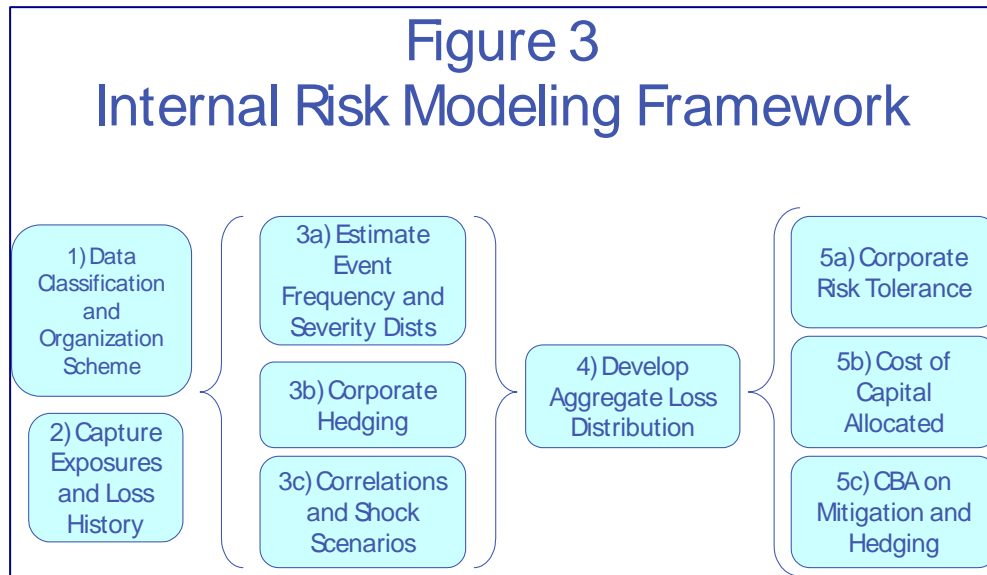
While this theory is appealing and has many advocates, it is short on practical advice for use in operational risk management:

- Those managing the firm have no way to identify which of the risks they face are firm-specific, and which are systematic. Finance theory offers no assistance beyond the aforementioned attempts to discern which decisions produced a noticeable impact on stock price, and which did not.
- One of the more common market-based risk signals—the risk-adjusted discount rate—only reflects risk if there is a time lag. For many operational risks, the time aspect is unimportant—the risk is essentially instantaneous (what is called a *jump risk*).
- Market-based risk signals often lack the refinement and discriminatory power that managers need to make cost-benefit and tradeoff decisions for mitigation efforts.

Neither approach appears sufficient on its own. There may be a perspective from which both approaches are seen as complementary parts of a single whole. Shareholders want market value maximized. Market value consists of *book* value—the recorded value of held assets—plus *franchise* value—the present value of future earnings growth. Risk management aims to both facilitate future earnings growth and *prevent loss of future earnings growth*—that is, to protect franchise value. So it appears shareholders and managers may be aligned in this regard: both want risk management in place to protect the franchise value. Therefore, both camps should support a corporate risk policy to make risk management decisions in a more consistent, transparent manner. This provides some degree of support for the evolution to Step 3.

3.2 Internal Risk Modeling (IRM) Framework

We will now outline an Internal Risk Modeling (IRM) framework well-suited for operational risk. Figure 3 shows the basic components of the IRM framework. Each component will be discussed in detail below.



(1) Data Classification and Organization Scheme

This is the collection framework for categorizing prior loss history for purposes of determining (event) frequency and (event) severity distributions. Care must be taken to capture all the necessary data elements to allow adjustments for differences in conditions. Typical data elements include exposure base, cohort (for example, by year), event cause code, etc. Much additional research is needed to determine the best schema for collecting operational risk information.

(2) Capture Exposures and Loss History

Once the collection scheme is ready, the database is populated with exposure and event loss information. This database will be the basis for the analyses in Step 3.

(3a) Estimate Event Frequency and Severity Distributions

This is arguably the most difficult step in the entire process. Frequency distribution represents the likelihood of various numbers of events per

exposure unit, per time period. Expressing frequency per exposure unit facilitates scaling for different BU size. Severity distribution is the range of possible loss amount outcomes per event. Statistical distribution fitting (known among actuaries as *curve-fitting*) is a complicated exercise with many potential pitfalls and issues. A thorough treatment is beyond the scope of this paper. Interested readers should consider *Loss Models: From Data to Decisions*, by Klugman, Panjer and Willmot,⁸ an industry-standard actuarial text.

(3b) Corporate Hedging and Mitigation

It is important to explicitly reflect the benefits of any corporate-level hedging or mitigation. Examples include insurance, derivatives, process improvement and control plans.

