

A New Approach for Managing Operational Risk

Addressing the Issues Underlying the 2008 Global Financial Crisis

Sponsored by:
Joint Risk Management Section
Society of Actuaries
Canadian Institute of Actuaries
Casualty Actuarial Society



Prepared by:



Originally Published: December 2009
Revised: July 2010

© 2009, 2010 Society of Actuaries, All Rights Reserved

The opinions expressed and conclusions reached by the authors are their own and do not represent any official position or opinion of the sponsoring organizations or their members. The sponsoring organizations make no representation or warranty to the accuracy of the information.

Table of Contents

1. Executive Summary	3
2. Introduction	7
2.1 Project Background and Scope	7
2.2 Project Team	8
3. What is Risk?	10
3.1 Origin and Use	10
3.2 Traditional and Modern Interpretations	10
3.3 A Practical Definition of Risk	13
4. Key Risk Concepts	15
4.1 Likelihood vs. Frequency	15
4.2 Expected Loss and Unexpected Loss	17
4.3 Risk Measurement and Assessment	20
4.4 Risk Assessment/Measurement Under Traditional ORM	21
4.5 The Educational Challenge	23
5. What is Operational Risk?	25
5.1 The Nature and Magnitude of Operational Risk	25
5.2 Wrong Turn: From Operational Risk to Operations Risk	25
6. Risk Architecture and Taxonomy	28
6.1 The Traditional Risk Universe	28
6.2 The Modern Risk Universe	29
6.3 Modern Operational Risk Taxonomy	31
6.4 Criminal Risk and Principal-Agent Risk	33
7. The ORM Business Problem	36
7.1 Roles and Responsibilities	36
7.2 Modern ORM in Practice	37
7.3 The Modern ORM Infrastructure	40
7.4 Traditional ORM in Practice	42

8. Measuring/Assessing Operational Risk	45
8.1 Goals of Measuring/Assessing Operational Risk	45
8.2 The Actuarial Approach	45
8.3 Data Requirements	47
8.4 Frequency and Severity Distributions.....	50
8.5 How to Model Frequency	51
8.6 How to Model Severity	52
8.7 Combining Internal and External Loss Data.....	57
8.8 Risk Assessment, Scenario Analysis and Stress Testing.....	60
8.9 Calculating Value at Risk	62
8.10 Modeling the Operational Risk Component of Other Risks	65
9. The Business Case for Modern ORM	67
9.1 The Current State of ORM.....	67
9.2 Key Differences between Traditional and Modern ORM	68
9.3 The ORM Evolutionary Path	69
9.4 The ORM Roadmap.....	70
9.5 The Economics of Modern ORM.....	71
10. Conclusions and Recommendations	73
11. References	77
12. Appendix A — Principal-Agent Risk and the 2008 Global Financial Crisis	79
The Sub-Prime Credit Crisis	79
AIG and Credit Default Swaps	81
Summary and Conclusions	81
13. Appendix B — Modeling ORM: Key Concepts	83
14. Appendix C — Glossary of Operational Risk Terminology.....	86

1. Executive Summary

Operational risk, broadly speaking, is the risk of loss from an operational failure. It encompasses a wide range of events and actions as well as inactions and includes, for example, inadvertent execution errors, system failures, acts of nature, conscious violations of policy, law and regulation, and direct and indirect acts of excessive risk taking.

Operational losses can be caused by junior staff; but they can also be caused by mid-level officers, senior managers, C level executives and Boards of Directors. They are sometimes caused by individuals and in other cases by groups of people working in collusion. Many of the largest losses take place when operational failures are present at the senior-most level.

Operational failure has played a role in virtually every catastrophic loss that has taken place during the past 20 years. In fact, the 2008 global financial crisis was largely caused by a series of massive operational failures. The American Insurance Group (AIG) event, which was an example of *principal-agent risk*, may be the single largest corporate loss ever recorded. (Principal-agent risk and its impact on the financial crisis are discussed in Appendix A.) Therefore, a natural question is whether there is a better approach to managing operational risk – one that might have either prevented or mitigated many of these events. This paper outlines such an approach; it also explains why the methods commonly used today may not adequately meet this challenge.

For many years the conventional wisdom has held that operational risk was best managed through a traditional audit-based approach. Going forward, we will refer to this approach as the Traditional Approach or Traditional operational risk management (ORM). In the United States, the broad principles underlying this general approach have been incorporated into a set of standards that are referred to as COSO ERM.

Virtually all the major accounting firms worldwide recommend using the Traditional Approach for managing operational risk. Numerous consulting firms, rating agencies, industry bodies and independent experts also advocate using this approach or a customized version thereof. A majority of national and international bank regulators have also at least tacitly endorsed this approach. Finally, a large number of corporate CFOs believe that the Traditional Approach represents the standard for best practices in ORM. Consequently, virtually every organization that has implemented an ORM or ERM program has based the underlying framework on the principles of Traditional ORM.

The Traditional Approach has many useful features. It provides structure, governance standards and an intuitive approach to risk identification and assessment. But it also has some drawbacks. It is based on a conception of risk which is inconsistent with that used in the actuarial and risk management disciplines (described in section 3.3). One major discrepancy is that under the Traditional Approach risk is associated with the average loss. Yet, in the

actuarial and risk management disciplines, risk represents uncertainty with respect to loss exposure or the “worst case” loss.

This discrepancy has several implications. Specifically, because Traditional ORM focuses attention on the set of commonly observable threats and control weaknesses associated with routine losses it fails to reveal the largest risks. Therefore, institutions that follow the Traditional Approach may be unaware of their most significant risks. In addition, organizations that base risk-control optimization decisions on the results of Traditional risk and control self assessment (RCSA) can easily become over-controlled in the areas where they have the least risk, but remain significantly under-controlled in the areas where they have the most risk.

The Traditional Approach is very effective for preventing losses at a tactical level, but loss prevention addresses only one aspect of the ORM business problem — and not the most important one. In particular, Traditional ORM does little to mitigate exposure to the large catastrophic events, such as sales and business practices violations and acts of excessive risk taking, which are really the key drivers of operational risk.

One of the most important operational risks is “principal-agent” risk. Principal-agent risk refers to the risk that, in circumstances where there is separation of ownership and control, agents (who control or act on behalf of the organization) may pursue actions that are in their own interest, but are not necessarily in the best interest of the principals (the stakeholders). Principal-agent risk has been the driving factor behind many of the largest losses, including the AIG event. (Principal-agent risk is defined in section 6.4)

An analysis of the AIG credit default swap debacle reveals that AIG senior management (the agents) decided that because no regulation specifically required that they reserve capital for credit default swaps, they did not need to reserve any additional capital to cover the associated risk¹. By aggressively pursuing this business they were thus able to generate abnormally high returns on capital, for which they were very well compensated. However, they did this only by taking risks far in excess of tolerance standards of the stakeholders (the principals). Since AIG was deemed by the U.S. Treasury and Federal Reserve Board to be “too big to fail,” the ultimate stakeholders turned out to be the U.S. taxpayers.

In recent years, a new approach to managing operational risk has been introduced. This new approach is called Modern ORM. Modern ORM is a top-down approach, which focuses first on the major risks — within a comprehensive and mutually exclusive risk architecture — and drills down only in those risk areas where more granularity is required. This holistic and systematic approach allows practitioners to triage the risk management process. Because it is significantly less resource-intensive, it avoids focusing management attention and resources on immaterial risks. Modern ORM could also prove to be very effective in mitigating principal-agent risk.

¹ New York State Insurance Department News Release: “Allstate Should Report Any Illegal or Inappropriate Use of Credit Default Swaps”; (April, 2009)

The Solvency II regulations, which are scheduled to become effective in Europe in 2012, use language that is consistent with Modern ORM. These regulations describe risk as a measure of uncertainty (specified at a 99.5% confidence level for a one-year time horizon). In addition, the Solvency II “use test” requires that internal models must play an integral role in any organization’s system of governance and risk management as well as its economic and solvency assessment/capital allocation. Solvency II will directly impact North American companies with international operations and will eventually translate into a corresponding set of U.S. and Canadian regulations. Moreover, the rating agencies are also likely to evaluate insurance companies based on these or similar standards; some have already expressed an intention to do so.

Besides helping organizations satisfy compliance requirements, Modern ORM can also produce tangible value. Specifically, Modern ORM addresses the most critical risk management business problem, which is mitigating exposure to the large events — the events that have the greatest impact on financial performance and solvency. Modern ORM aims to:

- Facilitate the holistic management of all operational risks, based on a consistent definition of risk and a comprehensive risk architecture/taxonomy.
- Create a structured and transparent process for factoring risk into the business decision-making process — at both a tactical and strategic level. Specifically, provide managers, senior managers and C level executives the tools and information they need to optimize risk-reward, risk-control and risk-transfer in the context of cost-benefit analysis.
- Embed a risk culture that harmonizes the goals of key decision makers and external stakeholders.
- Reduce information asymmetries between managers and stakeholders to help confirm that managers are pursuing strategies that conform to the risk tolerance standards of the stakeholders — in other words, mitigate principal-agent risk.

For the reasons given above, the authors recommend that North American insurance companies consider developing formal ORM programs. These programs would benefit from the principles of Modern ORM. Traditional ORM has many useful aspects, and the authors recommend that some of these program features also be adopted or retained. However, companies that have already developed comprehensive Traditional ORM programs may find it beneficial to consider scaling back several of the highly resource-intensive program components and replacing them with Modern ORM equivalents. Transitioning from Traditional to Modern ORM, if implemented correctly, may not only improve risk management effectiveness, but may also significantly reduce cost.

The key differences between Traditional and Modern ORM are summarized in Exhibit 1.1 below.

Exhibit 1.1 — Summary of Differences between Traditional and Modern ORM

Traditional ORM	Modern ORM
<ul style="list-style-type: none"> ■ Definition: Risk is defined primarily as a kind of undesirable incident/event, such as a fraud or a system failure (Operative question: What/where are your risks?) 	<ul style="list-style-type: none"> ■ Definition: Risk is defined primarily as a measure of exposure to loss from undesirable incidents/events (Operative question: How much risk do you have?)
<ul style="list-style-type: none"> ■ Risk Identification Process: Ask managers to identify their major risks. (Risks include risk factors, controllable factors, events and effects; no restriction on overlaps; generally no differentiation made between risks and controls.) Leads to the creation of a huge and unmanageable set of risks 	<ul style="list-style-type: none"> ■ Risk Identification Process: First define the “risk” universe, consisting of a finite (comprehensive) set of mutually exclusive (non-overlapping) “risk” classes. Use hard or soft data to reveal where the large losses are taking place (where the largest risks actually exist)
<ul style="list-style-type: none"> ■ Risk Assessment/Measurement Method: Calculate risk by multiplying likelihood and impact for each risk type (conditional on one event), one “risk” at a time 	<ul style="list-style-type: none"> ■ Risk Assessment/Measurement Method: Use Monte Carlo simulation and frequency and severity distributions to calculate the cumulative loss potential from multiple events, across all risk classes simultaneously
<ul style="list-style-type: none"> ■ Aggregation: Likelihood cannot be aggregated, so results cannot be aggregated 	<ul style="list-style-type: none"> ■ Aggregation: Frequency can be aggregated, so results can be aggregated
<ul style="list-style-type: none"> ■ What is measured: Probability weighted loss from one specific incident (the routine loss) 	<ul style="list-style-type: none"> ■ What is measured: Cumulative loss for one or more risk classes; both the expected loss and unexpected loss, which are comparable to the average and “worst case”
<ul style="list-style-type: none"> ■ Goal: Day-to-day management of current threats arising from imminent operational failures: loss prevention through tactical intervention 	<ul style="list-style-type: none"> ■ Goal: Management of <i>key</i> risks, specifically the optimization of risk-reward, risk-control and risk-transfer in the context of cost-benefit analysis
<ul style="list-style-type: none"> ■ Cost: Generally very resource intensive 	<ul style="list-style-type: none"> ■ Cost: Relatively much less resource intensive

2. Introduction

2.1 Project Background and Scope

Operational risk is not a new risk, but hard evidence suggests that this risk is significant and may be growing. Virtually every catastrophic financial institution loss that has taken place during the past 20 years — including Barings Bank, Long Term Capital Management, Allied Irish Bank-All First, Société Générale, Bear Stearns, Lehman Brothers and American Insurance Group (AIG) — has been caused or exacerbated by operational failure. In fact, the 2008 AIG event, which was caused by excessive risk taking involving *principal-agent* risk, resulted in aggregate losses in excess of \$100 billion — making it one of the largest individual losses ever recorded. As a result, ORM is gaining greater visibility within the insurance industry as well as many other industries.

In advance of the Solvency II² regulations, which are currently scheduled to become effective in 2012, many large European insurance companies have established formal ORM programs. Even though there are no corresponding North American insurance industry regulations, several U.S. and Canadian insurance companies (particularly those with global operations) have also established formal ORM programs, and many others are considering doing so.

Developing an effective method of managing operational risk is proving to be a daunting task, however. At present there is little industry consensus regarding what constitutes a comprehensive ORM framework or how to implement a viable ORM program. Many insurance companies have been leveraging the experience of the banking industry, which has been focused on the operational risk problem for over ten years. It is not clear whether this approach is prudent, however; at present, very few banks are willing to claim that their ORM programs have met original expectations or have added tangible value.

This raises several important questions. For example, what is the value proposition for insurance companies to establish formal ORM programs? And, are traditional methods of managing operational risk sufficient to meet the desired objectives? If not, is there some alternative method that may allow them to do so? In particular, is there a framework/methodology that, if properly implemented, could have prevented or minimized the impact of the 2008 global financial crisis?

To explore these questions the Joint Risk Management Section of the Canadian Institute of Actuaries, Casualty Actuarial Society and Society of Actuaries (together, the Sponsors) decided to sponsor this research project to

² The European insurance authorities have promulgated a set of regulations, referred to as Solvency II, which require insurance companies operating within Europe to improve their overall approach to risk management. These regulations are similar to the Basel II banking regulations, which were introduced in 2004. Both the Basel II and Solvency II regulations include specific provisions for the management of operational risk, including the calculation of operational risk capital based on standard formulae or an internal model.

examine ORM practices and assess the viability of ORM as a formal discipline. A key goal of this initiative has been to determine whether the management of operational risk is feasible, and, if so, what issues need to be addressed in order to effectively implement ORM within a broader ERM context.

Towers Perrin and OpRisk Advisory (together the Research Team) were engaged to do this work. The Research Team suggested and the Sponsors agreed that in order to address these issues it was important to review some of the fundamental assumptions upon which most existing ORM frameworks are based and to determine whether these assumptions are appropriate. Therefore, key issues addressed in this report include:

- How should the term risk be defined within the context of ORM?
- What are the key risk concepts and how should they be applied in ORM?
- What is the true nature and magnitude of operational risk?
- What is the ORM business problem?
- How can one assess and measure operational risk using hard data, soft data and/or expert opinion (scenario analysis)?
- What are the drawbacks of the Traditional Approach? What would be the benefits of transitioning to an alternative approach (e.g., the Modern Approach), and what is the evolutionary path?

Developing a set of best practices for implementing ORM is not part of the scope of this project, but may be the focus of a future research initiative.

It is anticipated that this report might also be used as a basis for constructive dialogue with rating agencies and regulators concerning how future risk regulations, including risk capital requirements, should be developed.

2.2 Project Team

The Research Team consisted of the following Towers Perrin and OpRisk Advisory staff:

- Ali Samad-Khan — OpRisk Advisory and Towers Perrin
- Sabyasachi Guharay — OpRisk Advisory and Towers Perrin
- Barry Franklin — Towers Perrin
- Bradley Fischtrom — Towers Perrin
- Mark Scanlon — Towers Perrin

-
- Prakash Shimpi — Towers Perrin

Significant industry input was provided by the Project Oversight Group:

- Anver Kasmani — MetLife Services Limited
- Mike O'Connor — MetLife, Incorporated
- Wendy Hart — Nationwide Mutual Insurance
- Gerald Wilson — Old Mutual US Life
- Steven Siegel — Society of Actuaries
- Fred Tavan — SunLife Financial (Chair)
- Ginnie Welsman — SunLife Financial
- Hercules Gray — USAA
- Tim Hieger — USAA
- Mary Gardner — Zurich NAC

3. What is Risk?

3.1 Origin and Use

“The term ‘risk,’ as loosely used in everyday speech and in economic discussion, really covers two things which, functionally at least, are categorically different.”³ This conceptual ambiguity is the root cause of much of the confusion in ORM today. To a large extent, confusion about the meaning of risk and related concepts has also obfuscated communication between the risk management function and senior management in most organizations.

The term risk is most commonly used in a qualitative context. In this context one might say, for example: “I am exposed to fraud risk.” From this statement it appears that the term risk describes a type of unpleasant event, incident or condition, such as a fraud, a system failure or a lack of resources. However, this statement is actually an abbreviated form of the expression: “I am exposed to the risk of loss from fraud events.” Therefore, in the original, unabridged context, risk is not a type of event. Instead, risk is a metric that describes the level of exposure to an adverse consequence, for example, the level of loss exposure to a fraud event. Yet, because the English language supports the abbreviated use of the term, this common understanding of risk has become widely accepted (see Section 3.2. and Exhibit 3.3).

3.2 Traditional and Modern Interpretations

The fact that there is more than one operative understanding of risk is not a trivial or academic matter. Largely because of this ambiguity, there are now two very different schools of thought on how to manage operational risk: Traditional and Modern ORM. And each approach incorporates within it a completely different conception of risk, language and classification scheme, business problem, management framework, measurement methodology, etc.

Traditional ORM is predicated on an understanding of risk that might be expressed as follows:

“Risk is the possibility that an event will occur and adversely impact the achievement of the entity’s mission or business objectives.”


Thus, under the Traditional Approach, risk measurement has come to mean measuring the probability of a loss. And it naturally follows that, with respect to risk measurement, risk = probability x (loss) impact.

³ Knight, Frank, “Risk, Uncertainty and Profit,” 1921.

Under Modern ORM, risk is a measure of exposure to loss at a level of uncertainty. Under this definition, risk requires both exposure and uncertainty.⁴ Where the probability of loss is 100%, because the loss is certain, the level of risk is zero. Under Modern ORM, probability x impact is referred to as the “expected loss.”

Exhibit 3.1 illustrates that under the Traditional interpretation, maximum risk exists where the probability of loss is 100% — i.e., the loss is certain. But, under the Modern interpretation, high risk is characterized by low probability (or low frequency) and high severity. These two conceptions of risk measurement are not just materially different — they are contradictory.

Exhibit 3.1 — Traditional vs. Modern Conceptions of High Risk



Because confusion about the meaning of risk is ubiquitous in the financial services industry, many organizations have unknowingly implemented ORM frameworks where different elements — for example, risk and control self-assessment (RCSA) and risk capital measurement — are based on different and contradictory definitions of risk. Therefore, these programs cannot legitimately integrate risk information drawn from different parts of their frameworks and consequently cannot be used to support many of the stated goals and objectives of ORM.

To better understand the meaning of risk, consider the example on the next page.

⁴ Holton, Glyn A., “Defining Risk,” Financial Analysts Journal, Volume 60, No. 6, 2004, CFA Institute.

Exhibit 3.2 — Understand the Meaning of Risk

Suppose you have the following three investments and their associated risk-and-return information:

- Investment A: Guaranteed return of 10%.
Investment B: 50% probability of a 0% gain; 50% probability of a 20% gain.
Investment C: 50% probability of a 10% loss; 50% probability of a 30% gain.

Question 1: Which investment has the highest expected return?

If you sum up the probability-weighted returns, you can calculate that all three investments have the same average or expected return, which is 10%. A's return is fixed at 10%. B's expected return is determined as $0.5 * 0\% + 0.5 * 20\%$, which is equal to 10%. Likewise, C's expected return is $0.5 * -10\% + 0.5 * 30\%$, which again is equal to 10%.

Question 2: Which investment has the most risk?

Anyone with experience in portfolio analysis will recognize that investment A, because it offers a guaranteed return of 10%, has no risk. Investment B has no chance of a loss. Its worst-case outcome is a break-even position, but it offers a 50% chance of a return that is below the mean return. Therefore, investment B has some risk. Investment C, which has the largest downside (–10% in absolute terms and –20% from the mean return), has the most risk.

Therefore, one can see that risk represents the level of uncertainty surrounding an adverse consequence — not the adverse consequence itself.

Question 3: How much risk is there in each investment?

Risk can be expressed in several ways. While the relative level of risk can be determined as explained above, the absolute risk of loss cannot be measured without first specifying a probability (confidence) level for the distribution of possible outcomes (at which we want to measure risk), for example, 99%. So there is not enough information to answer this question. This confidence level can be used to express risk tolerance in monetary terms.

Question 4: Which is the best investment?

It is important to recognize that risk is neither inherently good nor bad. So there is not enough information to answer this question.

A risk-neutral person ignores variance. He or she evaluates investments purely on the basis of expected outcomes — irrespective of the level of uncertainty associated with these potential outcomes. Since all three investments offer the same average (expected) return of 10%, a risk-neutral person would regard all three investments to be of equal value.

A risk lover would prefer investment C. In fact, he or she would be willing to pay a premium for an investment that offers the potential for a 30% gain, which is 20% in excess of the mean return.