(3c) Correlations and Shock Scenarios

This is another difficult but important step. Any organization is exposed to potential shock scenarios, either by multiple locations suffering events from a common cause, or a single location suffering events from multiple causes. Correlations are admittedly difficult to estimate from actual loss experience. Some degree of correlation can be inferred from assessing structural linkages between risk sources.

Shock or stress scenarios impact company viability. In some cases, shock scenarios are defined and promulgated by an external authority (e.g., regulator or rating agency). Alternatively, a firm may define its own shock scenarios as part of disaster planning.

(4) Develop Aggregate Loss Distribution

An aggregate distribution is the distribution of a sum of several random variables, also known as a *convolution*. Methods for generating an aggregate distribution are well known in statistical modeling. Samad-Khan (2005: "Why COSO is Flawed") advocates an actuarial-type approach, including convolution, as operational risk modeling best practice. An excellent reference on aggregate distributions is Wang (1998).

⁸ *Loss Models: From Data To Decisions, 2nd Edition*, Stuart A. Klugman, Harry H. Panjer, Gordon E. Willmot, 2004, Wiley. See www.wiley.com/WileyCDA/WileyTitle/productCd-0471215775.html.

(5a) Corporate Risk Tolerance

We seek a mechanism that:

1. Takes an aggregate loss distribution, with many sources of risk;
2. Assesses (quantifies) the impact of the possible aggregate loss outcomes on the corporation;
3. Assigns a cost to each amount of impact; and
4. Attributes the cost back to the risk sources.

Corporate risk tolerance is used in steps 2 and 3. The impact (effect) of a certain loss amount depends on the organization's size, financial resources, ability and willingness to tolerate volatility. All of these combined represent the "risk tolerance." The translation of impact into cost requires a risk preference function of some form, either implicitly taken from an outside source (e.g., the capital markets), or explicitly derived based on firm management attitudes. The key point: any selected method implies some type of translation from impact to cost. The questions facing practitioners: should the choice be yours or someone else's (by default), and should it be implicit or explicit?

If a firm would like to develop its own explicit risk preferences, there are methods available. Spetzler (1968) showed how to identify the parameters of a *normative*⁹ corporate risk policy using a series of experiments involving indifference. One of the benefits of such a process is the transparent, objective, mathematical expression of the corporation's acceptable risk-reward tradeoffs. Such a function can improve cost-benefit analyses by quantifying, for example, the minimum decrease in risk (measured by any number of possible metrics) sufficient to justify a certain mitigation cost. Without such a function, cost-benefit analysis (CBA) decisions will still be made; however, the criteria will be inconsistent and opaque, driven to a large part by the decision-makers' intuitions.

Walls (2004) demonstrates this in the context of energy exploration and production (E&P). He applies modern portfolio concepts to a set of available E&P opportunities. He identifies an efficient frontier of *possible* portfolios (i.e., those minimizing risk for a given return, subject to a constraint on total investment funds), and plots risk (measured as standard deviation of net present value (NPV)) against

⁹ Decision analysis differentiates among three possible purposes for a corporate utility function: *descriptive*, describing the subject's current risk attitude; *predictive*, predicting the subject's future behavior in risk situations; and *normative*, as a tool for improving future decisions involving risk. See Spetzler 1968.

return. The firm's existing portfolio is sub-optimal, measured either in terms of risk (i.e., there are efficient frontier portfolios with the same return at lower risk), or return (i.e., there are efficient frontier portfolios with higher return at the same risk). So far, his approach is consistent with "standard" modern portfolio theory.

Walls' innovation is the notion that in order to *select* one of the efficient frontier portfolios, the firm must be able to answer the following questions:

- How much risk (standard deviation) are we willing to tolerate?
- How much reward are we willing to give up for a given reduction in risk, and vice versa?
- Are the risk-reward tradeoffs *available along the efficient frontier* acceptable to us?

The first question requires an answer to the firm's *risk tolerance*. The answer will indicate the riskiest portfolio choice that is tolerable. The second question requires the firm to express its *risk preferences*. The answer will allow the firm to select among available returns for given risk levels. The third question can then be answered.¹⁰ Walls shows how by using a utility function, "decision makers can incorporate their firm's financial risk propensity into their choices among alternative portfolios" [2004, p. 61]. Of particular importance is the idea that modern portfolio theory presents a set of choices among *possible portfolios*. Which of those portfolios to choose is still a firm-specific decision, requiring explicit expression of firm risk preferences.

(5b) Cost of Capital Allocated

Capital allocation is among the most significant open questions in actuarial science. A thorough treatment is well beyond the scope of this paper. Some references include Mango (2003), Meyers, Klinker and Lalonde (2003), Myers and Read (2001), Venter (2003) and Merton and Perold (1993).

There appears to be general agreement that what is really allocated is the *cost of risk capital* as opposed to capital itself. This is in contrast to the capital allocations that occur in manufacturers. Those allocations involve actual cash transfers of retained earnings, and the investment of that capital in the operation of the business—in salaries, materials, power, marketing, etc. Risk capital allocations, on the other hand, are

¹⁰ If the *market price of risk* is currently too low, the firm could opt to hold off and not tie up any capital at that time.

completely theoretical. When an insurer writes an automobile policy, no risk capital is transferred to that policy.

Perhaps by an unfortunate turn of history, the same manufacturing capital language was used to describe risk capital. This reinforced the adoption of capital decision analysis techniques based on manufacturing analogies, such as internal rate of return.

Risk capital for financial intermediaries provides a buffer that secures a certain counterparty status for the *firm in total*. In this regard, risk capital is a measure of the firm's total risk-bearing capacity. Because of the portfolio phenomena of diversification and accumulation, however, this capacity is solely an aggregate measure, only having meaning for the portfolio in total. In the case of a generic firm, the "portfolio" is the accumulation of risk exposures from all sources—operational, credit, market, etc. Because the elements of portfolios may have non-linear interdependence, the impact of any one element on the portfolio can also be non-linear. This makes a linear, proportional allocation of a total amount back to individual elements quite complicated. There are many possible ways this problem can be solved, each having its own compromises and limitations. Some examples of current research include Kalkbrener (2005), Venter, Major and Kreps (2005) and Mango (2005).

Many researchers continue to allocate risk capital, but use it as an interim step in assigning the cost of that risk capital to portfolio elements. The cost is the product of a risk-adjusted capital amount and a hurdle rate. Since the capital is risk-adjusted, this goes by the acronym RORAC (return on risk-adjusted capital).¹¹

Merton and Perold (1993) tried to address this by taking a fundamentally different tack altogether. They define risk capital for a financial intermediary as the amount needed to guarantee the performance of that intermediary's contractual obligations at the *default-free* level. Under this framework, the notion of return on risk capital has no meaning, since they have bypassed the two-step RORAC process and leapt straight to a cost. Mango (2005) extended Merton and Perold to insurance, further clarifying the true nature of insurance capital usage. The fundamental innovation is the recognition and treatment of the *entire* pool of risk capital as a shared asset or common pool resource.

¹¹ Many authors are imprecise in their use of the acronyms RAROC and RORAC interchangeably. RAROC, risk-adjusted return on capital, is a method of arriving at a cost of risk capital by using simplistic (non risk-adjusted) capital amounts, and risk-adjusting the hurdle rates. See Mango (2005).

Given *cost* of risk capital, with no allocated capital amount, a good candidate decision variable is *economic value added* (EVA).¹² The formula for EVA is:

$$\text{EVA} = \text{NPV Return} - \text{Cost of Capital}$$

EVA is typically expressed as an amount. An activity with a positive EVA is said to “add value,” while one with a negative EVA “destroys value.” The examples below will use EVA as a core decision variable.

(5c) Cost-Benefit Analysis (CBA) for Mitigation

Once the corporate risk tolerance and capital cost allocation are completed, the CBA is straightforward. For example, under the EVA approach, any mitigation effort resulting in a positive incremental EVA is worth doing. Under capital allocation, any mitigation project where benefit (reduced capital cost) exceeds costs should be undertaken.

Now that the framework has been laid out, the next section will discuss the decision of selecting a probabilistic decision metric.

4. Choices of Probabilistic Decision Metric

4.1 Probabilistic Decision-Making (PDM)

Risk management in an organization with multiple business units and risk sources involves analysis of a complex, multi-dimensional space. Effective decision-making in that space will be challenging based solely on a single risk metric. Statistical theory would recommend a *suite* of decision metrics that are (as much as possible) distinct and independent, reflecting different dimensions of the space, and responsive to different dynamics. However, corporate realities often mandate a compromise be struck in the interest of parsimony. Corporations need to make decisions, and effective decision-making often simplifies complex situations as much as possible, *but no more so*.

This section will focus on two major classes of risk metrics, tail-focused and centered, each with its own advantages and disadvantages.

¹² VA is a registered trademark of Stern Stewart & Co. (www.sternstewart.com).

4.2 Tail-Focused Risk Metrics

These metrics are fairly standard in introductory probability and statistics.

- Percentile (also known as value-at-risk or VaR) = outcome associated with a certain cumulative probability. Some commonly used values include the 50th, 80th, 90th, 95th, and 99th. Under the actuarial convention, a higher percentile represents a worse outcome (i.e., larger loss). The capital market convention is the opposite. Typical capital market analyses focus on extreme tail percentages (e.g., 99.9th percentile or 1-in-1,000 year event). It may be unwise to rely too heavily on such extreme indications for critical operational decisions. Those very severe scenarios are sensitive to the number of simulation trials, data scarcity in the distribution curve fitting process, scarcity of extreme scenarios for calibration and largely judgmental correlation assumptions. (See Section 5.1 below.)
- Tail-value-at-risk (TVaR) = conditional expected value beyond a certain VaR.
- Excess-TVaR (XTVaR) = TVaR – Mean.