A risk-averse person would choose investment A because it offers the same expected return as the other investments, but with less risk — in fact, none at all. Because the majority of investors and financial institutions are risk-averse, they demand higher levels of return for higher levels of risk. This explains why more risky (more volatile) investments, when priced accurately, pay higher expected returns.

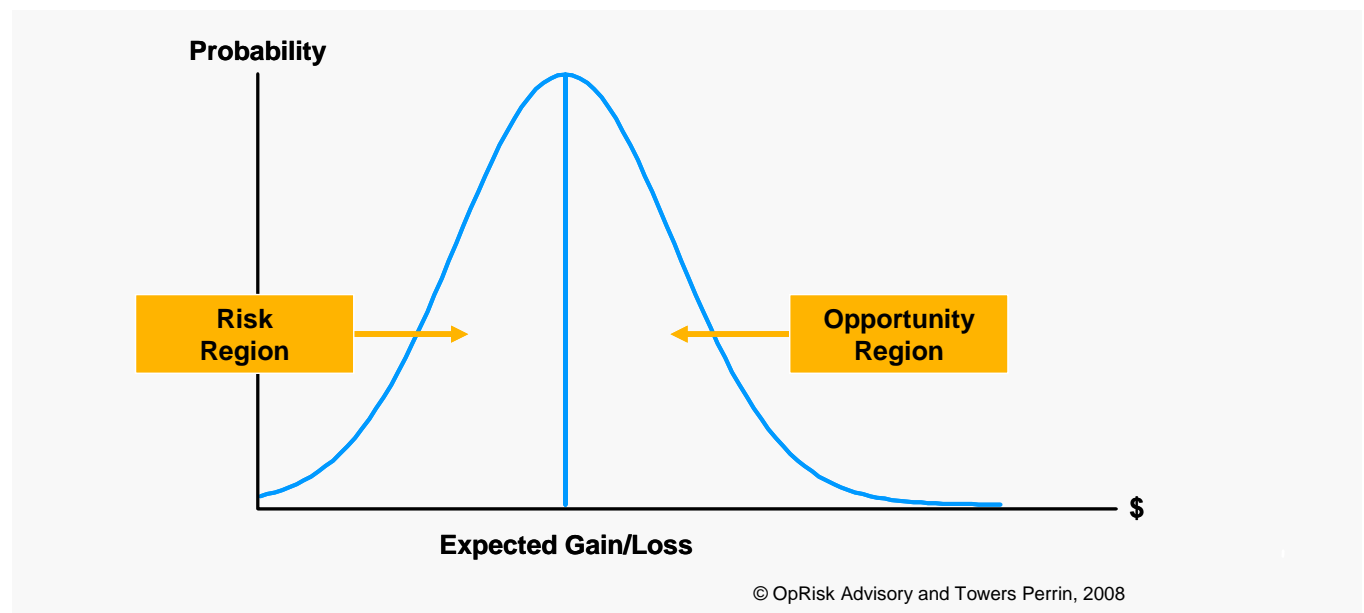
In summary, as mentioned above, risk is not a type of incident, it is a measure. It describes a level of variance or uncertainty with respect to some adverse consequence. Only where there is certainty is there no risk.⁵

© 2004 OpRisk Advisory

⁵ Samad-Khan, Ali, A. Rheinbay and S. Le Blevet, “Fundamental Issues in OpRisk Management,” *OpRisk & Compliance* (February 2006).

Exhibit 3.3 illustrates that risk is a measure of deviation from the expectation (mean). While in a general sense risk is equivalent to uncertainty, from an ORM perspective the “risk” region covers the negative variance, and positive deviations from the expected outcome define the “opportunity” region. These two regions are closely related; in fact, it is virtually impossible for opportunity (“upside risk”) to exist without the concurrent existence of the risk of loss (“downside risk”).⁶ However, people do not typically refer to the positive variance as risk, because we do not say “the *risk* of gain;” instead we say “the *opportunity* for gain.” Once again, within the context of risk management, risk is expressed in terms of an adverse consequence (such as loss), which is the convention we follow in this paper.

Exhibit 3.3 — Risk vs. Opportunity



3.3 A Practical Definition of Risk

In order to avoid confusion, the Research Team recommends that any organization interested in developing an integrated ORM program should adopt the following definition of risk:

“Risk is a measure of adverse deviation from the expectation, expressed at a level of uncertainty (probability).”

However, the Research Team acknowledges that, because the Traditional interpretation of risk, i.e., risk is an adverse or unpleasant incident/event, is now so firmly ensconced in the public vernacular, it would be infeasible to try to curtail its use. It should also be noted that this interpretation has certain practical benefits. In many respects it lends itself to more efficient communication. For example, it is easier to say “I am exposed to fraud

⁶ The U.S. actuarial profession has embraced this broad interpretation of risk in its branding campaign, “Actuaries: Risk is Opportunity.”

risk,” than “I am exposed to the risk of loss from fraud events.” Nevertheless, risk practitioners should use this informal definition with caution. In particular, the informal definition should always be regarded as a subordinate definition and should never be used in a manner that contradicts the formal definition. ORM practitioners should make every effort to be pedantic in risk communication. This is a critical issue, because confusion over the meaning of risk is the root cause of much of confusion in ORM as well as in many other areas of risk management.

4. Key Risk Concepts


4.1 Likelihood vs. Frequency

Under Traditional ORM, the terms likelihood and frequency are often used synonymously, but under Modern ORM these terms have very different meanings. Likelihood means *probability* and is generally used in the context of a single incident or scenario (e.g., the likelihood of getting into a car accident today is 5%). Likelihood is measured on a scale of 0 to 1 (or 0 to 100%).

Frequency describes the *number of events* (e.g., 10 losses per year). Frequency is measured on a scale of 0 to infinity. Mean frequency is the average number of events that have taken place or are expected to take place during a specified period of time. Note: The number of events that can take place is always an integer (e.g., 0, 1, 2, 3 ...), but mean frequency can be a fractional value (e.g., 23.84).

A frequency distribution is a discrete probability distribution for a specified time period, typically one year. A frequency distribution expresses probability values for each possible integer number of events. Exhibit 4.1 shows that a Poisson frequency distribution with a mean of 1 event per year has the following probability mix: 0 events: 36.8%, 1 event: 36.8%, 2 events: 18.4%; 3 events: 6.1%; 4 events: 1.5%, 5 and above events: 0.4%. The vertical axis of a frequency distribution represents probability (likelihood) and the horizontal axis represents the corresponding number of events. In a frequency distribution, as is the case with any probability distribution, total probability must sum to 100%.

Exhibit 4.1 — Frequency Distribution



One reason that likelihood and frequency are used synonymously is that for rare events the likelihood and mean frequency values are nearly equivalent. For example, a one in one thousand year event has a mean annual frequency of 0.001 and also a likelihood (probability) of approximately 0.001 per year. But this relationship does not hold true for the commonly occurring events. For example, an event that is expected to take place once a year, on average, has a mean annual frequency of 1, but the likelihood is usually less than 1, because likelihood of 1 means 100% probability of occurrence. (Note: If an event is expected to take place only once a year, on average, it is not true that that event will take place once a year with 100% certainty.)

Exhibit 4.2 illustrates the difference between likelihood and frequency when frequency follows a Poisson distribution. In this example, for low frequency values likelihood is approximately equal to $1/N$ (where N is the total number of years). But as mean frequency increases the values diverge, such that when mean frequency is one event per year, the corresponding likelihood of exactly one event taking place in any given year is approximately 0.3679 (36.79%); for one or more events it is 0.6321 (63.21%).

Exhibit 4.2 — Likelihood and Frequency

Years (N)	Mean Frequency if on Average 1 Event Occurs Every N Years	Likelihood of Exactly 1 Event in a Single Year	Likelihood of 1 or More Events in a Single Year
N	1/N	Prob. (1 Event)	1-Prob. (0 Events)
1000	0.001	0.000999	0.001000
500	0.002	0.001996	0.001998
200	0.005	0.004975	0.004988
100	0.010	0.009900	0.009950
75	0.013	0.013157	0.013245
50	0.020	0.019604	0.019801
40	0.025	0.024383	0.024690
30	0.033	0.032241	0.032784
25	0.040	0.038432	0.039211
20	0.050	0.047561	0.048771
10	0.10	0.090484	0.095163
5	0.20	0.163746	0.181269
4	0.25	0.194700	0.221199
3	0.33	0.238844	0.283469
2.5	0.4	0.268128	0.329680
2	0.5	0.303265	0.393469
1	1	0.367879	0.632121

4.2 Expected Loss and Unexpected Loss

Under Traditional ORM, the *expected losses* are the smaller or routine losses, and the *unexpected losses* are the large or rare losses. Modern ORM, once again, follows a different language. Under Modern ORM, the terms *expected losses* and *unexpected losses* do not exist⁷. Under Modern ORM the terms *the expected loss* and *the unexpected loss* are metrics and are used in a statistical context. Specifically, the expected loss is the average loss, or the probability weighted mean of a loss distribution, and the unexpected loss is the difference between the total exposure at the target risk tolerance level and the expected loss. The unexpected loss represents the risk.

⁷ Losses that occur within a commonly observed dollar range are often referred to as routine or high frequency events. The large, less commonly observed events are referred to as rare or low frequency events. The vast majority of execution errors fall into the commonly observed dollar range, whereas a lesser proportion of sales practices losses fall into this same range. Therefore, execution errors are often referred to as high-frequency, low-severity losses, and sales practices events are referred to as low-frequency, high-severity losses.

Exhibit 4.3 illustrates the meaning of the expected loss and unexpected loss in a statistical context, where unexpected loss is calculated at the 99% level (the target risk tolerance level).

Exhibit 4.3 — Expected Loss and Unexpected Loss

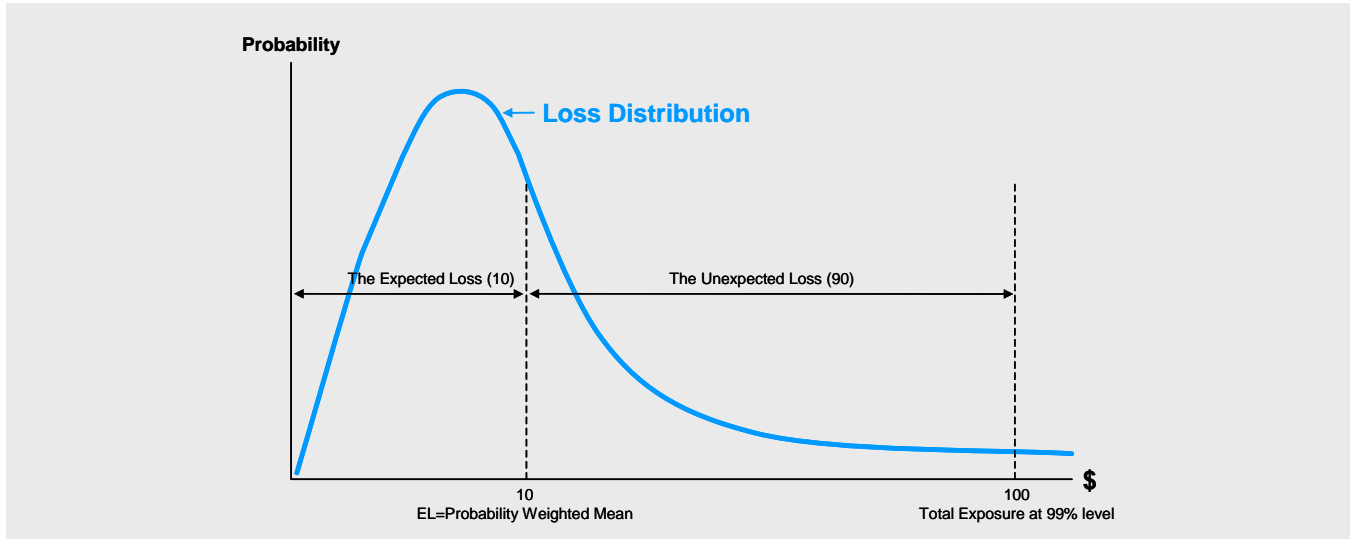
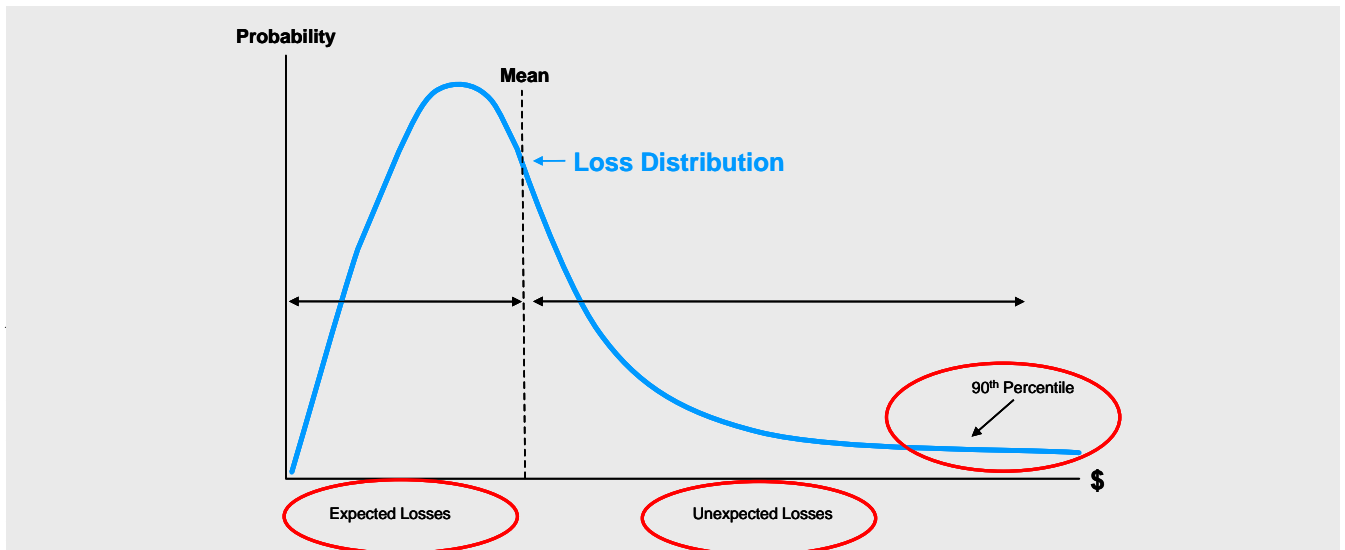


Exhibit 4.4 illustrates the common but incorrect conception of the terms expected loss and unexpected loss.

Exhibit 4.4 — Common Misconceptions: Expected Losses and Unexpected Losses



Unexpected Loss vs. Unexpected Losses: In risk management, just as there are no expected losses, there are also no unexpected losses. The unexpected loss is a specific number that represents the potential level of adverse deviation from the expected loss (the mean) up to the total exposure at the N% level (described below). The unexpected loss therefore measures the level of risk at the N% level. For example, as shown in Exhibit 4.3, if total

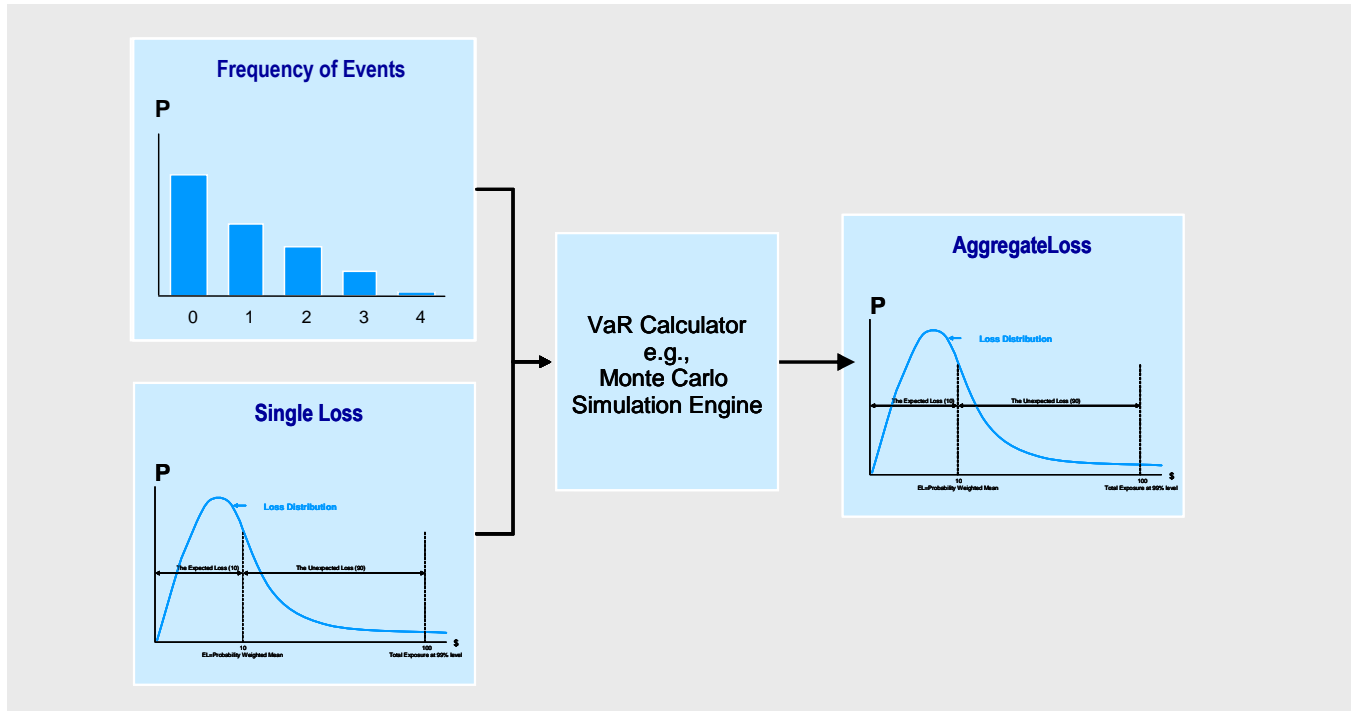
exposure at the 99% level is \$100 and the expected loss is \$10, then the unexpected loss or risk is \$90 (\$100 – \$10). However, many practitioners mistakenly confuse the term unexpected loss with “unexpected losses,” which they believe is used to describe the large losses or the losses above the mean (as shown in Exhibit 4.4).

Nth Percentile: As illustrated in Exhibit 4.3, the total risk exposure and the unexpected loss (risk) are always measured at a specific probability level. This is also the target risk tolerance level. For example, total exposure at the 99.9% probability level — with a one-year time horizon — represents the level of loss where a larger loss is expected to occur only once in a thousand years or has only a 0.1% chance of occurring in any given year.

The risk tolerance level is often set at the probability level associated with survival of the firm. For example, 99.9% tolerance level, with a one-year time horizon, indicates that the firm is only willing to tolerate a 0.1% (or 1/1000) chance of becoming insolvent in any given year. A 99% risk tolerance indicates a more aggressive risk profile. Here the firm is willing to risk becoming insolvent with a 1% (1/100) chance in any given year.

Exhibit 4.5. illustrates that the aggregate expected loss and aggregate unexpected loss are calculated by combining individual frequency and severity distributions. The frequency distribution shows the probability of events occurring based on a one-year time horizon. The severity distribution shows the probability associated with loss magnitude and has no time element. The aggregate distribution, which describes cumulative loss exposure for the specified time horizon, is generally derived through Monte Carlo simulation. These topics will be explored further in Section 8.

Exhibit 4.5 — Aggregate Loss Distribution



4.3 Risk Measurement and Assessment

Risk measurement and risk assessment are very similar concepts. Most risk frameworks do not draw a major distinction between these two terms. But many would agree that the term “assess” implies an estimation process, while “measure” suggests a more precise method of quantification.

Under Traditional ORM, however, risk measurement and risk assessment mean two completely different things, because they are based on two different and contradictory definitions of risk. Under Traditional ORM, risk measurement means estimating risk capital figures at specified probability level (e.g., 99.5%). Therefore, in a measurement context, the term is used in a manner consistent with the formal definition. However, when the term risk is used in an assessment context, risk = likelihood x impact. Again, likelihood x impact yields the level of expected loss, not the level of risk.

Under Modern ORM, risk measurement and risk assessment represent two ways of achieving the same objective – to calculate expected loss and unexpected loss figures. Only the method of calculation is different. In a measurement context, the process is generally based on the use of hard data⁸ and sophisticated methods while in

⁸ Hard data means empirical information that has been collected through a robust process. Soft data means empirical information that has been collected through some other reliable process. These types of data are described in more detail in section 8.3.2.

an assessment context the process is often based on soft data and expert opinion and/or a less complex set of calculation techniques.

Thus, under Modern ORM, the primary differences between risk assessment and risk measurement are the types of data used and the way parameters are derived. Where sufficient hard data are available, risk measurement is often more reliable than risk assessment. However, when sufficient hard data are not available, risk assessment based on soft data may produce more reliable results. Risk assessment techniques can also be used for conducting scenario analysis and stress testing.

It is important to recognize that developing a theoretically valid method of combining hard data and soft data is a very difficult process. In most cases, simply adding soft data to hard data is not theoretically valid and should be studiously avoided. This topic is explored further in Section 8.

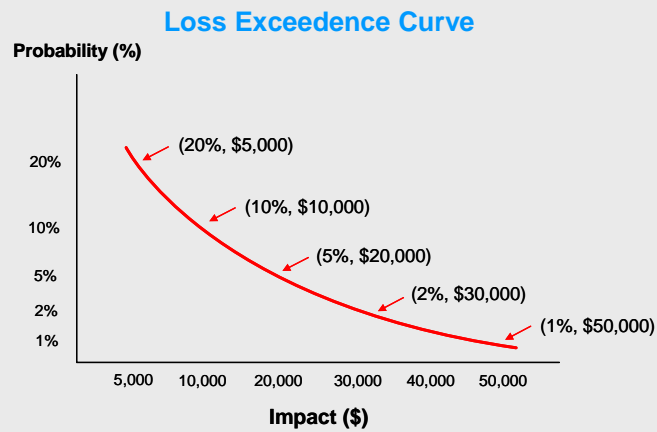
4.4 Risk Assessment/Measurement Under Traditional ORM

Under Traditional ORM, risk is assessed by multiplying likelihood and impact. As explained above, likelihood x impact is not equal to risk, but likelihood-impact analysis can yield metrics that may be useful for tactical decision making. However, the Traditional ORM method of expressing likelihood and impact as a single likelihood value multiplied by a single impact value is often not appropriate, for reasons explained in Exhibit 4.6 below.

Exhibit 4.6 — Likelihood and Impact

Consider the possibility of your having an automobile accident over the next year. You first apply likelihood-impact analysis and estimate that there is a 10% chance of having an accident that does more than \$10,000 of damage, but then you recognize that there is also a 1% chance of having an accident that does more than \$50,000 damage. In fact, there are an infinite number of likelihood values and corresponding impact values.

This full range of possibilities can be represented as a probability distribution — which is referred to as a single-event loss exceedance curve that shows the probability of at least one loss exceeding a given value during a specified time period.



The sum of all the different likelihood x impact combinations results in the probability weighted mean (mean severity). This is also referred to as the conditional expected loss — specifically, the expected loss conditioned on one event.⁹

In practice, under Traditional ORM, likelihood-impact analysis is generally conducted by calculating the likelihood (probability) of at least one event occurring¹⁰ multiplied by the average severity. However, there are problems associated with this type of analysis, because in most cases practitioners think of the impact in terms of the most likely outcome (the mode of the distribution) or perhaps the 50th percentile (median), not the true mean (the probability weighted mean). This point is further explained in Exhibit 4.7.

⁹ There are two ways of expressing the conditional expected loss. In this example, because the analysis is conditioned on one event, the event is assumed to take place. Thus, likelihood — by definition — is 100%. So likelihood x impact is 1 x mean severity, or just mean severity.

¹⁰ The probability of at least one event occurring can be calculated as 1 – Probability (0 Events) occurring.

Exhibit 4.7 — Problems with Traditional Risk Assessment

The following example explains how Traditional Operational Risk Assessment often provides misleading results.

Suppose there is a 50% chance of your experiencing a loss event over the next year, so:

$$\text{Prob. (1 Event)} = 50\%$$

And suppose severity has only two potential outcomes, such that:

95% of the time the loss is about \$.01 (essentially zero)

5% of the time the loss will be \$1,000,000

Therefore the mode of the severity distribution is about \$.01 and the mean is about \$50,000.

If you assume the commonly observed loss (the mode) represents the average impact, your analysis will be as follows:

$$\begin{aligned} \text{Likelihood x Impact} &= \text{Prob. (1 Event)} \times \text{Mode} \\ &= 0.50 \times \$0.01 = \$0.005 \end{aligned}$$

However, since the mean is a more pragmatic measure of the average loss, a more practical approach would be as follows:

$$\begin{aligned} \text{Likelihood x Impact} &= \text{Prob. (1 Event)} \times \text{Mean} \\ &= 0.50 \times \$50,000 = \$25,000 \end{aligned}$$

As can be seen from the above example, where severity is represented by a positively skewed (fat-tailed) distribution, as is common in operational risk, there is a significant difference between the routine loss and the mean loss. In such cases, likelihood x impact analysis generally yields no meaningful statistic.


The expected loss is a part of the cost of doing business. Using likelihood-impact analysis to estimate the expected loss, where impact is estimated as the commonly observed loss, can lead to artificially low cost estimates, and correspondingly, highly inflated profitability estimates. Furthermore, by using this information in risk-reward analysis, value-destroying investments can be made to appear profitable.

Given the fact that practitioners frequently estimate impact as the median or mode instead of the mean, one might ask why likelihood x impact is used at all. It turns out that for certain types of operational activities, losses do follow a Poisson frequency distribution and are characterized by a normal (or some other symmetrical) severity distribution. In this case, the mean, mode and median are virtually identical. These types of events are very common in manufacturing, transaction processing and other business activities involving large volumes of essentially identical trials with small variations in loss amount — the very areas for which Traditional ORM was designed. For most other risk management applications, however, event frequencies are not well-behaved and severity distributions are often positively skewed. Under these conditions, likelihood x impact analysis produces little or no actionable information — and often produces misleading information.

4.5 The Educational Challenge

In 2008, Towers Perrin conducted a survey of insurance industry CFOs. Exhibit 4.8 shows that according to their CFOs, the majority of insurance companies included in the survey (N = 31) have based their ORM frameworks on the traditional definition of risk.

Exhibit 4.8 — Towers Perrin CFO Survey Results (Operational Risk)



Eighteen of the 31 CFOs surveyed identified high probability (likelihood) and high impact (severity) events as those posing the greatest risk. This raises some interesting questions.

- If high impact events are considered to be those that threaten the ability of a company to execute its strategy (or perhaps even survive), can these occur with a high degree of probability in the first place? If so, would the company even reasonably expect to remain in business?
- If these events occur with a high degree of regularity (which one would expect given the high probability), wouldn't the company have implemented specific risk mitigation techniques to avoid loss, and/or have explicitly contemplated the losses produced by these events in its operating plans and budgets?

Improving the level of risk education must be viewed as a strategic imperative for all organizations across all industries. This also applies to regulators and rating agencies, many of whom have had difficulty consistently applying some of these basic concepts.

5. What is Operational Risk?

5.1 The Nature and Magnitude of Operational Risk

Operational risk, broadly speaking, is the risk of loss from an operational failure. It encompasses a wide range of events and actions as well as inactions, e.g., the failure to take appropriate action in a timely manner. When operational failures result in losses they are referred to as operational loss events. These losses include events ranging from unintentional execution errors, system failures and acts of nature to conscious violations of law and regulation as well as direct and indirect acts of excessive risk taking.¹¹

As previously indicated, virtually every catastrophic financial institution loss that has taken place during the past 20 years — including Barings Bank, Long Term Capital Management, Allied Irish Bank-All First, Société Générale, Bear Stearns, Lehman Brothers and American Insurance Group (AIG) — has been caused or exacerbated by operational failure.

Operational losses can be caused by junior staff; but they can also be caused by mid-level officers, senior managers, C level executives and Boards of Directors. They are sometimes caused by individuals and in other cases by groups of people working in collusion. The largest losses often take place when operational failures are present at the senior-most level. This might include situations where senior executives are themselves engaged in inappropriate risk-taking or even outright fraud, or perhaps more commonly, where executives intentionally overlook such actions by junior staff because they themselves are benefiting in the form of short-term financial incentives.

5.2 Wrong Turn: From Operational Risk to Operations Risk

Starting in the mid-1990s, following the news that several major financial institutions had experienced catastrophic operational losses, the leading banks began taking operational risk seriously. In fact, some of these institutions began allocating 20% – 40% of their total capital to operational risk. The importance of operational risk was formally recognized by bank regulators as early as 1999, when the Basel Committee for Banking Supervision (Basel Committee) published a consultative paper entitled *A New Capital Adequacy Framework*. In a later paper, the Basel Committee formalized its concerns, as follows:

“In recent years, supervisors and the banking industry have recognized the importance of operational risk in shaping the risk profiles of financial institutions. Developments such as the use of more highly automated technology, the growth of e-commerce, large-scale mergers and acquisitions that test the viability of newly integrated systems, the emergence of banks as very large-volume service providers, the increased prevalence of outsourcing and the greater use of financing techniques that reduce credit

¹¹ Barings Bank and Société Générale are examples of both direct and indirect excessive risk taking; the latter because senior officers consciously looked the other way.

and market risk, but that create increased operational risk¹², all suggest that operational risk exposures may be substantial and growing.”¹³

Around 2000, the Basel Committee decided that to draw attention to this major risk it was important for banks to separately reserve capital for operational risk, as well as for market risk and credit risk. However, this created a “double counting” problem because operational risk overlapped with the other two primary risks. After examining the capital allocation practices of a few large banks, in January 2001, the Basel Committee originally determined that the target capital level for operational risk ought to be about 20%¹⁴ of total bank capital. In retrospect, this 20% figure appears to have been too low, but it reflected the common perceptions of the largest banks at that time. Of course, in 2001 the vast majority of banks were largely unaware of the true nature and magnitude of operational risk, so general industry response was that this 20% figure was too high. In addition, because many banks were concerned that Basel II would eventually raise overall capital requirements, most were generally opposed to the operational risk capital charge.

Separately, the proposed introduction of Basel II also happened to coincide with another major piece of legislation — the Sarbanes Oxley Act (SOX). The original goal of SOX was to improve the integrity of the financial reporting process by creating greater transparency, but it soon took on a much broader interpretation. Faced with the prospect of numerous compliance initiatives, banks were under pressure to find cost-effective solutions. The argument that both SOX and Basel II ORM were essentially addressing the same issue and could therefore be jointly solved by implementing a framework based on the Traditional Approach was well received by the industry¹⁵. The fact that the Traditional Approach and Basel II were based on different and inconsistent definitions of risk and were designed to address altogether different business problems was generally ignored.

As bank regulators deliberated the specifics of the proposed Basel II regulations, they invited comments from the industry. Many industry groups, including rating agencies, software and data vendors, market and credit risk practitioners, etc., expressed concern that the introduction of this new overlapping risk would have an impact on the modeling of credit and market risk. They argued that any such decision would unnecessarily burden their firms because it would require them to reclassify all their data and recalibrate numerous metrics, such as credit default probabilities and loss given default values. In order to assuage their concerns and to simplify the modeling and data issues, the Basel Committee deemed that operational risk was a unique and distinct class of risk,

¹² This paper is prophetic in the way it warns of “financing techniques that reduce credit and market risk, but that create increased operational risk.” These are properties of collateralized mortgage obligations, bond insurance and credit default swaps. The inappropriate use of these financial instruments significantly contributed to the 2008 global financial crisis.

¹³ Basel Committee on Banking Supervision, *Working Paper on the Treatment of Operational Risk*, September 2001.

¹⁴ Basel Committee on Banking Supervision (Secretariat of the Basel Committee on Banking Supervision), “The New Basel Capital Accord: an explanatory note” (January 2001).

¹⁵ Skinner, Tara, “In Defense of AMA Methodology,” *OpRisk & Compliance* (February 2006).

independent of the other types of risk. In keeping with historical precedence, they also determined that credit losses driven by operational failure ought to be treated as credit losses for capital adequacy purposes.¹⁶ (Simply stated, operational risk + “pure” credit risk = credit risk.) And lastly, to ensure consistency with the other parts of the proposed new regulation, the Basel Committee, in September 2001, adjusted the target capital for operational risk downwards to 12%.¹⁷

As noted above, loss data reveals that operational risk is perhaps the most significant risk faced by financial firms. But the abovementioned Basel Committee decisions, which were based on precedence and expedience, set in motion a series of events that changed the definition of operational risk, not just within banking but in all other industries that followed suit. Soon banks began “*calibrating*” their data and models to produce low operational risk capital figures. These actions cemented the perception that operational risk was a minor risk. As a result, operational risk was transformed from a major risk to a minor risk. Shortly thereafter, it also moved from a front-office to a back-office issue and eventually became perceived as just *operations* risk.

Operations risk is only a subset of operational risk. Operations risk is characterized by unconscious execution errors and processing failures. Because these events stem from “normal” operating failures, the consequential single-event losses are relatively small — rarely in excess of a million dollars. Because risk is defined to be a measure of adverse deviation from the expectation, risk (the measure) is driven by the largest losses. Therefore, in Modern ORM, it is the “abnormal” operational failures — particularly conscious violations of a professional or moral standard and excessive risk taking that often result in sales practice violations, unauthorized trading acts and principal-agent events — that drive operational risk.

Modern ORM is concerned about operational losses. And because the top 1% of events account for about 60% – 70% of the total financial losses, under Modern ORM the largest losses are most relevant. Operations risks (such as execution errors and transaction processing failures) are a low priority issue in ORM because these small, frequent losses are well understood and can be managed through the ordinary audit/control process.

Ironically, because the pejorative conception of operational risk is now deeply engrained in normative risk standards, when a catastrophic operational loss takes place the knee-jerk response from many Traditional ORM practitioners is to say: “It’s a one in a hundred year event — so it won’t happen for another hundred years,” or “These types of events are beyond our control, so we just need to accept them as part of the cost of doing business,” etc. This demonstrates not only a fundamental misconception about the nature of operational risk, but also a deep misunderstanding of the real ORM business problem.

¹⁶ Basel Committee on Banking Supervision, “International Convergence of Capital Measurements and Capital Standards” (June 2004); Paragraph 673.

¹⁷ Basel Committee on Banking Supervision, “Working Paper on the Regulatory Treatment of Operational Risk” (September 2001).

6. Risk Architecture and Taxonomy

A prerequisite to managing risk is developing a comprehensive risk architecture and taxonomy (risk classification scheme). Classification is very important for management purposes. Therefore, one key criterion is that each set of risks must consist of like items that are relevant to management decision making. And in order for any classification scheme to be viable, the risk architecture must be based on scientific principles. In other words, there must be clear rules to support consistent classification.

6.1 The Traditional Risk Universe

Exhibit 6.1 shows a typical Traditional ERM risk universe (list of major risk issues) for an insurance company, which includes five top-level risks: credit, market, insurance, operational and strategic.

Exhibit 6.1 — An Example of a Typical Traditional Risk Universe

Credit	Market	Insurance	Operational	Strategic
<ul style="list-style-type: none"> ■ Default ■ Disputes ■ Sovereign ■ Downgrade ■ Settlement lag ■ Concentration 	<ul style="list-style-type: none"> ■ Equities ■ Concentration ■ Liquidity ■ Other assets ■ Basis ■ ALM ■ Currencies ■ Reinvestment ■ Interest rate sensitivity 	<ul style="list-style-type: none"> ■ Underwriting process ■ Basis ■ Mortality and morbidity ■ Pricing ■ Frequency and severity ■ Policyholder optionality ■ Reserve development ■ Lapse ■ Concentration ■ Product design ■ Longevity ■ Economic environment 	<ul style="list-style-type: none"> ■ Monetary controls ■ Distribution ■ Training ■ Financial reporting ■ IT systems ■ Turnover ■ Legal controls ■ Regulatory ■ Data capture 	<ul style="list-style-type: none"> ■ Competition ■ Rating downgrade ■ Availability ■ Demographic/social change ■ Customer demands ■ Technological ■ Negative publicity ■ Regulatory/political capital

Exhibit 6.1 represents an imprecise view of operational risk. Operational risk is generally viewed to be operations, execution or back-office processing risk, or the risk associated with routine employee misdeeds. In this illustration, operational risk does not include principal-agent risk, sales and business practices risk, or unauthorized activities risks, which are perhaps the most important operational risks. In addition, operational risk includes a category called legal risk, but legal (litigation) risk is an effect, not a type of risk. For example, sales practices violations could result in a lawsuit, but the lawsuit itself is not the risk.

Another example of imprecision is financial reporting risk, which could mean many things. For example, where a financial reporting loss resulted from an unconscious execution error, it would represent execution risk. Where it resulted from a deliberate act of wrongdoing (in which the perpetrator was trying to deceive the investor community), it would represent either business practices or fraud risk. These distinctions are very important for management purposes, i.e., associating risks with corresponding controls.

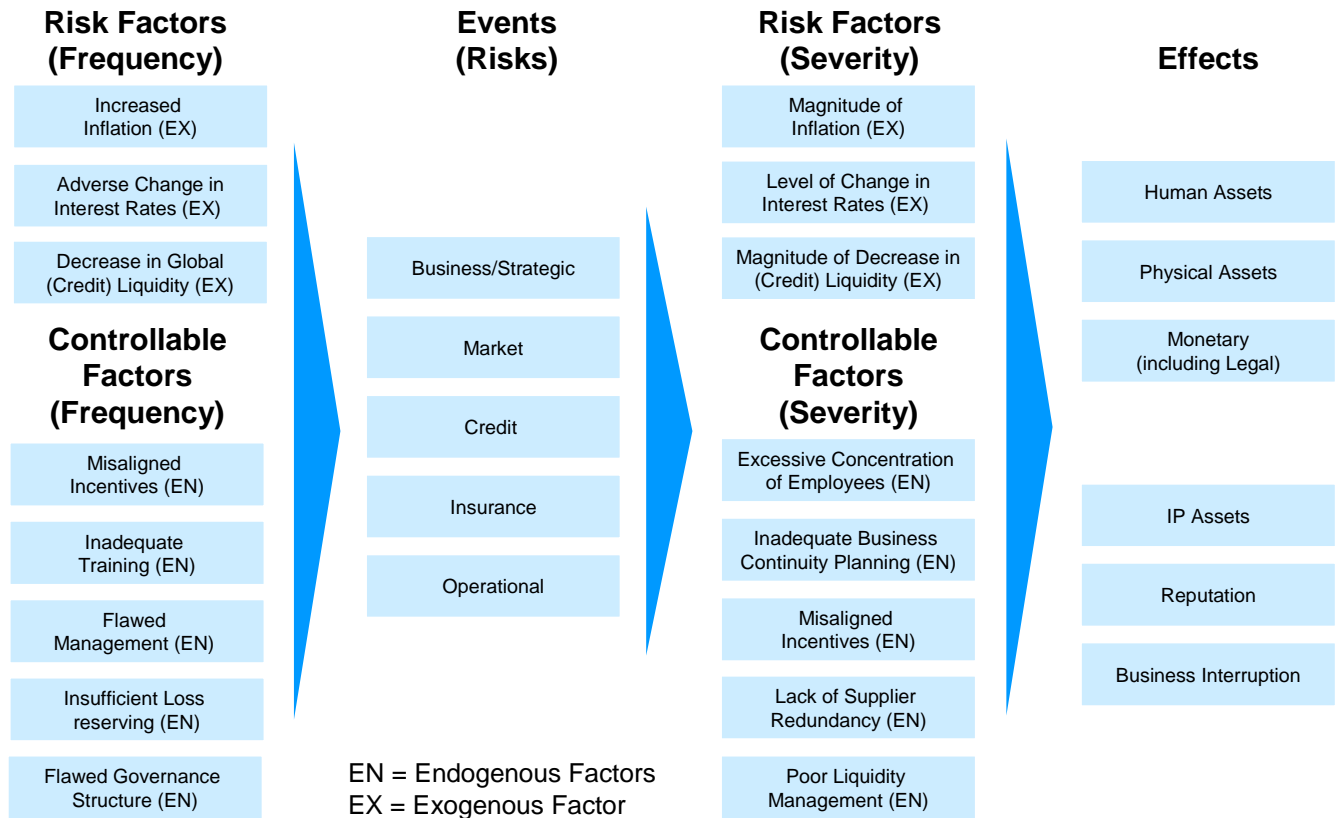
This illustration also fails to recognize that operational risks are embedded in the other risks. For example, concentration risk, which is listed under credit risk and also insurance risk, arguably represents operational risk. (The excessive concentration of insureds, assets, products or resources is an operational failure.)

Because of the complexity of this problem, the insurance industry has not yet been able to adopt a stable, uniform risk architecture/taxonomy. It is not uncommon for a company to introduce new risk taxonomy every two or three years, where the new approach is just another arbitrary allocation scheme. In order to be viable, a risk taxonomy must be based on a comprehensive understanding of the basic elements of ERM and their relevance to the ERM/ORM business problem. This topic is discussed in detail in the next section.

6.2 The Modern Risk Universe

The Research Team recommends that insurance companies adopt a risk universe with the following top-line risks: market risk, credit risk, operational risk, insurance risk and business/strategic risk. However, operational risk represents a special case. Because operational failures can manifest themselves in market, credit, insurance and business/strategic losses, operational risk permeates all aspects of the risk universe. This is illustrated in Exhibit 6.2. Therefore, operational risk is embedded in the other risks and must be recognized as such. Absent such an approach, event classification becomes arbitrary and does not adequately support management decision making.

Exhibit 6.2 — A Modern ERM Universe for the Financial Services Industry



© 2009 OpRisk Advisory and Towers Perrin

The Modern ERM Risk universe consists of four dimensions: risk factors, risks (events), controllable factors and effects (impacts). Where risk is defined to be the “risk of loss,” only events and effects can logically be considered categories of risk, because only these two classes are measured in terms of losses. However, classifying losses by effect has little management benefit, so few organizations use this dimension to define their risks.

Risk factors and controllable factors are referred to as contributory factors because they can both contribute to loss frequency and/or loss severity. For example, misaligned incentives, is a controllable factor; failing to align incentives with risk-adjusted performance can contribute to an increase in the frequency and severity of losses. Because controllable factors represent operational failure they naturally fall under operational risk. Exogenous factors, which are beyond the control an organization and do not represent operational failure, are considered risk factors. For example, interest rates are risk factors. Because changes in interest rates are not measured in terms of loss, interest rates are not risks in this context. However, because an adverse change in interest rates can lead to market losses, interest rates are risk factors (with respect to market risk). Liquidity can be a risk factor or a controllable factor, but it is not a risk because liquidity is not measured in terms of loss. For example an exogenous change in the macro-economic environment can cause a liquidity squeeze and lead to market and

credit losses; however, mismanaging the level of liquidity represents an endogenous operational failure (a controllable factor) and should be considered part of operational risk. It is very important to note that operational failures often manifest themselves in market, credit, insurance and business/strategic losses.