One desirable characteristic shared by both TVaR and XTVaR is *additivity*. For example, the sum of the TVaRs for the components of a portfolio equals the TVaR of the entire portfolio. See [18] for more details.

Many economic capital approaches use tail-focused metrics—for example, TailVaR at the 99.9th percentile. These extreme outcomes are grounded in reality, often being selected to correspond to the probability of default for a target credit rating.

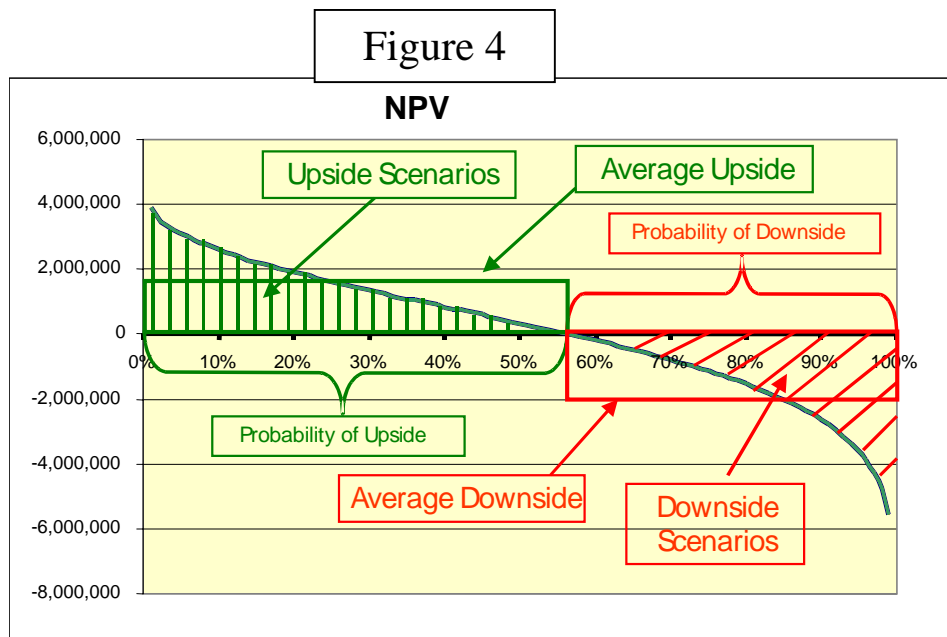
However, one concern with these approaches is their reliance on extreme outcomes from internal risk models. These worst-case scenarios are typically driven by shock scenarios or high correlations. These scenarios are typically the most difficult to estimate from data, due to small sample size and scarcity of data. The selected event frequency and severity parameters necessarily include a large degree of expert judgment. Once business unit advocates realize that this expert judgment is flowing through the IRM to their risk cost allocation, driving capital charges and impacting the performance assessment of their product line, there may be an organizational backlash.

4.3 Centered Risk Metrics

We introduce four metrics based on a decomposition of the mean into two sub-components: an upside and a downside.

- Upside (U) = conditional expected value of positive outcomes.
- Downside (D) = conditional expected value of negative outcomes.
- Probability of upside [P(U)] = probability of positive outcomes.
- Probability of downside [P(D)] = probability of negative outcomes.

Using expected value as a decision metric reduces the information available in the distribution to a single value. These metrics reduce the full distribution to two outcomes and their associated probabilities. While admittedly simplistic, this approach has appeal in facilitating organizational learning and acceptance of PDM. Business unit leadership may be more comfortable transitioning from expected values to a two-valued approximation (note that $\text{mean} = U \cdot P[U] + D \cdot P[D]$), rather than a full probability distribution. Upside and downside provide a stepping-stone for the organization to gradually introduce the notion that key values are not fixed estimates but distributions. Figure 4 demonstrates these concepts pictorially.



Two additional metrics can be derived from D, U, P(D) and P(U).

- D/U ratio = ratio of downside to upside. This is a good measure of the “leverage” of an outcome distribution.
- Reward-to-risk ratio (R2R) = $[U * P(U)] / [D * P(D)]$. This metric was first introduced in the actuarial literature (in a slightly different form) by Ruhm (2001), who referred to it as the “risk coverage ratio” or (RCR).¹³ This metric has high information content; it is a composite risk and reward measure, combining the leverage aspects of the D/U ratio with the associated probabilities. Figure 5 presents the meaning graphically. R2R is particularly good as an indifference measure, for use in CBA (i.e., make decisions such that the R2R stays constant).

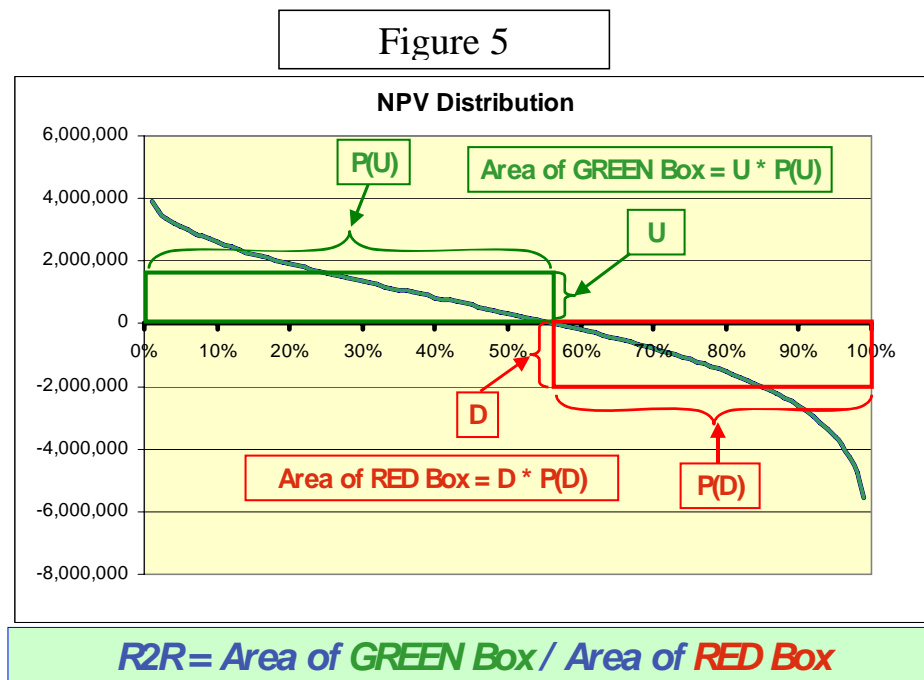


Figure 6 shows a simple demonstration of tail sensitivity. The results of 10,000 simulation trials are shown for three random variables with lognormal distributions.¹⁴ Each has a mean of \$25,000, differing only in the estimate of σ , the parameter associated with dispersion or variance, ranging from 0.28 to 0.32. With the limited sample size

¹³ Keating and Shadwick (2002) introduced the *Omega* function, a close counterpart to RCR and R2R. For a random variable x , $\Omega(x)$ is the ratio of the *magnitude* of conditional expected value of outcomes above x to the *magnitude* of conditional expected value of outcomes below x . It is a unique, one-to-one transform of a probability distribution, similar to the moment-generating or characteristic functions. Ω has been advocated as a superior decision variable for investment portfolio selection.

¹⁴ The natural logarithm of a lognormal variable has a normal distribution with parameters μ and σ .

commonly found with some operational risks, this degree of variation in parameter estimate is not atypical.

Figure 6
TVaR Sensitivity Testing

10000 Trials
LogNormal Variates

Mu	10.087	10.082	10.075	
Sigma	0.28	0.30	0.32	
Mean	\$ 25,000.00	\$ 25,000.00	\$ 25,000.00	
VaR 99.9%	\$ 54,405.73	\$ 62,044.20	\$ 64,704.18	
TVaR 99.9%	\$ 56,919.36	\$ 64,645.66	\$ 70,933.27	
Capital	\$ 31,919.36	\$ 39,645.66	\$ 45,933.27	=TVaR - Mean
Diff	-19.5%		15.9%	
VaR 95%	\$ 37,918.88	\$ 39,498.30	\$ 40,253.85	
TVaR 95%	\$ 42,565.05	\$ 45,223.47	\$ 46,270.23	
Capital	\$ 17,565.05	\$ 20,223.47	\$ 21,270.23	=TVaR - Mean
Diff	-13.1%		5.2%	
R2R	1.258	1.268	1.325	
Diff	-0.8%		4.5%	

The tail sensitivity using TVaR 99.9 percent (i.e., 1-in-1,000 outcomes) for capital determination is dramatic: swings of +16 percent to -19 percent. Using a slightly less extreme TVaR 95 percent metric does temper the effect, but the more stable R2R metric shows the least sensitivity.

This example is not meant to be a rigorous statistical test. It merely aims to alert the practitioner to the fact that tail sensitivity in a risk metric is a real issue that should be thought through and tested carefully.

5. Examples: Risk Mitigation and Cost Assessment

5.1 Considerations in Choosing an Approach

Cost of risk allocation is a complex problem with no single “right” or “best” answer. Decision-making under uncertainty within an organization involves a mixture of theory and politics. Modern finance theory dictates that the portfolio view of risk is the only view that matters. However, a portfolio view requires a portfolio risk model with dependencies and correlations. Correlation links the fate of multiple product areas

and business units. A political debate can ensue, as affected parties realize that these parameters are driving cost allocations, portfolio mix planning, bonuses, and future prospects. Also, correlations are often among the most difficult parameters to estimate. Data is rarely sufficient, and a fair amount of subjectivity and expert opinion is necessary.