A simple example can help define and contrast these four dimensions. An electrical fire represents a potential event (risk) — something that could happen. Fires can be measured directly in terms number of events and associated financial losses. The insulation around electrical wires is a frequency factor. This is because insulation can prevent certain electrical fires — in other words insulation can contribute to a reduction in the frequency of fires. A sprinkler system is a severity factor. A sprinkler system will not prevent a fire from occurring, but it may help extinguish the fire more quickly and mitigate potential damage — in other words a sprinkler system can contribute to a reduction in severity of damage from a fire. Loss of physical assets and business interruption are effects/impacts, because they represent the way a fire can affect the value and financial viability of the firm.

As one can see, the Modern ERM architecture is logical and intuitive. This framework can be used to bring together competing and confusing SOX/audit/COSO classifications and definitions. For example, many auditors mistakenly define missing controls as risks. Under Traditional ERM events, risk factors, controllable factors and effects are all considered risks, and the associated loss data are lumped together without regard to distinctions between cause and effect. This not only causes communication problems, it also constrains the development of an integrated ERM framework.

It is very important to understand the subtleties of risk classification when developing a viable risk management framework. This is because for management purposes, it is very important to differentiate between ordinary market, credit, insurance and business/strategic losses and those that were driven or exacerbated by operational failure. Not doing so may obscure the underlying cause(s) of many of the largest losses. Without a clear picture of the causes of the major losses, one cannot implement an effective risk management/mitigation strategy. For example, treating the AIG loss as an ordinary insurance loss (because credit default swaps are insurance products) makes little sense from a management perspective. This is because in the absence of excessive risk taking (operational failure) these losses would have been minimal.

Since operational risk is not a unique and distinct risk, it may be more appropriate to measure operational risk in terms of its contribution to total loss. This can be accomplished by modeling the marginal contribution of operational risk to each of the other classes of risk. This topic is explored further in Section 8.10.

6.3 Modern Operational Risk Taxonomy

Modern ORM is based on a risk architecture/taxonomy that is designed to address the primary operational risk business problem — the management of the key operational risks. To accomplish this goal it is necessary to first identify the key risks and their core characteristics — *from a management (i.e., risk and control) standpoint*.

Events can be caused by people and by acts of nature. They can also be caused or exacerbated by negligence and incompetence or through conscious and deliberate acts of wrongdoing. In some cases the perpetrators may intend to benefit one or more parties; in other cases they may intend to harm one party and/or benefit another. And in certain other cases they may not intend to harm anyone, even though the expected outcome is one that would result in harm to another party. Classification is very important for management purposes, because it is very important to group together those events that have similar (homogenous) risk characteristics and corresponding controls. For example, the controls associated with inadvertent events are very different from those for events caused by conscious acts of wrongdoing.

A viable classification scheme needs to be able to express the key conceptual differences between the major classes of risk in clear and precise terms. A study of major operational losses reveals that *the most important factor one should use to delineate the risks and controls is intent*.

The Research Team suggests one possible set of top-level event risk categories could be the set described below. These risks are similar to the top-level Basel II risk categories.

- **Accidents:** Events that represent damage to human/physical assets, where human involvement, if any, was inadvertent (unintended). Example: Automobile accidents.
- **Acts of Nature:** Events that resulted directly from acts of nature. Examples: Hurricanes, floods.
- **Criminal and Malicious Acts¹⁸:** Events where the perpetrator(s) engages in a conscious act of wrongdoing, where he/she intends to benefit him/herself at the expense of another party. Criminal acts involve events where the perpetrator expects to receive a monetary benefit. Examples: Theft, fraud. Malicious acts involve events where the perpetrator also expects to benefit, but the benefit is of a non-monetary kind. Examples: Vandalism, terrorism.
- **Execution Errors:** Events caused by inadvertent human acts, excluding those events that primarily involve damage to human/physical assets. Examples: Transaction processing errors.
- **Principal-agent:** Events where the perpetrator(s) engages in a conscious act of wrongdoing, which may nominally benefit his/her firm, but which are not in the firm's best interest Example: falsifying or misrepresenting underwriting information to secure additional clients.
- **Sales, Business and Employment Practices:** Events where the perpetrator(s) engages in a conscious act of wrongdoing, where he/she intends to benefit the firm at the expense of a third party. Examples of Sales and

¹⁸ Criminal and Malicious Acts should generally be disaggregated. These subcategories should be further divided into Internal and External. (If even one employee was knowingly involved, it is treated as an internal event). The above mentioned categories could also be disaggregated into first order and second order events.

Business Practices include: Improper disclosure, account churning, mis-selling and underreporting firm income to tax authorities. Similarly, individual managers might believe their firm’s success in a particular market would be enhanced by taking actions that are detrimental to the rights of certain employees. Examples of Employment Practices include: promotion based on age/race/gender.

- **System Failures:** Events caused by hardware or software failures, where human involvement, if any, was inadvertent and incidental.
- **Unauthorized Activities:** Events where the perpetrator engages in a deliberate act of wrongdoing, but intends or expects to benefit all parties — at least in nominal terms. Example: Unauthorized underwriting, unauthorized trading, unauthorized approvals.

6.4 Criminal Risk and Principal-Agent Risk

One important aspect of Modern ORM is the distinction between “criminal” risk and “principal-agent” risk. Most people recognize criminal events, such as thefts or frauds. But to accurately define criminal acts one must describe the activity in terms of its payoff matrix. As can be seen from the payoff matrix in Exhibit 6.3, a criminal act is one where the perpetrator intends to benefit himself/herself and intends to harm another person to achieve his/her objective. Specifically, in order for the perpetrator to succeed in this “zero-sum” game, the counterpart *must* be harmed. Therefore, we define a criminal act as one where the perpetrator *intends* to benefit himself/herself and also *intends* to harm another party.

Exhibit 6.3 Payoff Matrix for Criminal Acts

Criminal		Intended Beneficiary			
		Perpetrator	Firm	Counterparty	No One
Intended Loss Sufferer	Perpetrator				
	Firm				
	Counterparty	X			
	No one				

© 2004 OpRisk Advisory

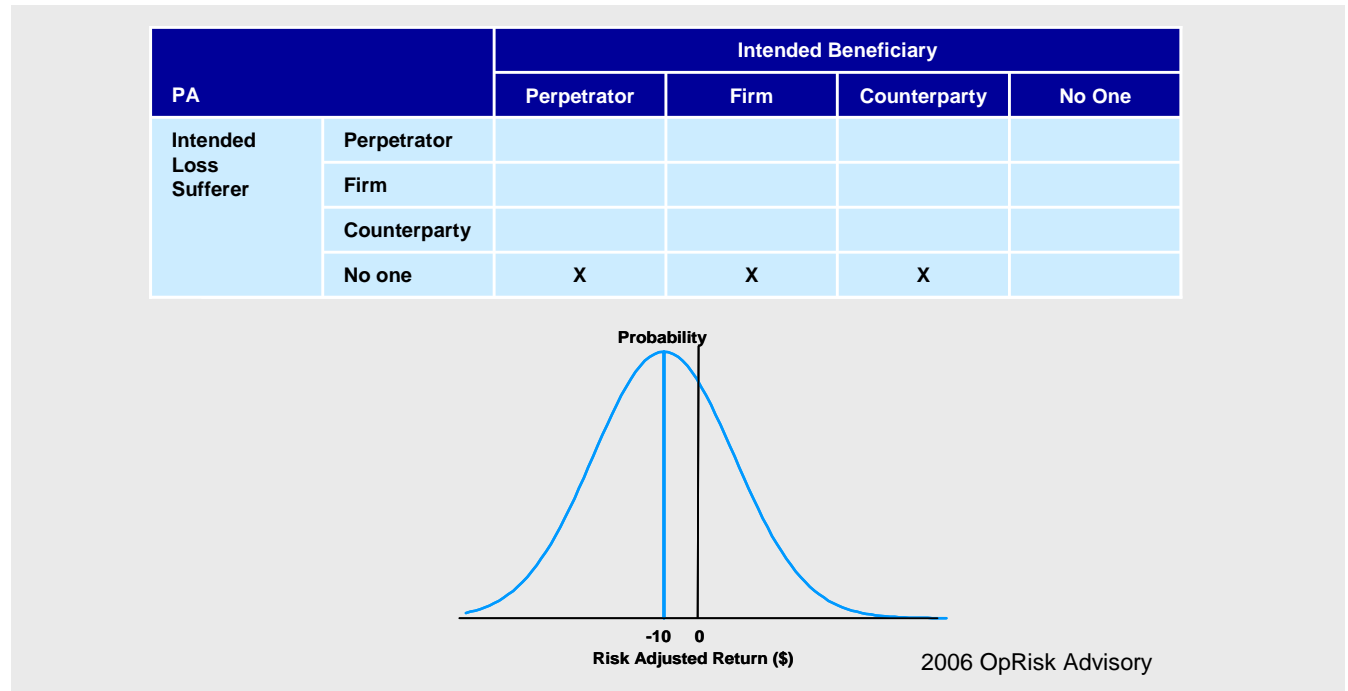
The difference between criminal risk and principal-agent risk is that in the latter case, the perpetrator does not intend to harm his/her firm. However, he/she knows (or should know) that the *expected outcome* of his/her actions is that the firm will be harmed, but he/she is unconcerned.

Principal-agent risk is logically represented in two ways: through a payoff matrix and a distribution of outcomes. The payoff matrix for principal-agent risk is shown in Exhibit 6.4 below. In this case the perpetrator intends to

benefit himself/herself, but he/she does not intend to harm anyone else. In other words, to succeed in this venture it is not necessary that anyone else be harmed.

The second illustration in Exhibit 6.4 shows the distribution of all expected outcomes for the agent's firm. Note that the expected outcome is a loss (which is identical to that of criminal risk). Thus, even though the perpetrator does not specifically intend to harm his/her firm, he/she knows that on a risk-adjusted basis the expected value of all outcomes is negative. In legal parlance, this may be referred to as willful or criminal negligence.

Exhibit 6.4 Illustrating Principal-Agent Risk



By virtue of its importance to ORM, principal-agent risk deserves to be formally recognized as a tier one event risk category. (Operational risk events are generally represented in a three tiered hierarchy.) Where information asymmetries and misaligned incentives exist, principal-agent risk is a key risk. This topic is explored further in Appendix A.

A classification framework based on the principles described above can be used to develop a risk taxonomy for any industry. For each industry, the risk categories could be expanded into a relevant hierarchical structure. For example, in manufacturing, where accidents are very prominent, this category would be broken down into numerous subcategories.

Developing a comprehensive set of Modern ORM event-based risk categories is beyond the scope of this project, but could be the focus of a future research initiative.

7. The ORM Business Problem

7.1 Roles and Responsibilities

As a general rule, the risk management function is not supposed to directly manage risk; that responsibility lies with business line managers, senior managers and C level executives. The role of the risk function is to facilitate effective risk management. Therefore, the primary goal of the ORM function is to do the following:

- Embed a risk culture that harmonizes the goals of key decision makers and external stakeholders.
- Provide the framework, infrastructure, tools and methodology to allow key decision makers to manage operational risk as part of their overall portfolio of risks, in conformity with cost-benefit analysis, within the risk tolerance standards of the stakeholders. Risk tolerance is determined by the Board of Directors.
- Ensure that there is transparency in the decision analysis process, such that independent observers can verify that key decision makers are in fact optimizing risk-reward, risk-control and risk-transfer in conformity with the risk tolerance standards of the stakeholders, i.e., mitigate principal-agent risk¹⁹.

Therefore, the ORM function's primary responsibility is to provide key decision makers with data, tools and techniques to accomplish the following:

1. Determine the magnitude of exposure to each major operational risk — in the context of the business's existing control environment — to confirm that it is in line with the risk tolerance standards of the stakeholder.
2. For each significant operational risk, determine whether the business has optimized the risk-control and risk-transfer relationships in the context of cost-benefit analysis. Given the business' overall portfolio of risks, determine whether one can implement any specific operational risk mitigation strategies that will cost-effectively improve financial performance on a risk-adjusted basis.
3. For new business opportunities, determine what impact investing in a new project will have on the business' level of operational risk and whether such a move is in the best interests of the organization. Specifically, determine whether the incremental profits are commensurate with the incremental risk.

¹⁹ Where the firm and the industry are systematically underestimating risk, principal-agent risk may exist. In such situations, the risk function needs to be prepared to take appropriate action to prevent a repeat of the circumstances that brought about the 2008 financial crisis.

7.2 Modern ORM in Practice

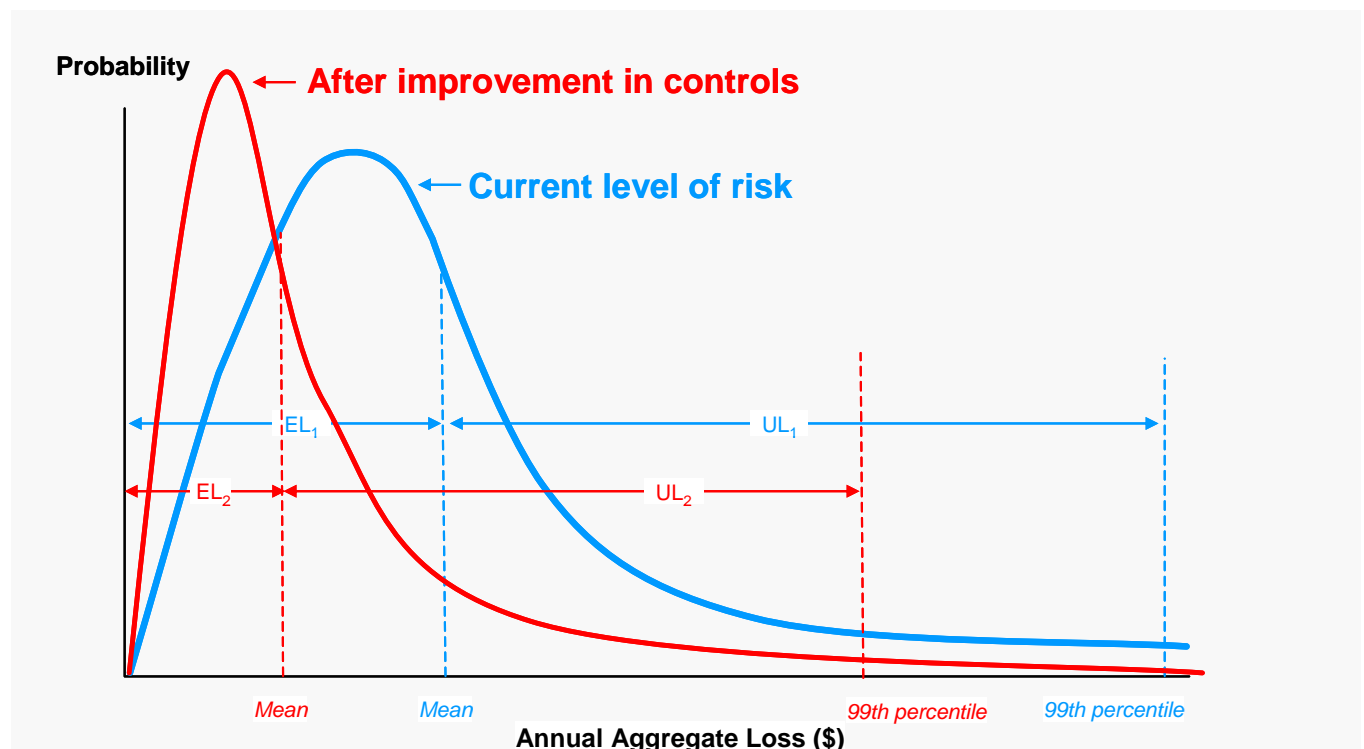
The first two subsections describe how key decision makers manage risk within a Modern ORM framework. The third subsection describes how the risk function can use these same methods to identify potential incidents of principal-agent risk.

7.2.1 Business Problem 1: Strategic Risk-Control Optimization

An executive vice president (EVP) is looking at past claim experience and observes that in the past five years the organization has experienced numerous fraudulent claims. By employing an actuarial model, using internal and external data, she estimates that her aggregate expected loss is \$20,000,000 and the aggregate exposure (99%) is \$100,000,000. (Therefore the aggregate unexpected loss = \$80,000,000).

The EVP subsequently learns about a new fraud prevention system designed to identify fraudulent claims much earlier in the process. She expects that by implementing this new system the aggregate expected loss will fall to \$5,000,000 and aggregate exposure (99%) will fall to \$25,000,000. (Thus, aggregate unexpected loss = \$20,000,000). Exhibit 7.1 shows the hypothesized change in risk profile with the new system.

Exhibit 7.1 — Risk-Control Optimization



The EVP needs two more pieces of information to make an informed risk-based decision: the fully amortized annual cost of the proposed new system and the firm's cost of capital, which are \$5,000,000 and 10%, respectively. All relevant information is summarized below. (In isolation, the unexpected loss is assumed to be the level of capital set aside for large losses.)

Exhibit 7.2 — Strategic Risk-Control Optimization Information Summary

Level of Loss	Current	Hypothesized
99% level aggregate loss	100,000,000	25,000,000
Expected loss	<u>20,000,000</u>	<u>5,000,000</u>
Unexpected loss (99%)	80,000,000	20,000,000

Cost of system	\$5,000,000
Cost of capital	10%

© 2007 OpRisk Advisory

The next step is to calculate the change in the Cost of Risk²⁰. The Cost of Risk is a measure of the economic value of the risk-taking activity and is calculated as follows:

$$\text{Cost of Risk} = \text{Expected Loss} + \text{Cost of Capital}^{21} \times \text{Unexpected Loss}$$

Therefore, the change in Cost of Risk that would result if the EVP were to implement the new system can be calculated as follows:

$$\begin{aligned} \Delta \text{ Cost of Risk} &= \Delta \text{ Expected Loss} + \text{Cost of Capital} \times \Delta \text{ Unexpected Loss} \\ &= (20,000,000 - 5,000,000) + 10\% \times (80,000,000 - 20,000,000) \\ &= 15,000,000 + 10\% \times 60,000,000 \\ &= 21,000,000 \end{aligned}$$

In the above example, the cost of the system (\$5,000,000) is less than the expected reduction in Cost of Risk (\$21,000,000). Therefore, in the context of cost-benefit analysis, it can be said that purchasing this system will optimize the risk-control environment, at the risk tolerance level of the stakeholders.

²⁰ The term Cost of Risk has many applications. For the purpose of this research paper, the term is defined as referenced above.

²¹ Cost of Capital should be a firm's target return, its borrowing rate, or its weighted average cost of capital.

7.2.2 Business Problem 2: Strategic Risk-Reward Optimization

The following paragraphs describe the process for optimizing the risk-reward relationship in an ORM context.

The senior manager of a major bank is considering launching a new loan product. Using a risk assessment model and soft data, the manager estimates the expected loss and unexpected loss associated with this new business venture. Based on this analysis he determines that at a volume of \$1,000,000,000 the expected loss is \$50,000,000, the unexpected loss is \$500,000,000 and the level of loss experienced in most years is \$5,000,000. The firm's cost of capital is 20%, and at the deal volume the estimated net profit (not including Cost of Risk) is \$100,000,000.

The senior manager's cost of risk and profitability estimates are summarized below.

	Accounting Profits	Risk-Adjusted Profits
Baseline	\$100,000,000	\$100,000,000
With only EL included	\$100,000,000	\$50,000,000
With total Cost of Risk included	\$100,000,000	(\$50,000,000)

The business manager realizes that on a risk-adjusted basis this is not a value-adding proposition, so he chooses not to pursue this business.

7.2.3 Business Problem 3: Analyze Principal-Agent Risk

The following paragraphs describe what the ORM function needs to do in a Modern ORM framework to mitigate principal-agent risk.

Because other firms are pursuing this business, the business manager in the above example decides to enter the fray. And because information asymmetries exist, he manipulates his model to produce aggregate expected loss and unexpected loss figures of \$5,000,000 (which is close to the observed loss or the mode) and \$100,000,000, respectively. He subsequently receives senior management approval to launch the product. Since this product appears very profitable, he makes this product his primary focus and generates \$5,000,000,000 in volume. He consequently earns a \$2,000,000 bonus each year for the next five years. In the sixth year macro-economic factors completely change the viability of this business. The firm loses \$3,000,000,000. The manager is asked to resign — which he does, but keeps the \$10,000,000 he has earned during the past few years.

Many of the largest financial institutions' losses have been caused by principal-agent risk, which is a key driver of operational risk. But it is unreasonable to assume that this risk can be managed through the same methods as one uses to reduce *operations* risk, which is driven by routine process failures. Instead, managing this risk requires a fully verifiable and transparent method of determining whether business managers are taking risks that are in conformity with the risk tolerance standards of the stakeholders, along with a set of disincentives that reduce the manager's associated payoff to a level below zero economic profit.