Politics may necessitate compromises, sequential adoption and evolution of practice. For political expediency, (and change acceptance), practitioners may want to consider starting with a standalone (SA) approach, then moving to a portfolio approach over time, possibly with parallel testing. SA approaches only use the risk metrics for each product in isolation. While this may appear theoretically inferior to a portfolio approach, it has the advantage of being less politically contentious. Business unit advocates may be more willing to accept the approach and the risk-modeling framework, as well as the indications, if the metrics are under their control.

5.2 Capital Allocation

There are numerous references available demonstrating capital allocation techniques. See for example Meyers, Klinker and Lalonde (2003), Myers and Read (2001), Venter (2003) and Merton and Perold (1993). Since many other authors cover the approach thoroughly, it will not be treated in any detail here.

5.3 Example with Standalone (SA)

This section will demonstrate one SA approach with a robust metric (R2R). Figure 8 shows the simplified example. A business unit (BU #1) has modeled 10 possible scenarios for its operational loss before any mitigation effect. For purposes of assessing BU #1 for its operational risk, corporate risk management will assume (take on) any loss amount in excess of BU #1's expected loss. For this risk reduction, BU #1 must pay corporate an internal insurance *premium*. Corporate calculates the premium such that its net position (premium minus any excess loss transferred from BU #1) distribution has an R2R = 3.

BU #1 wants to lower its premium, so it implements a mitigation program. This program limits any operational loss > 80 to a maximum of 80. The resulting modified operational loss distribution is shown in the "Mitigated Loss*" column of Figure 7.

Figure 7
Operational Risk Cost Assessment Example
Standalone BU Mitigation Reward

Scenario	Starting			Mitigated		
	Loss BU #1	Excess Loss	Corporate Position	Loss* BU #1	Excess Loss*	Corporate Position*
1	\$ 56.00	\$ -	\$ 19.67	\$ 56.00	\$ -	\$ 16.50
2	\$ 24.00	\$ -	\$ 19.67	\$ 24.00	\$ -	\$ 16.50
3	\$ 13.00	\$ -	\$ 19.67	\$ 13.00	\$ -	\$ 16.50
4	\$ 55.00	\$ -	\$ 19.67	\$ 55.00	\$ -	\$ 16.50
5	\$ 89.00	\$ (28.50)	\$ (8.83)	\$ 80.00	\$ (23.00)	\$ (6.50)
6	\$ 77.00	\$ (16.50)	\$ 3.17	\$ 77.00	\$ (20.00)	\$ (3.50)
7	\$ 27.00	\$ -	\$ 19.67	\$ 27.00	\$ -	\$ 16.50
8	\$ 78.00	\$ (17.50)	\$ 2.17	\$ 78.00	\$ (21.00)	\$ (4.50)
9	\$ 90.00	\$ (29.50)	\$ (9.83)	\$ 80.00	\$ (23.00)	\$ (6.50)
10	\$ 96.00	\$ (35.50)	\$ (15.83)	\$ 80.00	\$ (23.00)	\$ (6.50)

Exp Loss	\$ 60.50	Exp Loss*	\$ 57.00	
Premium	\$ 19.67	Premium*	\$ 16.50	-16.1%

U	\$ 14.81	U*	\$ 16.50	<i>Improved</i>
D	\$ (11.50)	D*	\$ (5.50)	<i>Improved</i>
P(U)	70.0%	P(U)*	50.0%	
P(D)	30.0%	P(D)*	50.0%	
R2R	3.00	R2R*	3.00	Target = 3.00

Corporate takes the modified loss distribution and calculates the new premium such that its new net position also has an R2R = 3. BU #1's premium goes from \$19.67 to 16.50, a -16.1 percent reduction. This savings can be compared with the cost of the mitigation for purposes of BU #1's CBA.

5.4 Example with Portfolio

Now we move up a level to a portfolio method we will call the "RISK X-RAY" (see Ruhm (2004) for an accessible overview article on the methodology). This method is a transparent way to distribute the costs of risk and benefits of mitigation to the components of a risk portfolio. It has been published in the actuarial literature using several names, including Ruhm-Mango-Kreps or "RMK," conditional risk charges, riskiness leverage models and co-measures.¹⁵

¹⁵ See Ruhm and Mango (2003) and Kreps (2005) for detailed discussion.

The method begins by evaluating risk at the aggregate or portfolio level, and ends by producing risk charges for individual components, effectively allocating the total portfolio risk charge by each component's contribution to total portfolio risk. The result is an internally consistent allocation of diversification benefits. The method extends risk valuation from the aggregate portfolio level down to the individual risks that comprise the portfolio, reflecting each component's contribution to total portfolio risk.