Therefore, one effective method of mitigating principal-agent risk would be to ask business managers to specifically incorporate the Cost of Risk into their profitability estimates. As a general rule, profitability and compensation should be measured on a risk-adjusted basis, not an accounting basis. All profitability estimates should be validated by *independent* risk experts. Finally, changing the incentive structure such that business managers share in the downside as well as in the upside may also prove useful. This topic is explored further in Appendix A.

7.3 The Modern ORM Infrastructure

Modern ORM requires a specific infrastructure. To create this infrastructure, the ORM function should do the following:

1. Explain to key decision makers the goals and objectives of ORM, how Modern ORM meets these objectives and specifically how Modern ORM methods can be used in practice. If this is not done correctly, it is unlikely that any of the other program goals will be met.
2. Facilitate the development of a viable risk management architecture, or more specifically, a mutually exclusive and comprehensively exhaustive operational event risk taxonomy. Under Modern ORM the risk taxonomy is designed to delineate risks in a manner that facilitates risk management
3. For each risk class, provide key decision makers with the ability to measure the expected loss and unexpected loss, in the context of the existing control environment, according to the risk tolerance standards of the stakeholders.
4. Assist key decision makers in developing/acquiring tools and methodology to assess and monitor internal control quality on a periodic basis.

A sample Modern ORM Risk Assessment template, which allows risk to be expressed as Expected Loss and Unexpected Loss, is shown below. (Note: The taxonomy shown in these examples is the Traditional Basel II taxonomy.)

Exhibit 7.3 — Example of Modern ORM Risk Assessment/Measurement Template

Business Unit A	Internal Fraud	External Fraud	Employment Practices and Workplace Safety	Clients, Products and Business Practices	Damage to Physical Assets	Execution, Delivery and Process Management	Business Disruption and System Failures	Total
Unexpected Loss	36,000,000	21,000,000	45,000,000	75,000,000	24,000,000	20,000,000	18,000,000	239,000,000
Expected Loss	4,200,000	3,500,000	5,000,000	4,000,000	3,000,000	12,000,000	4,000,000	35,700,000

Exhibit 7.3 shows a sample Modern ORM Risk Assessment/Measurement template, which allows for risk scores to be calculated in the context of the current control environment. This framework enables a risk manager to analyze and prioritize all operational risk classes in terms of both cost (EL) and risk (UL).

Exhibit 7.4 shows a sample Modern ORM control assessment template, where raw control scores are calculated using a specified set of control standards. These control standards are aligned with the risk tolerance standards of the organization at a specified confidence level (e.g. 99%). In this illustration, individual scores are based on a 1-9 scale. They are then weighted (by relevance) and aggregated on a normalized 0-100 scale, so that they can be incorporated into an integrated risk-control assessment template.

Exhibit 7.4 — Example of Modern ORM Control Assessment Template

Score	Control Name	Weak Attributes			Moderate Attributes			Strong Attributes		
		1	2	3	4	5	6	7	8	9
6	Segregation of Duties	<ul style="list-style-type: none"> Segregation of duties is managed by policy and procedures 			<ul style="list-style-type: none"> Capabilities limited through system profiles System profiles are not administered by independent team Profiles reviewed annually or less 			<ul style="list-style-type: none"> Capabilities limited through system profiles System profiles are administered by independent team Profiles reviewed semiannually 		
8	Ethics Code	<ul style="list-style-type: none"> An Ethics Code does not exist or exists but has not been effectively communicated to employees 			<ul style="list-style-type: none"> An Ethics Code has been established by management Employees are required to read and affirmatively note their acceptance of Ethics Code 			<ul style="list-style-type: none"> An Ethics Code has been established by management Employees are required to read and affirmatively note their acceptance of Ethics Code Anonymous Hotline has been established to report unethical behavior 		
4	Employee Activity Reports	<ul style="list-style-type: none"> Some employee activity reports have been developed but are not monitored consistently 			<ul style="list-style-type: none"> Some employee activity reports have been developed Reports are monitored periodically by management 			<ul style="list-style-type: none"> Critical employee activity reports have been developed Reports are monitored by an independent team on a continuous basis An escalation process exists 		

Source USAA FSB

Exhibit 7.5 shows a sample Modern ORM Integrated Risk and Control Template, where risk measures (expressed as loss exposure) and control measures (expressed as normalized scores) can be viewed concurrently in a logical and consistent format. Using this framework, controls improvements can be prioritized according to risk. In

addition, hypothesized improvements in control quality can be used to estimate the reduction in the overall level of risk.

Exhibit 7.5 — Example of Modern Integrated Risk and Control Assessment Template

Property/ Casualty Insurance	Internal Fraud	External Fraud	Employment Practices and Workplace Safety	Clients, Products and Business Practices	Damage to Physical Assets	Execution, Delivery and Process Management	Business Disruption and System Failures	Total
Current Risk	36,000,000	21,000,000	45,000,000	75,000,000	24,000,000	20,000,000	18,000,000	239,000,000
Change in Control Scores	50 55	65 65	70 72	53 55	55 60	70 75	64 68	59 61
Hypothesized Risk	32,000,000	21,000,000	44,000,000	72,000,000	22,000,000	19,000,000	17,000,000	227,000,000

7.4 Traditional ORM in Practice


For firms that have implemented Traditional ORM, the centerpiece of their framework is generally Traditional RCSA, which is generally implemented as follows:

- Identify your risks.
- Quantify your risks through likelihood and impact analysis before and after hypothesized improvements in control quality; risk level after controls are net risks or residual risks.
- Accept those net/residual risks that are within your tolerance.
- Develop actions plans to address those risks that need to be mitigated.

The Traditional Approach has many useful features. It provides structure, governance standards and a simple approach to risk identification and assessment. But it also has one very important flaw. It is based on a definition of risk that is inconsistent with the formal definition — the definition used in the risk management and actuarial profession. As shown in Exhibit 7.6, under the Traditional Approach, risk is associated with the average, or expected loss. Yet, under the formal definition, risk represents the unexpected or “worst case” loss.

This discrepancy has huge implications. Specifically, Traditional ORM fails to reveal the real operational risks. Instead, it typically focuses attention on the set of commonly observable threats and control weaknesses associated with routine losses — independent of risk. Therefore, the largest risks generally go unrecognized. Institutions that follow the Traditional Approach are often unaware of their most significant risks. In addition, organizations that base risk-control optimization decisions on the results of Traditional RCSA can easily become over-controlled in the areas where they have the least risk, but remain significantly under-controlled in the areas where they have the most risk.

Exhibit 7.6 — Traditional vs. Modern Conceptions of High Risk



The Traditional Approach is very effective for preventing losses at a tactical level, but loss prevention addresses only one aspect of the ORM business problem — and not the most important one. Specifically, Traditional ORM does little to mitigate exposure to the large catastrophic events, such as sales and business practices violations and acts of excessive risk taking, that are really the key drivers of operational risk.

Another practical concern with the Traditional Approach is the lack of cohesion. The risk identification process — naming your risks — appears intuitively appealing and can be useful when it is used to identify a few key imminent threats. But, it is very challenging and resource-intensive to implement across the enterprise. In particular, because risks overlap (risk factors, controllable factors, events and effects), and because the list of possible causes is virtually unlimited, it is possible for a conscientious practitioner to identify thousands of risks. Needless to say, it is very difficult to prioritize and/or actively manage such a large panoply of risks.

A typical RCSA under the Traditional Approach is shown in Exhibit 7.7 below. Three key drawbacks are described in red.

Exhibit 7.7 Example of Traditional RCSA Template

Risk	Description	Likelihood	Impact	Gross Risk	Control	Affect of Control	Net Risk
1	Adequate resources may not be available when required.	1					
2	New products may be introduced without adequate due diligence.	3	2	6	Comprehensive new product approval procedures are in place.	3	3
3	Product knowledge is concentrated within a small group of people, which may lead to significant IP loss if certain individuals were to leave the firm.	2	1	2	Each key staff member has a back up who has almost complete knowledge of all critical tasks.	1	1
4	Sales staff may inadvertently fail to fully disclose the risks associated with investment securities to retail customers.	1	3	3	All sales representatives are trained on the risks associated with investment securities.		
5	Contract terms for an investment banking assignment may be recorded incorrectly.	2	2	4	Entry and verification of all relevant information are performed by different individuals.	1	3
6	A change in law or regulation in Asia may cause the firm a loss.	3	2	6	The legal department conducts an annual risk review of all Asian countries to assess the chance of such a change occurring	2	4
7	Transaction volume may increase to a point that exceeds capacity levels and leads to transprocessing losses.	3	2	6	Senior management reviews resource and staffing requirement on a monthly basis.	4	2

Managers can identify the “risks” that they observe, but these threats generally don’t represent the major risks.

Likelihood x Impact is not equal to Risk; so the rank order is wrong.

These “risks” are a conglomeration of risks, risk factors, controllable factors and effects (so not mutually exclusive). Through this hugely resource intensive process, conscientious staff can identify thousands of “risks.”

8. Measuring/Assessing Operational Risk

8.1 Goals of Measuring/Assessing Operational Risk

The purpose of measuring and/or assessing operational risk is to solve the key business problems described in Section 7. To address these problems, key decision makers must know their business' exposure to each key operational risk, in the context of the business' existing control environment, at the target risk tolerance level. Specifically, they require the following information:

1. Total Aggregate Risk Exposure for each business at the target tolerance level (N%)
2. Aggregate Expected Loss, which is the probability weighted annual aggregate loss
3. Aggregate Unexpected Loss (N%), which is the Aggregate Risk faced by the business

— $\text{Unexpected Loss (N\%)} = \text{Total Aggregate Risk Exposure (N\%)} - \text{Aggregate Expected Loss}$

4. Total Cost of Risk.

Operational risk models can also be used to estimate regulatory and economic capital for operational risk. However, using a risk model only to calculate regulatory capital represents perfunctory compliance and will not meet the standards of the Solvency II or Basel II “use tests²².” Similarly, rating agencies will ultimately want to see evidence that risk models are being used to guide decisions and not simply for compliance.

8.2 The Actuarial Approach

As explained in the previous section, strategic risk-based decision analysis requires one to estimate aggregate risk metrics. However, directly modeling the annual aggregate distribution requires many years of data. This is because under such an approach the aggregate or cumulative loss for an entire year represents just one data point, five years of loss data yields only five data points, and so on.

To overcome the data paucity problem, actuaries decompose the aggregate loss distribution into its two integral components: frequency and severity. By doing so, each individual loss can be viewed as a single data point. This is the reason the process for modeling aggregate or cumulative loss distributions is often referred to as the actuarial approach.

Under the actuarial approach, frequency represents the number of events, and severity represents loss magnitude per event. Therefore, we can represent aggregate loss as follows:

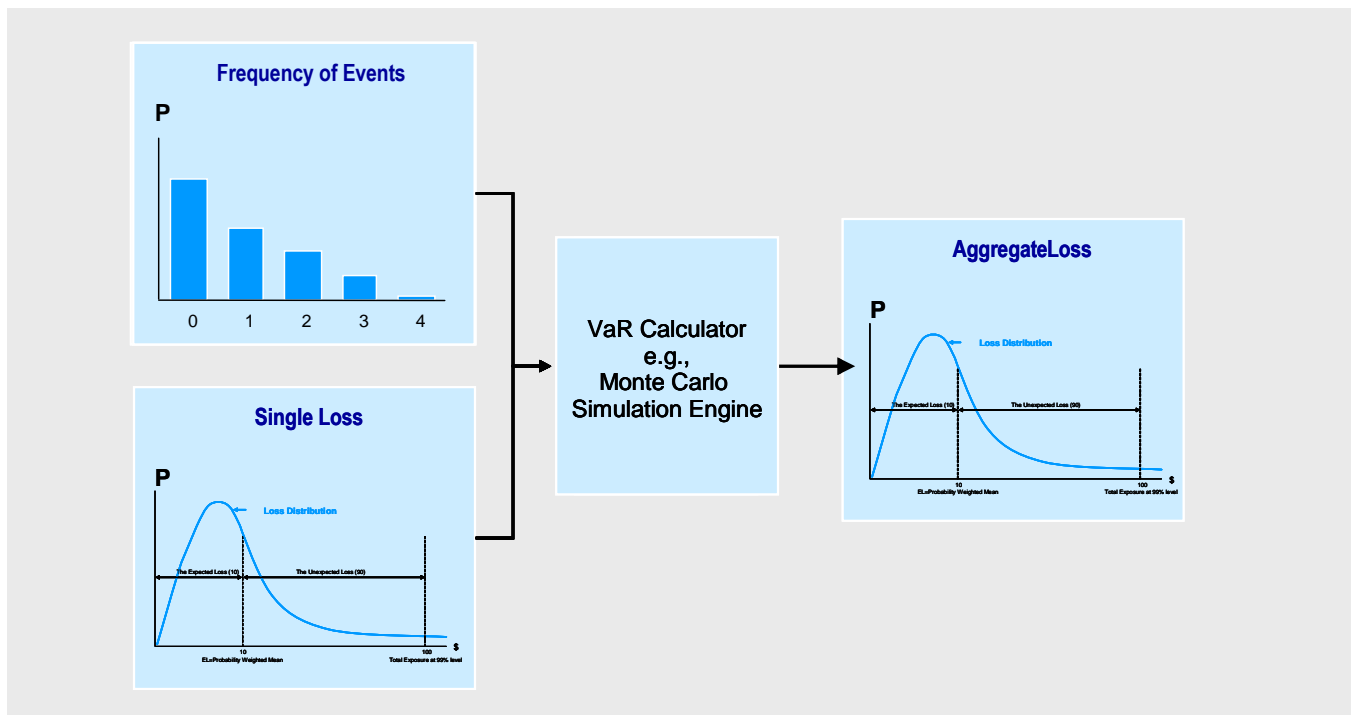
²² At present, the standards for the use test are low or undefined, but it is likely that even after these standards have been introduced, they will become more stringent over time.

$$\text{Aggregate Loss} = \sum_{i=1}^N X_i$$

where each X_i represents a single event loss and N is the number of events.

Here N is often referred to as the frequency, and each X_i is referred to as the loss severity. A distribution of N is called a frequency distribution. Similarly, a distribution of X_i is referred to as a severity distribution. One frequency and one severity distribution are needed to create an aggregate loss distribution for one risk class, as shown in 8.1 below.

Exhibit 8.1 — Aggregate Loss Distribution



Although there are no restrictions on the types of frequency and severity distributions one can use, a theoretical requirement is that the X s be independent and identically distributed (i.i.d.) and that the frequency and the severity be independent.

While mathematical or analytical approaches are available to determine the aggregate loss distribution directly for many combinations of frequency and severity distributions (see Panjer or Heckman-Meyers, for example²³), it is nonetheless sometimes difficult to describe the aggregate loss distribution in mathematical form, even when the

²³ Panjer, H.H. (1981). Recursive evaluation of a family of compound distributions. ASTIN Bulletin 12, 22-26; Heckman, Philip E. and Glenn G. Meyers, “The Calculation of Aggregate Loss Distributions from Claim Severity and Claim Count Distributions,” PCAS LXX, 1983, pp. 22–61.

frequency and severity distributions are both well known, and particularly when dealing with empirical distributions. As a result, numerical methods such as Monte Carlo simulation are often used to calculate the aggregate distribution and corresponding Expected Loss and Unexpected Loss figures.

There is a relatively simple relationship between the mean of the Aggregate Loss distribution (the Aggregate Expected Loss) and the frequency and severity distributions, as expressed below:

Aggregate Expected Loss = Mean Frequency x Mean Severity

For example, if Mean Frequency is ten events per year and Mean Severity is \$100,000

Aggregate Expected Loss = 10 x \$100,000 = \$1,000,000

8.3 Data Requirements

8.3.1 Internal and External Data

One common misconception in ORM is that internal modeling means modeling primarily with internal loss data. External data are almost always necessary for modeling severity and are often needed for modeling frequency as well.

The use of external data in operational risk is based on the assumption that companies in an external peer group are very similar to the firm being modeled, i.e., that these institutions all have a similar risk/control profile. In such cases, loss data from external institutions can be deemed relevant for risk analysis. To use external data, one must also assume that operational failures are independent, such that ten years of data from 20 companies is roughly the equivalent of 200 years of data from one company. Where the goal is to assess or measure risk at the 99% level, the top few losses from a relevant 200 company-year data set is much more valuable than even a million hard data points from one institution collected over a five-year period. Therefore, for operational risk assessment/measurement, data requirements should be defined in terms of years of data, rather than number of data points.

8.3.2 Hard Data and Soft Data

There are two kinds of data, hard data and soft data. Hard data is empirical information that has been collected through a systematic process on a prospective basis. Soft data is information that is based on empirical observations, but where the data may not have been collected through a robust process and/or where the data may represent a proxy variable. For example, suppose you were to set up a scientific process for measuring the height of ocean waves at a certain beach. The data you would collect from this process – as of today – would be hard data. If you were to collect this data for the next five years, you may end up with several million hard data points. Alternatively, if you were to check the geological record, you may discover scientific evidence of tsunamis during the past 500 years. While you may not be able to estimate the exact height of these tsunamis, assuming you were to use advanced scientific methods (e.g., measure how far inland the waves traveled) you could make reasonable

inferences about the height of these waves. This data would be soft data. However, this process may yield only five or ten data points.

So, which data is more useful for risk analysis? If your goal is to estimate a one in a hundred year event (i.e., 99% level wave, with a one-year time horizon), hard data may be almost irrelevant even when you have millions of data points. For certain kinds of analysis hard data are essential, but for risk analysis sometimes soft data can be much more valuable. Many models rely exclusively on hard data even when this type of data is known to be insufficient. To a large extent this explains why we witness one in a hundred year events every 15 to 20 years. For a concrete example refer to Section 8.6.5.

8.3.3 Outliers

In most statistical analysis, where the goal is to understand central tendency, the mean is not a reliable measure, because the mean is affected by outliers. Therefore, most statisticians prefer using the median (the middle value) or the mode (the most common observation). However, because the mean has many useful properties, many statisticians still use this metric. But to make the mean a better representation of central tendency, it has become acceptable to throw out the outliers.

In risk analysis the opposite is true, particularly in ORM where the major risks are characterized by large rare events. In operational risk modeling it is the outliers that are most relevant. As one can see from the tsunami example in Section 8.6.5, in some cases the so called outliers are the most relevant data. In that case all the “non-outliers” could have been thrown out.

The inability to reliably differentiate outliers from relevant data in small samples is perhaps the best argument in favor of using external data. For example, suppose you collected internal loss data for three years and have 1,000 losses in your database. And also suppose all the losses are in the \$10,000 to \$1,000,000 range, except one loss that is over \$100,000,000. If you include this loss in your data set, it might cause you to significantly overestimate both the expected loss and the unexpected loss estimates, unless a \$100,000,000 loss should be expected to occur in any given three year period. However, by excluding this event your estimates will almost certainly be too low, since the fact that the event occurred in the past suggests there is a non-zero likelihood of a similar event occurring in the future. Only by supplementing internal data with external data from several peer institutions will the true shape of the tail become evident. This will help you estimate how often a \$100,000,000 loss is likely to take place.

8.3.4 Homogeneity and Data Classification

As discussed above, in order to model losses with a small data sample, it is imperative that the data be identically distributed. As we have already seen, if you combine ordinary wind-driven waves with tsunami waves in a grand set called ocean waves, the data will exhibit a multi-modal shape and will be difficult to model. The same would be true if you were to combine routine execution errors and unauthorized activities in one grand set called operational risk. Accurate data classification is critically important for operational risk modeling.

8.3.5 External Data Sources

There are two good sources of external loss data: external public data and consortium data. Both have advantages and disadvantages. In addition, there are significant differences in size and quality of these products and initiatives. Where the goal is to manage operational risk under a Traditional Approach, external data requirements are not very stringent. Under the Traditional Approach, external data are used only for qualitative analysis — to support awareness and to identify new threats. However, under Modern ORM good quality external data are of critical importance.

External public data are data that have been collected from publicly available information, such as media reports, legal settlements and judgments and corporate filings. These data tend to be well documented. Some vendors provide relevant information for scaling losses. One problem with public data is that it suffers from a reporting bias, in that not all losses are reported. This is particularly true for the smaller losses (under \$10,000,000). Because of the reporting bias, directly modeling these data using traditional or modified Maximum Likelihood Estimation (MLE) methods is not advisable. However, using the largest losses, which are less prone to reporting bias, as “soft data” in a risk assessment process may be feasible — depending on the quality of the data. This topic is discussed further in Section 8.8.