The portfolio risk measure can be essentially any choice that represents the management point of view, so that many attitudes toward risk can be expressed. The total risk charge is allocated down to all risk sources simultaneously using just one formula, and the allocations are perfectly additive across arbitrary subgroups and the whole enterprise. Thus, attention can be turned to what is a plausible risk measure for the whole, if one is prepared to use the indicated allocation technique for the pieces.

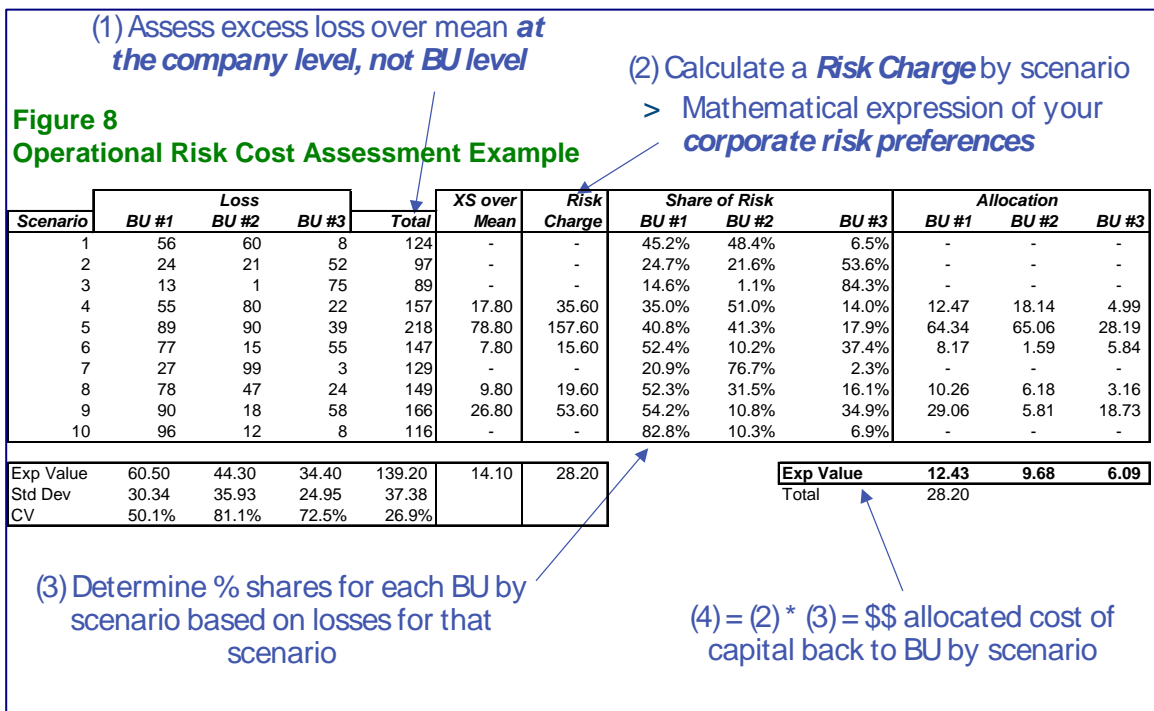


Figure 8 shows the output from a full IRM for three BU's. It includes the outcome distribution for each BU and the company total. This application of the Risk X-Ray will follow four major steps.

1. Assess the excess loss above the mean *at the company level*.
2. Calculate a risk charge by scenario using that excess loss amount. This is a mathematical expression of corporate risk preferences. In the example, the

risk charge is simplistically set to be two times the excess loss above the mean.

3. Determine each BU's share of the scenario risk charge based on that BU's losses for that scenario.
4. Allocate the risk charge amount by scenario back to each BU. This allocation is a simple product of the total risk charge and the BU % share.

The BU risk charge is the expected value of its allocations over all scenarios.




Figure 9 shows the same exhibit with BU #2's loss distribution reflecting its mitigation efforts. Similar to the SA example, their mitigation caps any loss amounts > 80 at a maximum value of 80. The Risk X-Ray follows the same four steps. The key results are highlighted.

Every BU actually sees a reduction in risk charge. This is because BU #2's mitigation efforts reduced the *aggregate distribution*, which lowered the risk charges for everyone. But BU #2 enjoys the largest percentage reduction. This highlights the potential political aspects of moving from an SA to a portfolio approach:

- Portfolio approaches reward for diversification. This means BUs that naturally hedge each other will mutually benefit from their consolidation into one portfolio.
- Portfolio approaches share benefits and penalties. Part of the improvement from mitigation will go to peer BUs. Similarly, bad outcomes from one BU can bleed over and impact others.

6. Conclusions

This paper has provided a high-level overview of the potential synergies in applying actuarial techniques to operational risk analysis. It also demonstrated practical techniques, grounded in actuarial science, to solve operational risk management problems such as capital cost allocation and mitigation cost-benefit analysis. It is hoped that this paper will be the first step in an ongoing integration of actuarial techniques into all aspects of risk analysis.