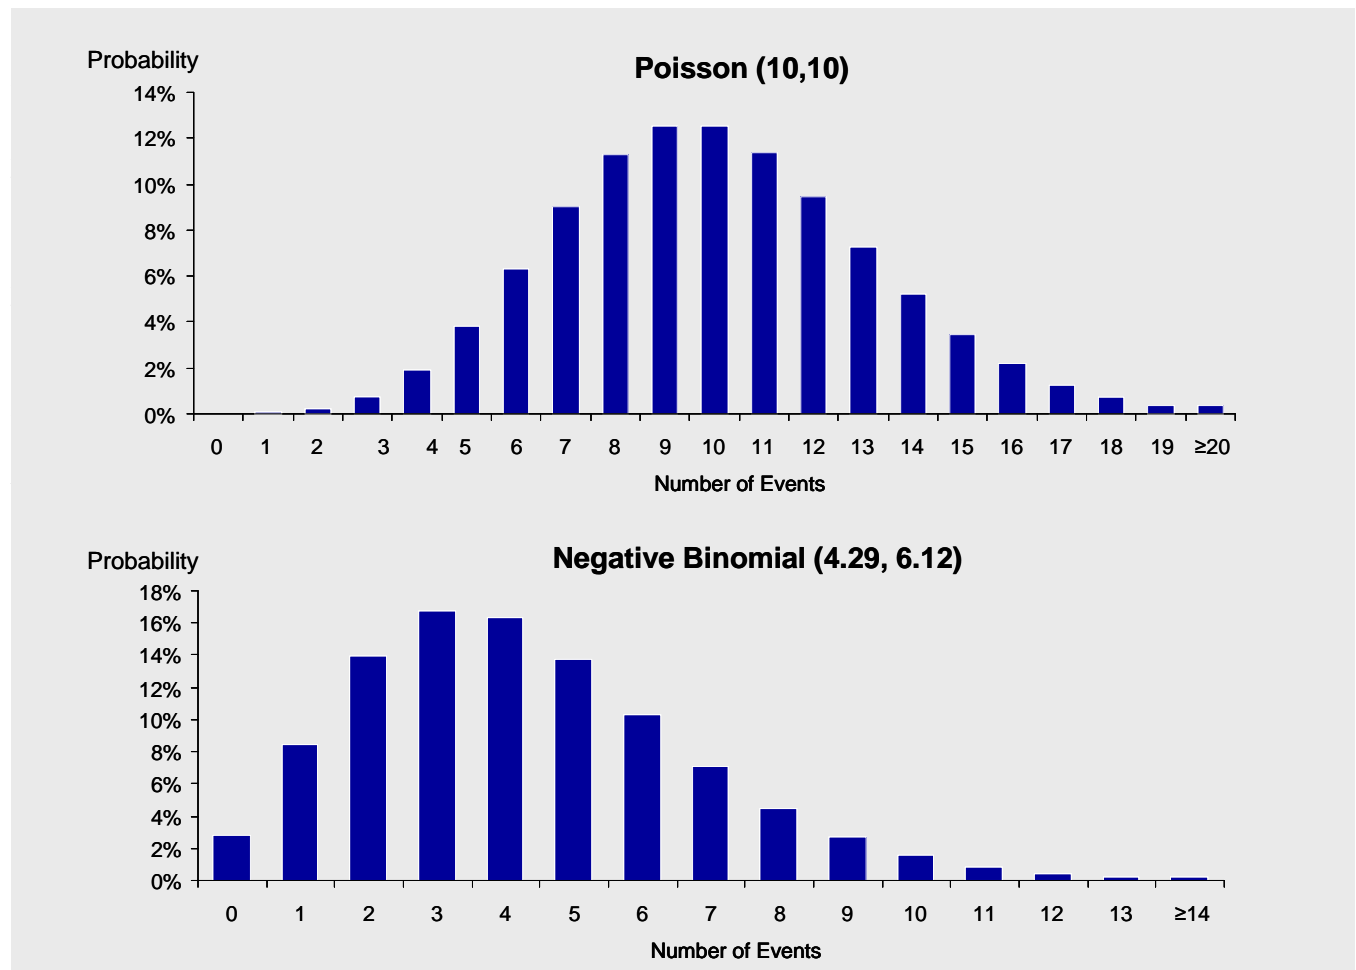
Consortium data represent pooled data from member institutions who have agreed to share their internal loss data on an anonymous basis. Some of these initiatives have a large member base. One common problem with consortium data is that, to preserve anonymity, very little descriptive information, if any, is provided. This presents a challenge, because where the data may have been misclassified or even classified under a different scheme it difficult to use the data. However, with effective controls over data quality and consistency, consortium data can be a very useful source of information.

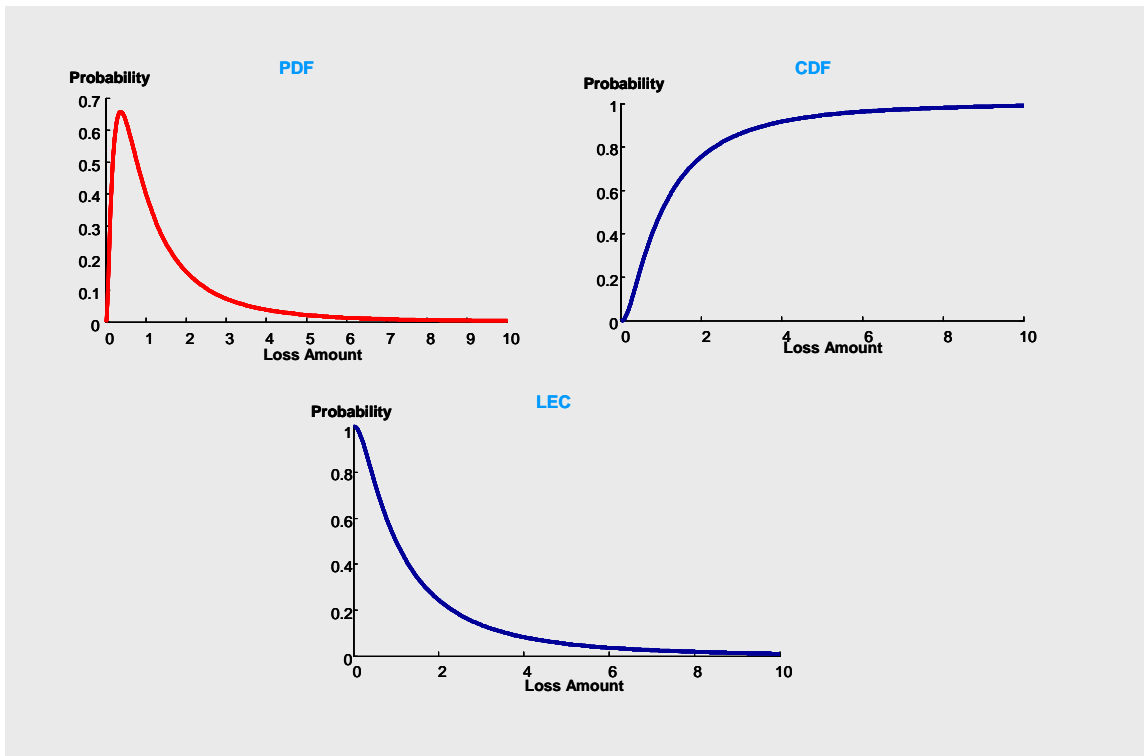
8.4 Frequency and Severity Distributions

8.4.1 Frequency

Frequency refers to the number of events that occur within a given time period. Empirical evidence suggests that events tend to follow a Poisson process. For such processes, distributions for the number of events in a fixed period of time include the Poisson, binomial, and negative binomial. Figure 1 shows a Poisson distribution with a mean of 10. Figure 2 shows a negative binomial distribution with an r of 10 and p of 0.7. This corresponds to a mean of 4.29 and a variance of 6.12. The Binomial distribution with a mean of 10 would appear virtually identical to the Poisson shown below.

Exhibit 8.2 — Examples of Frequency Distributions: Poisson and Negative Binomial





8.4.3 Annualized Loss Exceedence

As mentioned above, a severity distribution has no time element, so the probability values are not related to time. This means one cannot use the severity distribution alone in any form (PDF, CDF or LEC) to estimate the probability of a given loss occurring during a given time period.

By combining an annual loss frequency distribution and an LEC, one can create an annualized LEC. An annualized LEC shows the probability of a loss exceeding any given value within one year. Annualized LECs can be very useful in risk-based decision analysis, because the annualized LEC represents probability information in a manner that has intuitive meaning, i.e., it describes the level of loss associated with a one in N year event.

8.5 How to Model Frequency

In an actuarial model, a frequency distribution is a stochastic distribution, expressed as a discrete probability density function, where the X values consist of non-negative integers $\{0, 1, 2, \dots, n\}$.

Significant empirical evidence supports the use of the Poisson distribution for modeling loss frequency. Specifically, it can be shown that if the frequency satisfies the so-called Poisson postulate, then it will also follow the Poisson process. Although any distribution of non-negative integers can be used for modeling loss frequency, three distributions are commonly used. These are the Poisson distribution, the negative binomial distribution and the binomial distribution (mathematical representations are shown in Appendix B).

The Poisson distribution has a unique attribute: the mean of this distribution is equal to its variance. Thus, the Poisson distribution is effectively a one parameter distribution. Modeling annual frequency using a Poisson distribution requires much less data than does modeling with many other distributions because for the Poisson one needs only enough data to estimate the mean — the average number of events expected to take place in a year.

One can estimate mean annual frequency by dividing the number of events observed by the number of years in the observation period. For example, 50 observations over a five-year period indicate a mean annual frequency estimate of ten events per year (50/5). Alternatively, 200 observations over the same five year period would result in a mean annual frequency estimate of approximately 40 events per year (200/5).

Mean frequency can also be estimated using either the method of moments or maximum likelihood estimation (MLE) method. In property and casualty actuarial work, the Poisson distribution is often selected for modeling loss frequency because of its convenient properties, especially where the mean and variance are expected to be approximately equal.

The variance of a negative binomial distribution is always greater than the mean, a feature that differentiates the negative binomial distribution from the Poisson distribution (recall that the Poisson mean is always equal to the variance). Where empirical evidence suggests that there is excessive variability in the frequency relative to the mean, one should consider using the negative binomial distribution rather than the Poisson.

The variance of a binomial distribution is always less than the mean. Where empirical evidence suggests that there is very low variability in the frequency, one should consider using the binomial distribution rather than the Poisson.

8.6 How to Model Severity

Modeling severity is the most challenging aspect of modeling operational risk. This is true for several reasons, including the lack of sufficient data, the poor quality of data, the existence of truncated or censored data, sensitivity to low probability/high severity loss events, classification issues (commingling of loss data from different non-homogenous distributions) and the need to incorporate both internal and external data. Because the Expected Loss and Unexpected Loss results are very sensitive to changes in loss severity, accurately modeling severity is the most important task in modeling operational risk.

8.6.1 Simple Empirical Methods — Lack of Sufficient Data Points

Let us assume the target risk tolerance level is based on the Solvency II requirements, which are 99.5% for a one-year time horizon. This equates to the level of aggregate loss associated with a one in two-hundred year event²⁴.

²⁴ As previously mentioned, severity has no time element. The term “one in two hundred year event” is used to convey the relative frequency with which a loss of a given size might be expected to occur, and does not ascribe a

Accurately estimating this figure, using raw empirical analysis, requires about 1,000 years of relevant loss data — and the data must be drawn from a static risk/control environment. In 1,000 years this approximate level of loss would be exceeded five times. Thus, the 99.5% loss would be estimated as the fifth lowest loss. With exactly 200 years of data the 99.5% loss would be the largest loss. Thus, even with 200 years of data, the entire analysis would rely on a single data point — and would be prone to serious estimation error.

8.6.2 Fitting Data to Distributions

Many organizations model risk by fitting data to common parametric probability distributions. By using such methods one can significantly reduce data requirements. Provided that the loss data are independent and identically distributed, one can model a 1/200 year event with only a few years of relevant loss data. This can be accomplished by using a probability distribution to extrapolate the loss value at the target tail probability level — assuming, of course, one has sufficient data to determine the form of the distribution and reliably estimate the required parameters.

This method is reliable when the loss data are of good quality, the severity distribution used for the analysis is a “good fit” and the data are identically distributed. Where this is so, techniques such as MLE and statistical goodness of fit tests can be used to select the best fit severity distribution and parameters. Under these circumstances risk modeling is a very straightforward process. Unfortunately, this is rarely the case, because numerous data-related issues make operational risk modeling a very difficult task. A viable operational risk model must address these data issues in a theoretically valid manner. Where this is not done the results may be of limited value.

8.6.3 Fitting Truncated Data

As mentioned above, modeling operational losses presents many challenges. Very rarely are losses collected at a \$0 threshold. As a result, many common distributions, such as the lognormal distribution, which are specified from the \$0 threshold, cannot be used — at least not in a conventional manner. This is because the lognormal distribution, which is parameterized in terms of mean and standard deviation, requires that the mean and standard deviation be specified as if the data were collected from \$0. Where loss data are collected from a non-zero threshold, the mean is biased upwards and the standard deviation is biased downwards. Using these biased parameters as if they were the actual parameters would produce meaningless result.

There are two methods of addressing the truncated data problem. One is to use severity distributions, which are designed to be used with data collected from a non-zero threshold, such as the Generalized Pareto distribution (GPD). However, this distribution has a very heavy tail and often overestimates the level of risk. In addition, because the rate of decay of the tail is so low, the tail of GPD may not converge to zero, in which case the GPD has an infinite mean. As a result, the GPD can be used to estimate the Aggregate Expected Loss only by

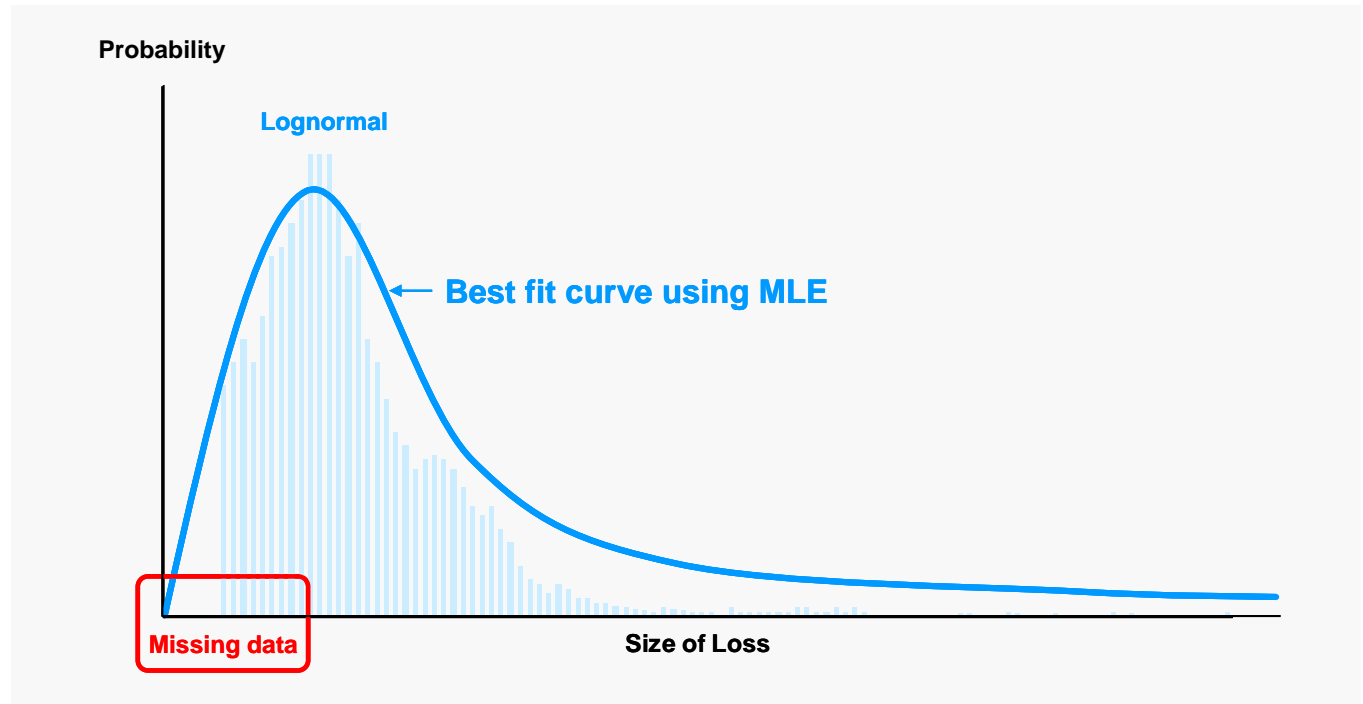
time element to loss severity that does not otherwise belong. This type of description is common in insurance and reinsurance and is referred to as a “return period”.

incorporating an upper bound. The subjectivity built into this process is problematic. Consequently, the Pareto family of distributions is not appropriate for modeling operational risk.

Another approach is to use less heavy tailed distributions and a modified MLE process — one specifically designed for fitting truncated data. This approach is explained in Appendix B. It is very easy to test whether a modified MLE fitting routine is working. Simply generate loss data from a known theoretical distribution to see if the fitting routine recovers the original parameters. For example, generate 1,000,000 data points from a lognormal distribution, with parameters of mean = 10 and standard deviation = 3, from a \$0 threshold. Then, test the routine at different fit thresholds (e.g., \$0, \$10,000, \$100,000). The fitting routine should recover parameters very close to the original parameters (10, 3).

This process of fitting a theoretical distribution to truncated data is illustrated graphically in Exhibit 8.4 below.

Exhibit 8.4 - Fitting a Truncated Dataset Using Modified MLE



8.6.4 Fitting Heavy Tailed Truncated Data

In certain situations the above mentioned process breaks down and the fitting routine returns unrealistic parameters. This often happens when one tries to fit, for example, a lognormal distribution using a modified MLE fitting routine to heavy tailed, truncated data. In such cases the routine will typically return an unreasonably low mean and an unreasonably high standard deviation. This routine usually fails because even with a modified MLE routine, a traditional two parameter distribution generally cannot be used to model certain types of heavy tailed data. Unfortunately much of operational risk is characterized by heavy tailed truncated loss data.

To model such data it is generally necessary to use distributions that consist of more than two parameters (two degrees of freedom) — for example, the Lognormal-Gamma (LNG) distribution or the Burr distribution. Both the LNG and the Burr can be reliably fit using the modified MLE methods described above. One other innovative method can also be used. This involves fitting the tail portion of the data set to an annualized LEC (ALEC). (See Section 8.8. for more information.)

Modeling operational loss data using severity distributions that have more than three degrees of freedom may result in over-fitting, which can be problematic. In addition, because many of these distributions cannot be fit using MLE methods modified for truncated data, the fitting process may be prone to estimation error.

8.6.5 The i.i.d. Assumption

Modeling risk with only a few years of loss data is now very common, particularly in market risk management, but it is essential to understand the key assumptions underlying this approach. One critical assumption underlying all such models is the i.i.d. assumption — in particular, the assumption that the loss data are identically distributed. (The term i.i.d. stands for independent and identically distributed.) When this assumption is not valid, i.e., where the data are not homogenous (identically distributed), the models can produce spurious results unless appropriate modifications are made. Consider the following example.

Suppose you want to determine the height of the “one in a hundred year” ocean wave. You begin by collecting wave data. Suppose after a five year study you have gathered data on millions of waves; however, during this time period only ordinary wind-driven ocean waves have been observed. In order to model the overall “one in a hundred year” wave event you would have to assume one of the following: (1) that ordinary wind-driven waves are the only kinds of waves that exist, or (2) that wind-driven waves have the same properties as all other waves, including earthquake-driven tsunami waves and asteroid-driven tsunami waves — in other words all waves are i.i.d).

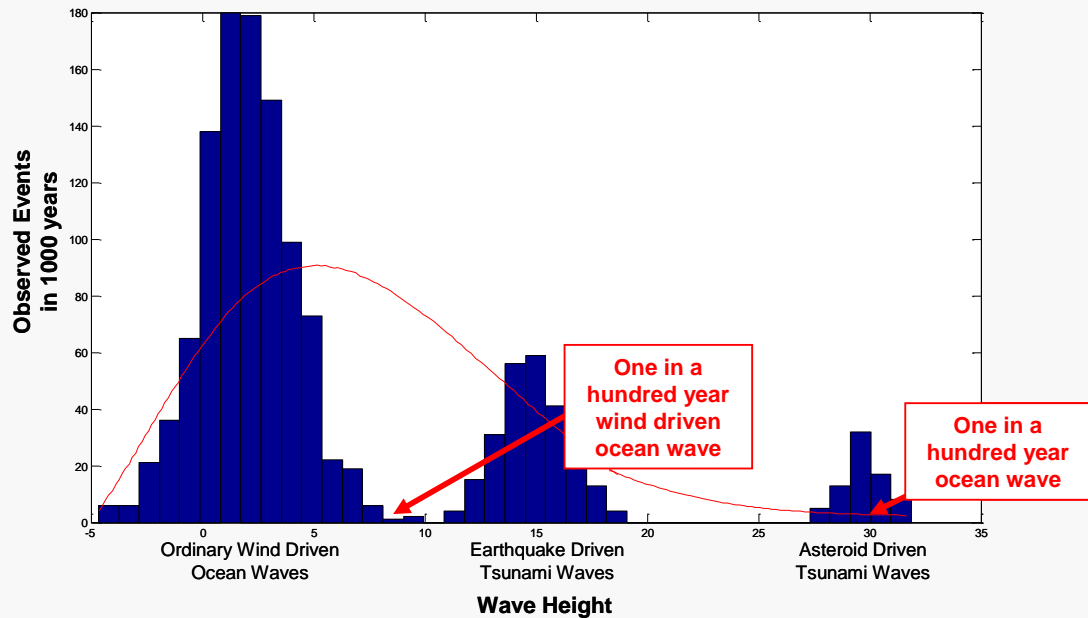
As illustrated graphically in Exhibit 8.5 below, both assumptions are invalid. And, as a result, your model would significantly underestimate the height of the largest waves. If you had access to 1,000 years or more of data, and this data included a representative sample of earthquake-driven and perhaps asteroid-driven tsunami waves, you would observe that the “one in a hundred year” wind-driven wave is significantly smaller than the overall “one in a hundred year” wave. (Note: In this figure, for illustrative purposes, wave heights are shown as normally distributed within each category of waves. In addition, the relative proportion of tsunami waves to wind driven waves is overrepresented.)

This generally explains why so many models systematically underestimate risk, i.e., why we observe a “one in a hundred year event” every 10-15 years. This issue is particularly relevant in market and credit risk modeling,²⁵ where the analysis is often based on huge amounts of data collected over a small period of time. It is important to recognize that one cannot reliably estimate long term volatility from short-term data unless the short-term data accurately represent the full range of possible outcomes.

As one can see from this example, where the goal is to measure risk at a 99% level for a one-year time horizon and where the data are not i.i.d., specifically when they are not identically distributed, millions of “hard” data points collected over a five-year period may provide less relevant information than five or ten “soft” observations over a 100-year period. This simple example illustrates the dangers of ignoring the i.i.d. issue in all areas of risk modeling.

²⁵ Flaws in market and credit risk models represent operational risk.

Exhibit 8.5 – A Graphical Illustration of Model Risk (Underestimating the Tail)



8.7 Combining Internal and External Loss Data

While internal data may be the most relevant data for modeling operational risk, it is insufficient for modeling heavy tailed data sets. In such cases, internal data is insufficient to estimate even the expected loss.

Various methods have been developed, each based on its own set of assumptions, to model operational risk by combining internal loss data, external loss data (data from other institutions) and expert opinion. Some methods are based on scientific principles and are feasible and practical to implement; many others are easy to implement but have questionable theoretical validity. In addition, many approaches reveal a fundamental misunderstanding of the basic business problem. Because many organizations view operational risk modeling as a compliance exercise, and because the bar currently established for achieving regulatory compliance is very low, virtually any method of combining internal and external data tends to be accepted in practice.

8.7.1 Frequency

There are a number of common methods of using external data for frequency, including using the data directly and using the data after scaling for size through a normalization algorithm or through proportionality.

Using external data directly, for example by using the external Poisson mean, can be problematic. This is because the average external firm may be larger or smaller than the institution being modeled. Using these data directly would be the equivalent of modeling the industry average institution.

Using scaled data may sound appealing, but deriving robust scaling algorithms can be a very challenging task, and there is often very little empirical data to support such analysis.

A straightforward and objective method of scaling external data is through proportionality. In this case, the institution estimates, for a specific business line, the Poisson mean for at least one risk category (provided there is sufficient data to estimate the mean, and assuming that the external source business line is similar in nature to the internal business line being analyzed). This is referred to as an anchor category. Then for categories where the organization does not have sufficient data, it uses external data to calculate the ratio of the Poisson means for the anchor category and other risk categories. The Poisson means for the other categories can then be calculated through proportionality, i.e., simple arithmetic. A simple example is shown below:

Exhibit 8.6 – Example of Frequency Proportionality

	Poisson Mean for:	
	Execution Errors	Business Practices
Internal Data	60	?
External Data	300	30

Since the ratio of Poisson means between Execution Errors and Business Practices is 300:30 based on external data, the institution can estimate a 10:1 proportional relationship for loss frequency between these two categories. This ratio would represent the inherent frequency relationship between Execution Errors and Business Practices. One can then extrapolate a Poisson mean for Business Practices as follows:

$$60 \times 30/300 = 6$$

Note: If the institution had used external data directly, the Poisson mean for Business Practices would be 30, which is unreasonable because it would not reflect the size of the target institution. Also note that the application of a simple method such as this requires many assumptions. In this example, we are implicitly assuming that all firms within a peer group have the same relative risk-control profiles across all categories.

8.7.2 Severity

There are several methods for combining loss data for severity. As with frequency one can use the data directly, after scaling for size using a robust scaling algorithm or through applying proportionality. Again, using external data directly can be problematic. Scaling for size may also be difficult. Proportionality appears to be the most objective method, but severity is not likely to be scalable using a simple linear scale factor. For severity, a more

complex method of proportional scaling is necessary. However, deriving such a methodology is beyond the scope of this paper.

Several other methods of combining internal and external data are common in the industry. Many of these methods are easy to implement, but most lack theoretical validity and offer little, if any, practical value. One such method involves selecting “relevant” losses from an external database and then modifying each individual loss to make it reflective of the particular firm’s internal environment. The adjusted loss is then added to the firm’s internal loss database. This method is invalid for the reasons given below.

First of all, with respect to loss severity, the purpose of collecting loss data is to determine the loss probability for each level of loss. For example, if 50 out of 1,000 losses in a particular data base are over \$100,000, then one can say, with respect to that particular organization, the probability of any given loss exceeding \$100,000 is 5%; and if 10 are over \$1,000,000 then the probability of exceeding that level is 1%. Consequently, losses carry two distinct pieces of information: the loss magnitude and its associated probability of occurrence — as measured by its proportionate representation in a loss database. Even without modification, once you remove a loss from its associated data set it loses the corresponding probability information. Thus, it loses all practical value.

Secondly, it is impossible to know how to adjust a loss to a particular firm’s control environment. For example, if one were to ask a large insurance company to adjust the recent \$100 billion+ AIG loss to reflect its own control environment, the answer would probably be unsatisfying. Finally, the process of subjectively adding data to a finite data base is theoretically invalid. Suppose you had a database of 1,000 data points, and none of the losses were over \$1,000,000, and then you added ten relevant external losses, all of which were from a database with a \$1,000,000 threshold. Doing so would likely increase probability mass in the tail region by roughly 1%. Suppose you then subscribe to another database and choose to add ten other data points over the same \$1,000,000 level. This would have a similar effect. As one can see, this process can easily be used to artificially manipulate probability mass until the model produces results that are in the “acceptable” range. Similarly, one could keep subtracting data points until the model produces the “right” answer. While this may be surprising, several large European banks use this method for combining internal and external data. What may be even more surprising is that their respective bank regulators, some of whom are considered leading regulators, have approved their models.

Another popular method used by several major banks is to use internal data for the body and external data for the tail. There are several variations of this approach, some of which use MLE and other fitting techniques. A key assumption in this approach lies in determining where the body (the internal data section) ends and the tail (the external data section) begins. In the absence of robust procedures for making this determination, organizations that use this approach are able to select the threshold in such a manner as to achieve their preferred result. Applied in this manner, this approach is almost identical to the flawed method described above, except that it offers the “appearance of validity” which makes it a matter of greater concern. Again, several banking regulators have approved models based on this methodology as well.

Within the banking industry, it appears that the scrutiny of models and data in most jurisdictions is such that as long as a bank is able to demonstrate the actual use of internal and external data in what appears to be an empirical model, they are likely to have their internal model approved. In many cases, the goal does not seem to be to try to understand what the data reveals about the organization's risk profile, but how to back into the "right" answer. This may to a large extent be a function of target capital range for operational risk, which bank regulators specified as 12%. Given the fact that virtually all of the largest bank losses have been driven by catastrophic operational failure, it is not clear how one could use actual loss data and a theoretically valid actuarial model and still produce operational risk capital figures in the 10% to 15% range.

8.8 Risk Assessment, Scenario Analysis and Stress Testing

As previously mentioned, under Modern ORM, the primary difference between risk measurement and risk assessment has to do with the types of data used and the way parameters are derived. Where sufficient hard data are available, risk measurement is more reliable than risk assessment. However, when hard data are not available, soft data and theoretically valid risk assessment methods may produce more reliable results.

Historical data are generally very useful, but in certain circumstances over-reliance on historical data can change behavior and bring about the very circumstances that cause the data to become irrelevant. For example, suppose a person who has never had an automobile accident believes in the infallibility of historical data. This person will naturally believe that he/she cannot have an accident, which could change his/her behavior such that he/she may start driving at 200 miles per hour.

One of the most important lessons of the sub-prime credit crisis is that historical data are only valid when they are representative of the current risk/control environment. Most real estate pricing models, which used historical data, calculated that a simultaneous housing price decline across all U.S. geographies was a zero probability event. This was based on loss experience going back 72 years. Recall that lax lending standards coupled with excess liquidity created the housing bubble. In this risk/control environment the historical data, which were collected at a time when credit standards were much higher and there were no additional factors contributing to a major housing bubble, were largely irrelevant. This topic is explored further in Appendix A.

Risk assessment methods are best applied when the goal is to determine risk at a 99+% level, for a one-year time horizon, and where hard data are not available, or if the hard data are not likely to be identically distributed. As we observed from the waves example in Section 8.6.5, millions of hard data points collected over a five-year period are much less useful than a few soft observations over a 100-year period. However, the critical question is, how can one use just a few data points to estimate the Expected Loss and Unexpected Loss levels? After all, fitting a severity distribution is supposed to take many data points — even when the data are i.i.d.

8.8.1 Risk Assessment Using Annualized Loss Exceedence

An innovative method for modeling risk with soft data involves fitting severity data not to the PDF, CDF or LEC, but to the annualized LEC. By fitting either hard or soft data directly to the tail portion of the annualized LEC,

one can mathematically derive the underlying severity distribution. This process requires expressing the loss data in the form of 1 in N year exceedences (or “return periods”). Since this involves fitting points to a curve, not a histogram of data to a density function, much less data are required — only one data point for each degree of freedom. This is a huge advantage when modeling the tail, because there are typically very few observations in this region. Furthermore, because one is fitting the tail directly, the resulting curve can usually provide a good representation of the tail. This is important because, for heavy tailed data sets, the body and tail cannot ordinarily be described using a simple two-parameter severity distribution. An illustrative example of this method of severity fitting is shown in Exhibit 8.7 below.

Exhibit 8.7 – Risk Assessment Under Modern ORM

Suppose one had observed the following 315 hard or soft data points collected over 300 Company Years (60 firms over 5 years), as shown below. One could then count the number of observations at each loss threshold (the two leftmost columns) and fit the data to an annualized LEC. The input and fitted output results are shown below (two rightmost columns):

Loss Threshold	Number of Observations	1 in N years Input	1 in N years Output
\$1,000,000	315	0.9524	0.9524
\$5,000,000	96	3.1250	3.0272
\$25,000,000	19	15.7895	15.7895
\$50,000,000	8	37.5000	37.7968
\$100,000,000	3	100.0000	100.0000

Where the resulting fitted severity distribution is the following:

Severity: Lognormal: M: 12.6026 ; SD: 2.0888

Patent Pending (Stamford Risk Analytics)

8.8.2 Scenario Analysis and Stress Testing

The above mentioned method of estimating frequency and severity parameters by fitting soft data to an annualized LEC can also be used to support scenario analysis and stress testing. For example, if a business line manager wants to enter a new business for which there is no loss data, he or she could use educated guesses about potential litigation threats, acts of nature, or supply chain failures to estimate the level of risk and conduct risk-based cost benefit analysis (as described in Section 7).

This method could also be used for stress testing operational risk as well as market risk, credit risk, business/strategic risk, insurance risk, etc. In particular, a method like this could be used to supplement hard data by adding soft data in the tail of the distribution, where information is sparse or nonexistent. Recall that one of the factors underlying the 2008 financial crisis was that many credit models assumed the probability of a housing

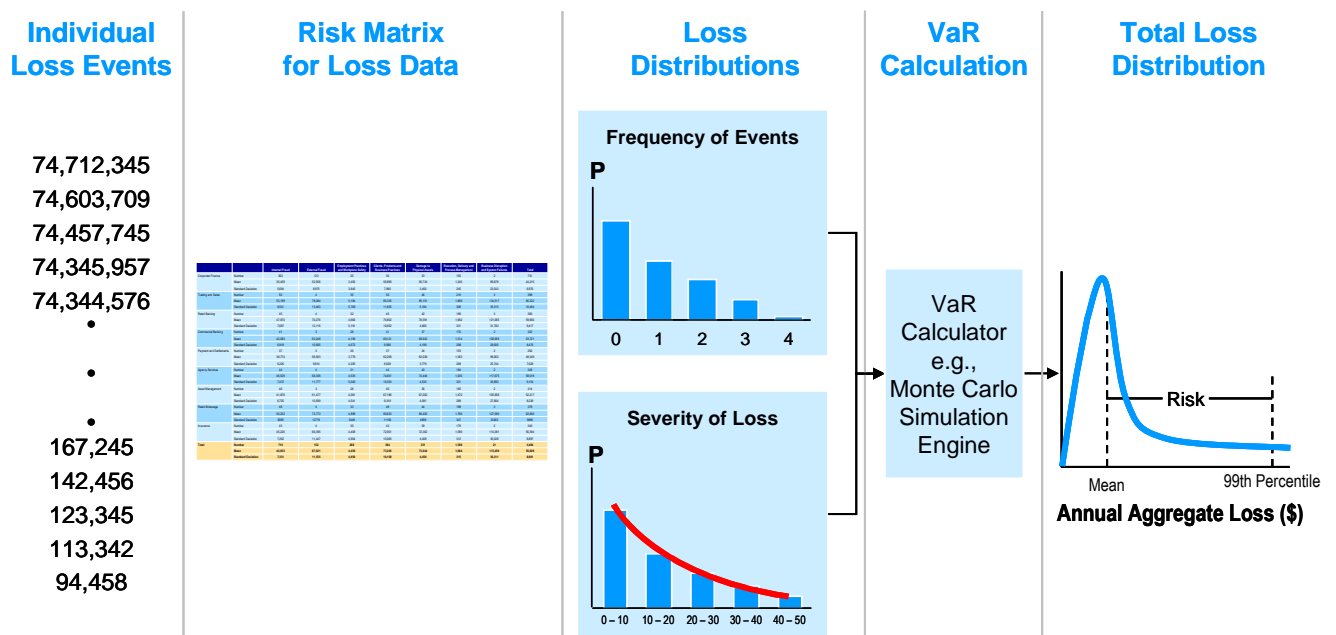
price decline across all U.S. geographies was 0%. This was supported by 72 years of hard data, during which period no such decline had taken place. Had this hard data been augmented by soft data derived from expert opinion, studies of prior financial market bubbles or other credible sources, some portion of the increased risk of collapse might have been recognized and priced into transactions. Again, this is not modeling in the traditional sense; the annualized LEC fitting method is an improvement over most other forms of scenario analysis.

8.9 Calculating Value at Risk

Value at Risk (VaR) is generally described as the Unexpected Loss at the organization’s risk tolerance level, though it is sometimes used to describe the Total Risk Exposure at that level. With the advent of high-speed computers, most organizations calculate VaR through Monte Carlo simulation.

In an actuarial framework, Monte Carlo simulation can be used to combine sets of individual frequency and severity distributions into aggregate loss distributions, as shown in Exhibit 8.8 below.

Exhibit 8.8 – Using Monte Carlo Simulation to Estimate the Aggregate Distribution



This is accomplished through a series of steps that are meant to simulate reality. Using the frequency distribution, the simulation engine randomly selects a value (N) that represents the number of losses the institution experiences in a given sample year. The simulation engine then randomly draws N losses from the severity distribution to determine the individual loss amount associated with each of the N events. The model then sums all the individual losses to calculate the cumulative or aggregate loss for that hypothetical year. By iterating this process thousands or millions of times, the simulation engine can create a hypothetical aggregate loss distribution that represents the

entity's theoretical loss exposure. The loss value that is 1% from the top represents the 99% level aggregate loss for that institution. Many simulation engines are able to incorporate insurance information and correlations (among business lines and risk categories) and can be designed to produce results with diversified and undiversified totals.

While correlation is a significant issue in market and credit risk modeling, it does not appear to be a significant factor in operational risk modeling. This is because most dependencies manifest themselves in terms of frequency, not severity. However, it has been demonstrated that when modeling with heavy tailed distributions (as are common in operational risk management), frequency correlations have little impact on the results.²⁶ This is because in operational risk modeling VaR results are driven not by frequency (a large number of small losses in different risk categories), but instead by severity (a single large catastrophic loss), such as the \$100 billion AIG CDS loss (which was one event) or the \$7 billion Société Générale Unauthorized Trading loss, both of which took place in 2008.

8.9.1 Issues with Value at Risk

In the wake of the 2008 financial crisis, VaR has been blamed for failing to accurately represent risk to investors, but much of the criticism directed at VaR may be unfounded or at least misplaced. Prior to the crisis, the public, regulators and even institutional investors were reassured by statements from major holders of collateralized mortgaged debt instruments that the risk of default was remote and that their level of risk as measured by their VaR models was well within institutional risk tolerances. Following the crisis, the warnings of long-time VaR critics, such as Nassim Taleb, author of *Fooled by Randomness* and *The Black Swan*, reverberated across the industry, and VaR became a convenient scapegoat.

So, what exactly is the problem with VaR? First of all, VaR is a threshold measure. That is to say, VaR represents the single value exceeded by a [typically very small] specified percentage of modeled outcomes. It provides no insight concerning the severity of outcomes beyond that point — just the likelihood of an outcome exceeding that point. Since VaR is rarely calculated for probability levels below 95% (and frequently calculated at much higher confidence levels like 99.5% or even 99.9%), the practical effect is to minimize concern regarding a loss of that magnitude occurring, let alone anything worse taking place. If a person is only told they can be 95% confident a given investment will not lose money (i.e., 95% VaR is a loss of \$0), they tend not to concern themselves with what happens the other 5% of the time. If, on the other hand, they are told that 95% of the time an investment will not lose money, but in the other 5% of cases they can expect to lose everything, this starts to sound more like gambling than investing, and they want to know more about the other 5%.

A more appealing, though still imperfect, measure of risk would be Tail Value at Risk, or TVaR. The difference between VaR and TVaR is that VaR provides the value beyond which a small percentage of outcomes lie,

²⁶ Frachot, Antoine, Thierry Roncalli and Eric Salomon, "The Correlation Problem in Operational Risk," Working Paper (January 2004).

whereas TVaR provides the expected (or average) outcome beyond that point. In essence, TVaR is the average of all outcomes beyond the VaR measure, or the conditional mean at the VaR% level. Under ordinary conditions, TVaR satisfies the requirements of a coherent risk measure, whereas VaR does not.²⁷

Even TVaR is subject to certain shortcomings, however. Consider the question of correlation, or dependency. It is generally acceptable to assume independence for many financial and economic variables within “normal” operating ranges of investment markets and macroeconomic conditions, and where independence is not appropriate in these circumstances it is usually possible to measure or estimate dependencies between key variables. For example, the probability of default for two randomly selected borrowers living in different states and working in different industries, but with the same average credit rating and loan characteristics, could normally be reasonably assumed to be independent and identically distributed.

Even with a large pool of i.i.d. risks such as we describe above, a problem arises when economic conditions worsen on a global basis. The debt rating of each pool of risks, which is a proxy for the perceived likelihood of default, will change in response to actual default experience for that pool as well as the experience of other such pools, particularly if the default experience is worse than initially expected. A perceived trend in downgraded credit ratings can reasonably be expected to result in tighter lending standards and increased interest rates for new borrowers, and possibly increased interest rates for existing borrowers. This tightening of credit slows existing home sales, new home construction and related purchases, and ultimately consumer spending, which adversely impacts economic indicators and leads to further credit tightening. Suddenly, all these i.i.d. risk pools are perceived as suffering from the same inherent problems and ultimately perception becomes reality. What we have just described is an example of conditional tail dependency — i.i.d. risks and risk classes behave independently up to a point, beyond which the degree of apparent dependency increases. Risk models that assume independence, or some measurable but constant degree of dependence across all possible economic states, will generally underestimate risk “in the tail.” As we have seen in the mortgage crisis, there is a much higher degree of dependence in extreme economic conditions than most of the prevailing VaR (or TVaR) models assumed.

To a large extent the issue surrounding VaR may have more to do with the assumptions underlying VaR models as applied in market and credit risk management. Specifically, in market risk management VaR is calculated based on daily price volatility and is then transformed into an annual VaR using “normal” statistical assumptions. Thus, annualized VaR represents a 1% probability of exceedence based on normal economic conditions, not the level of loss equivalent to a one in a hundred year event. If a 99% level VaR were actually meant to represent the latter, then it would be necessary to factor in the frequency of economic shocks every hundred years and their corresponding severity. In other words, market and credit risk would be modeled using an actuarial framework.

VaR or TVaR, as they are currently applied in market and credit risk, can be useful metrics for measuring short-term (daily/weekly) portfolio risk when one appreciates the limitations described above, but they are generally

²⁷ Artzner, Philippe et al. “Coherent Multiperiod Risk Measurement” (February, 2002)

inadequate for measuring long-term volatility. This is because one cannot extrapolate long-term risk measures from short-term risk measures using the methods in place today. Global economic and other macro forces do not fully manifest themselves in daily market movements, and “losers” are regularly factored out of market indices. For example, Enron, WorldCom, General Motors and Lehman Brothers are not reflected in current market return indices. The indices themselves are moving targets at times.

Under an actuarial framework, as shown in the tsunami example from Section 8.6.5, each economic environment would represent a different non-homogenous distribution, and the 99% level event across all economic cycles may be several times higher than the 99% level event in a “normal” economic environment. Recall that the 99% level ocean wave is a moderately large tsunami that is many times larger than the 99% level wind-driven wave.

A related issue with many risk measures concerns additivity. This refers to the inability to simply add results together for classes or sub-classes of risks to determine the aggregate level of risk for an institution, or an economy for that matter. The only time it is possible to add measures of volatility is when the individual risks within the underlying classes are 100% perfectly correlated with one another. What tends to happen in practice is that companies either assume independence between risk classes (which tends to understate risk in all scenarios), or assume some constant level of dependency that is intuitively appealing and easy to model (which may overstate risk under normal conditions and understate risk under stress scenarios). Neither approach is correct, and both can prove disastrous.