There are many areas of needed research:

- Improvement of internal risk modeling techniques and software;
- Industry event data collection and aggregation;
- Company event data collection and aggregation;
- Industry default frequency and severity distributions;
- Standards for correlations and shock scenarios; and
- Elicitation and elucidation of risk preferences.

References

Bouska, A. 1989. Exposure bases revisited. *Proceedings of the Casualty Actuarial Society* LXXIV: 1-14.

Available online at www.casact.org/pubs/proceed/proceed89/89001.pdf.

Bodie, Z., and Merton, R. 2000. *Finance*, 2nd Edition. Upper Saddle River, NJ: Prentice Hall.

Kalkbrener, M. 2005. An axiomatic approach to capital allocation. *Mathematical Finance* 15 (3): 425-437.

Keating, C., and Shadwick, W. 2002. A universal performance measure. The Finance Development Centre Working Paper.

Available online at http://www.isda.org/c_and_a/pdf/GammaPub.pdf.

Kreps, R. 2005. Riskiness leverage models. Accepted to the *Proceedings of the Casualty Actuarial Society*.

Available online at <http://www.casact.org/library/05pcas/kreps.pdf>.

Mango, D.F. 2003. Capital consumption: An alternative methodology for pricing reinsurance. *Casualty Actuarial Society Forum*. Winter, Ratemaking Call Paper Program, 351-378.

Online at <http://www.casact.org/pubs/forum/03wforum/03wf351.pdf>.

Mango, D.F. 2005. Insurance capital as a shared asset. Accepted to *ASTIN Bulletin*.

Available online at

ftp://ftp.math.ethz.ch/hg/users/astin2005/contributed_papers_astin/Mango.pdf.

Merton, R., and Perold, A. 1993. Theory of risk capital in financial firms. *Journal of Applied Corporate Finance* 6(3): 16-32.

Meyers, G., Klinker, F., and Lalonde, D. 2003. The aggregation and correlation of insurance exposures. *Casualty Actuarial Society Forum*, Summer, DFA Call Paper Program, 15-82.

Available online at <http://www.casact.org/pubs/forum/03sforum/03sf015.pdf>.

Myers, S.C., and Read, J.A. Capital allocation for insurance companies. 2001. *Journal of Risk and Insurance* 68:4: 545-580.

Available online at <http://www.aria.org/journals/jri/2001/december/myers.pdf>.

Ruhm, D. 2001. Risk coverage ratio: A leverage-independent method of pricing based on distribution of return. Presented at the 32nd *International ASTIN Colloquium*, July.

Available online at

<http://www.casact.org/coneduc/reinsure/astin/2000/paperspresented.htm>.

— — — . 2004. An X-ray for risk. *Actuarial Review*, November: 28-30.

Available online at <http://www.casact.org/pubs/actrev/nov04/latest.htm>.

Ruhm, D., and Mango, D. 2003. A risk charge calculation based on conditional probability. Presented at the 2003 *Thomas P. Bowles Jr. Symposium*, Georgia State University.

Online at www.casact.org/coneduc/specsem/sp2003/papers/.

Samad-Khan, A. 2005. Actuarial techniques in banking operational risk. Presented at *CAS Annual Meeting*.

Available online at <http://www.casact.org/coneduc/annual/2005/handouts/khan.pdf>.

— — — . Why COSO is flawed. 2005. *Operational Risk*, January.

Available online at <http://www.opriskadvisory.com/articles.html>.

Spetzler, C.S. 1968. The development of a corporate risk policy for capital investment decisions. *IEEE Transactions on Systems Science and Cybernetics*, SSC-4 (3) September: 279-300.

Venter, G. 2003. Capital allocation: An opinionated survey. *Casualty Actuarial Society Forum*, Summer: 279-306.

Available online at <http://www.casact.org/pubs/forum/03sforum/03sf279.pdf>.

Venter, G., Major, J.A., and Kreps, R. 2005. Additive marginal allocation of risk measures. Submitted to *ASTIN Bulletin*.

Walls, M.R. 1995. Corporate risk tolerance and capital allocation: A practical approach to setting and implementing an exploration risk policy. *Journal of Petroleum Technology*, SPE Paper No. 28281, April: 307-311.

Available online at www.econbus.mines.edu/walls-pubs.htm.

— — — . 2004. Combining decision analysis and portfolio management to improve project selection in the exploration and production firm. *Journal of Petroleum*

Science and Engineering 44: 55-65.

Available online at www.econbus.mines.edu/walls-pubs.htm.

Wang, S.S. 1998. Aggregation of correlated risk portfolios: Models and algorithms.

Proceedings of the Casualty Actuarial Society LXXXIII: 848-939.

Available online at www.casact.org/pubs/proceed/proceed98/980848.pdf.