While VaR and TVaR are imperfect measures, and misapplication of these metrics can provide a false sense of security, a more appropriate target of criticism might be the assumptions underlying the risk models from which these metrics are derived — specifically the lack of attention to conditional tail dependency between individual risks and between key economic variables. The solution to conditional tail dependency lies in the use of statistical modeling techniques utilizing copulas, which are basically functions that define dependency structures that vary by percentile of the underlying distributions. It is possible to define copulas that increase levels of dependency between variables in both “tails” of the distribution, or perhaps only one tail, as appropriate. A description of copulas is beyond the scope of this paper, but the concept is critical to appreciate.

8.10 Modeling the Operational Risk Component of Other Risks

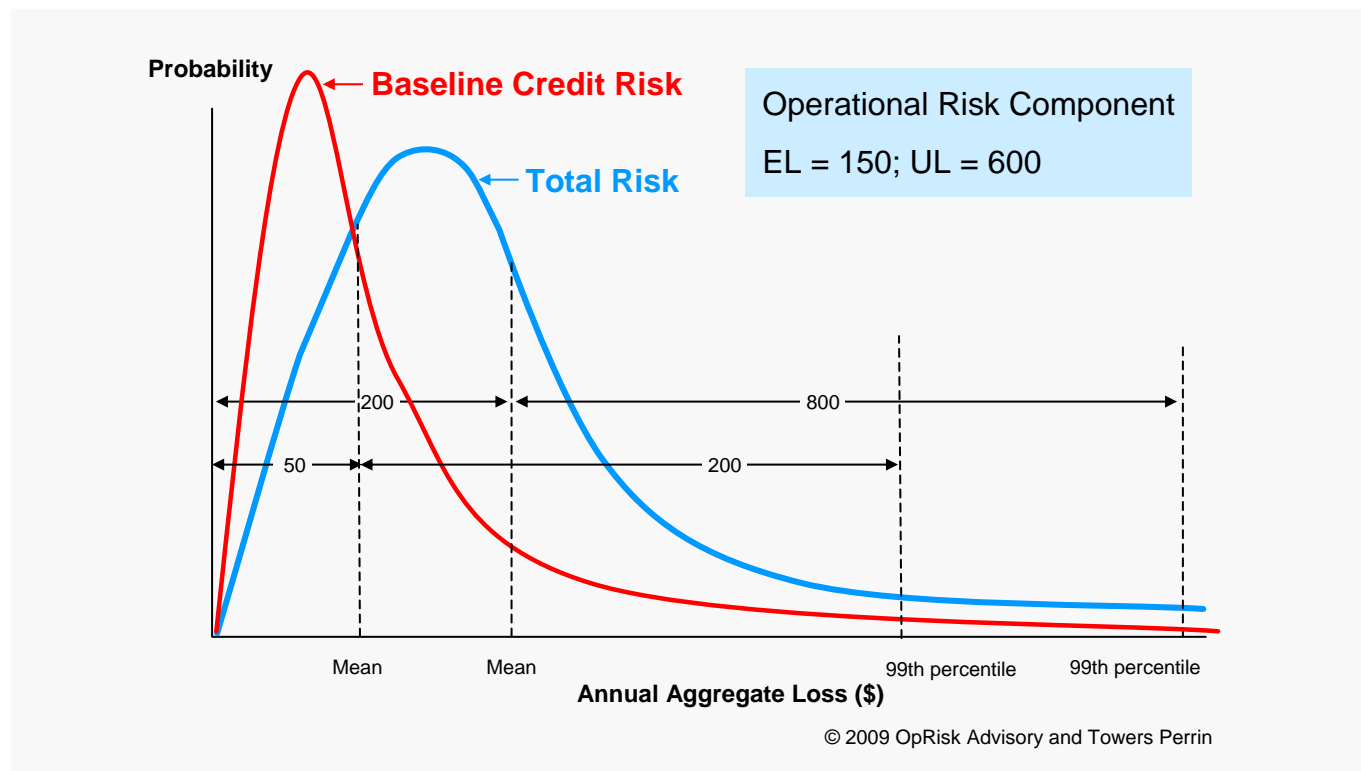
In the past operational risk has been defined to be a unique and distinct risk. However, as explained in Section 5, operational failures can manifest themselves in market, credit, business and insurance losses. It is important to recognize that a distribution of ordinary market and credit losses is very different from a distribution of losses that includes losses driven by operational failure. Treating such losses as ordinary market and credit losses obscures the underlying causes of the largest losses and leads to a misprioritization of risk management resources.

Because operational risk is embedded in the other risks, it is important to isolate and quantify the operational risk component of other risk, in order to focus attention on the underlying causes of the major losses. In the short term this could be accomplished by using methods described in Section 8.9.1, where hard data from the other risk areas

and soft operational loss data are used in conjunction with one another. In the long term, a more effective method would be to specify data requirements such that all major losses (in excess of \$1,000,000) were disaggregated by type of risk, based on contribution to total loss. To ascertain the operational risk component, the underlying question may be: What would have been the magnitude of this loss if no operational failure had been present? The remaining amount would be deemed to be the operational loss component.

Exhibit 8.9 illustrates conceptually the way one could measure the contribution of operational risk to other types of risk. This method does not require one to oversimplify the data and modeling issues. It is likely that a reasonable result could be obtained by modeling only those events where the loss value exceeds \$1,000,000.

Exhibit 8.9 — Measuring Incremental Operational Risk



9. The Business Case for Modern ORM

9.1 The Current State of ORM

In most organizations the role of the ORM department is not well defined. In the absence of tangible business objectives, common practices in ORM represent a mix of silo-based tasks designed to comply with regulatory, rating agency and/or audit requirements. These tasks generally include using Traditional ORM methods, such as RCSA, the tracking of open audit issues, process mapping, etc. Following the introduction of the Basel II banking regulations, these tasks have come to include collecting and analyzing internal and external loss data and “key” risk indicators and calculating/allocating operational risk capital.

However, many senior executives believe that a large number of the tasks performed by ORM fall naturally within the ambit of audit and compliance. They do not see what additional value is created by introducing another control function that performs activities that appear to them to be very similar to those normally performed by two other similar functions. Consequently, the ORM function constantly has to justify its existence.

In addition, many organizations have not been impressed with the results produced by ORM to date. To a large extent this is because many of the processes and procedures that reside under the banner of ORM are very similar to those performed by audit and compliance. These activities are hugely resource-intensive, place an excessive burden on the business lines and do not appear to produce much in the way of tangible benefits. But to some extent this also has to do with the fact that justifying the value of ORM is a Catch 22 situation: An effective ORM program will generally bring about a reduction in losses, but when no major losses are taking place there is less appreciation for ORM. Still, there is a legitimate case to be made that common practices in ORM are not best practices. In fact, if the true goal of ORM is to more effectively manage the key operational risks, there is much room for improvement in ORM. There may actually be a need for a complete paradigm shift in ORM practices. According to an October 2009 Towers Perrin survey of senior financial executives, only 15% of respondents indicated they were less concerned regarding the effectiveness of their company’s risk management function than they were at the height of the financial crisis, and only 4% felt that their company’s operational risk practices needed no improvement.

9.2 Key Differences between Traditional and Modern ORM

The key differences between Traditional and Modern ORM are summarized in Exhibit 9.1 below

Exhibit 9.1 — Summary of Differences between Traditional and Modern ORM

Traditional ORM	Modern ORM
<ul style="list-style-type: none"> ■ Definition: Risk is defined primarily as a kind of undesirable incident/event, such as a fraud or a system failure (Operative question: What/where are your risks?) 	<ul style="list-style-type: none"> ■ Definition: Risk is defined primarily as a measure of exposure to loss from undesirable incidents/events (Operative question: How much risk do you have?)
<ul style="list-style-type: none"> ■ Risk Identification Process: Ask managers to identify their major risks. (Risks include risk factors, controllable factors, events and effects; no restriction on overlaps; generally no differentiation made between risks and controls.) Leads to the creation of a huge and unmanageable set of risks 	<ul style="list-style-type: none"> ■ Risk Identification Process: First define the “risk” universe, consisting of a finite (comprehensive) set of mutually exclusive (non-overlapping) “risk” classes. Use hard or soft data to reveal where the large losses are taking place (where the largest risks actually exist)
<ul style="list-style-type: none"> ■ Risk Assessment/Measurement Method: Calculate risk by multiplying likelihood and impact for each risk type (conditional on one event), one “risk” at a time 	<ul style="list-style-type: none"> ■ Risk Assessment/Measurement Method: Use Monte Carlo simulation and frequency and severity distributions to calculate the cumulative loss potential from multiple events, across all risk classes simultaneously
<ul style="list-style-type: none"> ■ Aggregation: Likelihood cannot be aggregated, so results cannot be aggregated 	<ul style="list-style-type: none"> ■ Aggregation: Frequency can be aggregated, so results can be aggregated
<ul style="list-style-type: none"> ■ What is measured: Probability weighted loss from one specific incident (the routine loss) 	<ul style="list-style-type: none"> ■ What is measured: Cumulative loss for one or more risk classes; both the expected loss and unexpected loss, which are comparable to the average and “worst case”
<ul style="list-style-type: none"> ■ Goal: Day-to-day management of current threats arising from imminent operational failures: loss prevention through tactical intervention 	<ul style="list-style-type: none"> ■ Goal: Management of <i>key</i> risks, specifically the optimization of risk-reward, risk-control and risk-transfer in the context of cost-benefit analysis
<ul style="list-style-type: none"> ■ Cost: Generally very resource intensive 	<ul style="list-style-type: none"> ■ Cost: Relatively much less resource intensive

Traditional ORM was developed to prevent and avoid the normal losses — those stemming from routine process failures and arising out of basic control weaknesses. Modern ORM also focuses on losses, but Modern ORM gives much more weight to the types of events that cause the large losses — specifically the drivers of tail events that contribute to the expected loss and unexpected loss, and therefore have the most significant impact on financial performance and even threaten the survival of the firm.

Traditional ORM uses linear, audit-based techniques to prevent losses generally resulting from routine events. Modern ORM tries to address the more esoteric problems through the use of sophisticated classification methods, advanced control assessment processes and mathematical models that improve the effectiveness of risk-based decision making. Under Modern ORM the goal is not just to avoid losses; it is to facilitate the optimization of risk-reward, risk-control and risk-transfer decisions.

Under Modern ORM, the goal is not just to serve the needs of business line managers and senior officers. An important part of ORM is to mitigate principal-agent and business practices risk by increasing transparency in the business decision-making process and eliminating information asymmetries.

Both Traditional ORM and Modern ORM can add value, though as explained in this report, there is much room for improvement in Traditional ORM practices. Both Traditional ORM and Modern ORM are necessary, but neither is sufficient on its own. An effective ORM program should combine elements of the Traditional and Modern approaches in a manner scaled to their respective potential loss exposure for the organization in question. Therefore, many of the activities associated with Traditional ORM may need to be scaled back, with additional resources focused on Modern ORM as described above.

9.3 The ORM Evolutionary Path

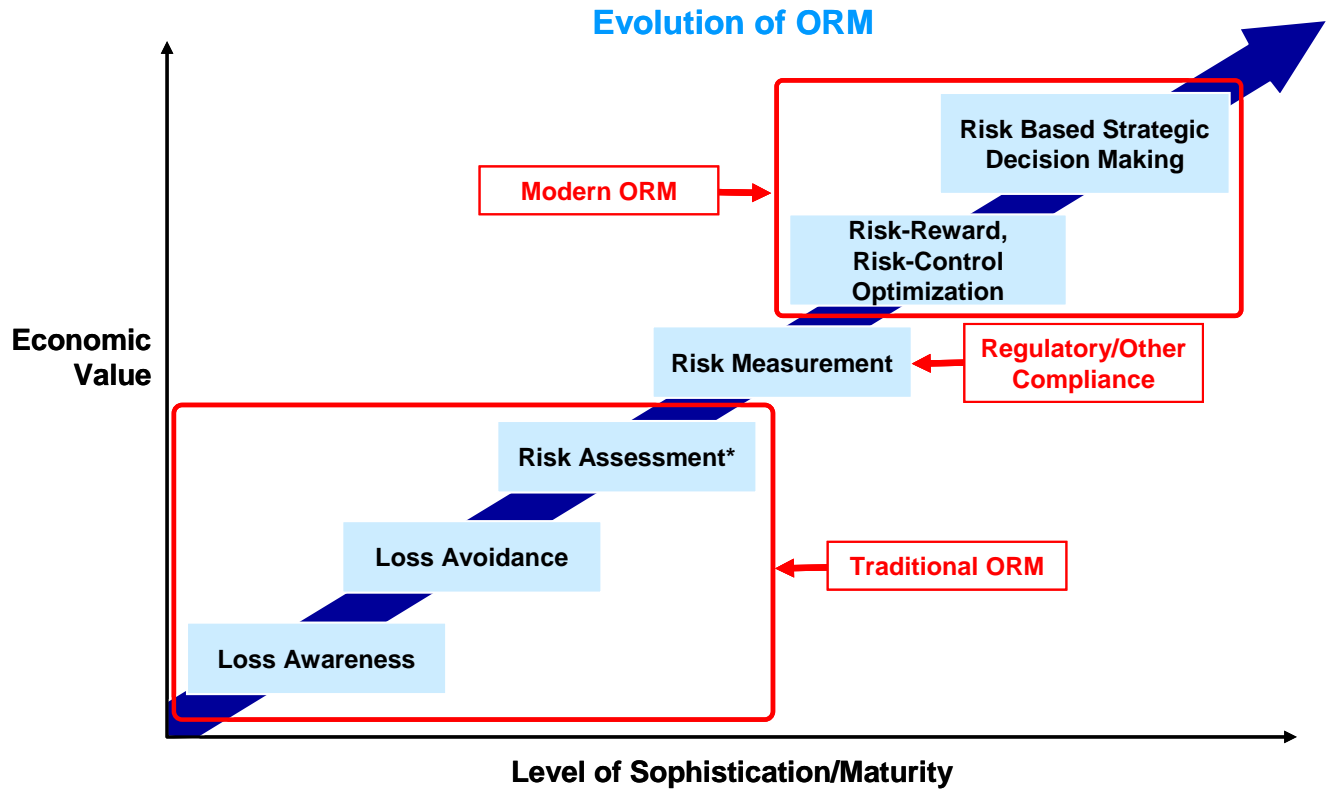
The evolutionary path from Traditional to Modern ORM is summarized in Exhibit 9.2 below. As shown in this figure, the set of tasks performed under Traditional ORM are best characterized as the early stages tasks in the ORM evolutionary process. The tasks are designed to avoid and prevent losses. Loss prevention and loss avoidance are part of the overall risk management process, but they represent basic forms of risk management. Because these activities do not actually involve risk measurement, concepts such as risk tolerance, risk-reward, risk-control and risk-transfer optimization do not really exist²⁸.

Modern ORM represents that next step in the evolution of ORM. When firms evolve to the Modern ORM stage, they are able to engage in the higher value-adding activities. At that point they can legitimately claim to be able to make tactical and strategic risk-based decisions, in the context of cost-benefit analysis, and in conformity with the risk tolerance standards of the stakeholders.

Organizations that measure/model operational risk only for purposes of calculating risk capital do not embrace the concepts of Modern ORM. In addition, modeling operational risk only for the purposes of satisfying regulatory or other requirements represents perfunctory compliance. Proving that the results are used in decision analysis is necessary to meet the requirement of the Solvency II and the Basel II “use test.” Where risk information is not used in risk-based decision analysis, the modeling process adds little value.

²⁸ Some practitioners have tried to apply these concepts to Traditional ORM. However, these terms bear very little resemblance to the terms used in Modern ORM.

Exhibit 9.2 — The Evolutionary Path from Traditional to Modern ORM



* Traditional ORM practitioners believe they are conducting "risk assessment," while they are actually performing "expected loss assessment."

Source: Adapted from SunLife Financial 2007.

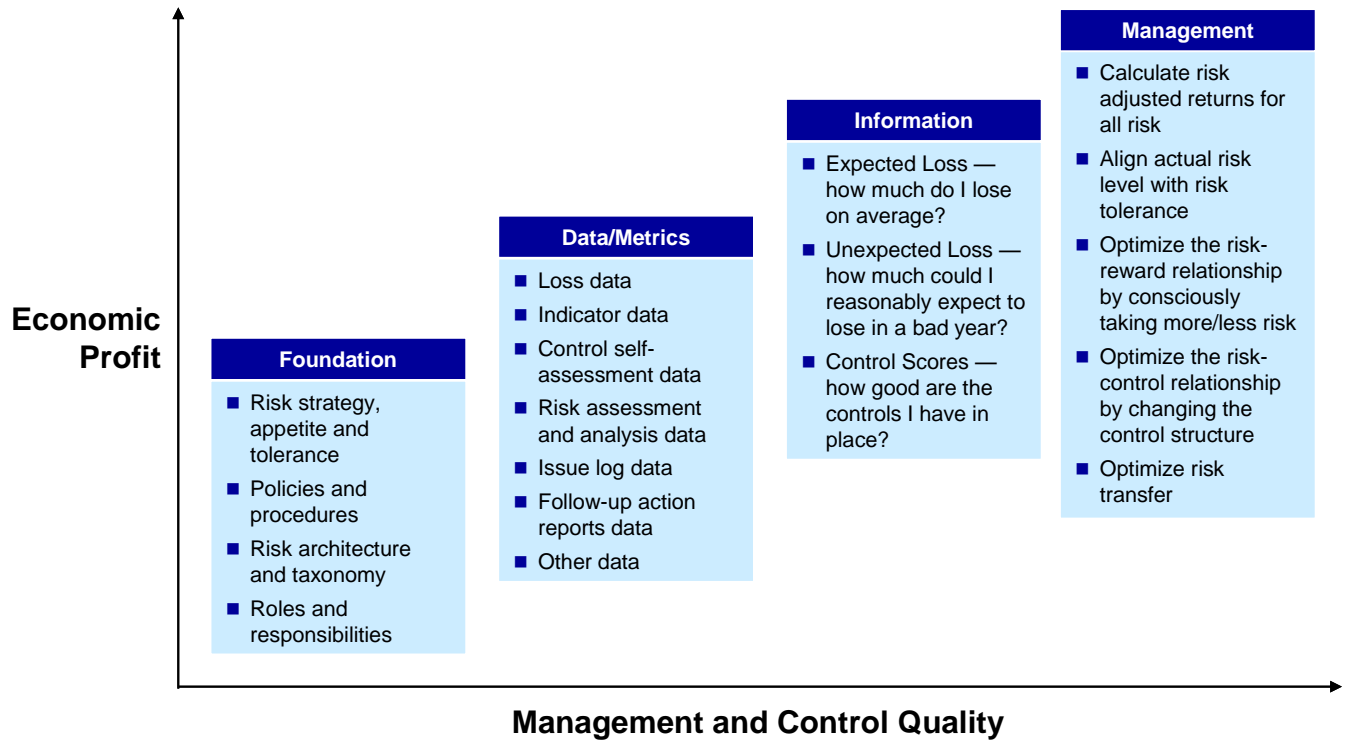
9.4 The ORM Roadmap

As explained above, the goal of ORM is to manage operational risk in a value-adding rather than just a compliance context. This involves establishing a process to transform raw operational risk data into risk information that supports educated risk management decision-making, as well as creating greater transparency in the decision-making process. This must be accomplished through an integrated framework, based on a consistent definition of risk and a common set of objectives.

As shown in Exhibit 9.3, implementing a Modern ORM framework requires a sound infrastructure. This involves a multi-step, sequential process, where each task builds on a set of previous tasks. For example, in order to conduct strategic risk-based decision analysis one must first develop a comprehensive, mutually exclusive risk taxonomy (Step 1), gather relevant internal and external loss data (Step 2), transform this data into expected loss and unexpected loss estimates (Step 3) and then optimize the risk-reward, risk-control or risk-transfer in the context of cost-benefit analysis in conformity with the risk tolerance standards of the stakeholders (Step 4).

Exhibit 9.3 describes the roadmap from foundational activities to educated decision making.

Exhibit 9.3 — The ORM Roadmap



9.5 The Economics of Modern ORM

During the past few years major North American corporations have grudgingly spent billions of dollars complying with a plethora of requirements brought on by Sarbanes-Oxley, COSO ERM, the SEC, the rating agencies as well as specific industry regulations. Understandably, most companies have had their fill of risk management and compliance, and in the current economic environment there is virtually no appetite for investment in any new risk management initiatives. Nevertheless, investing in Modern ORM can be a beneficial course of action.

The case for Modern ORM is based on the following points:

1. Transitioning from Traditional to Modern ORM should not result in an increase in resources. Quite to the contrary, adopting Modern ORM can allow organizations to scale back many highly resource-intensive Traditional ORM initiatives, such as RCSA.
2. Some large organizations have estimated that implementing a robust RCSA program requires a 200-person effort annually. The Modern ORM equivalent, Risk Assessment and Control Assessment, can be very effectively implemented at a significantly reduced cost.

-
3. If implemented correctly, adopting a Modern ORM approach will create value, because Modern ORM is focused on identifying, understanding and mitigating exposure to the largest risks — those risks that have the greatest propensity to impact financial performance and solvency.

One reason that it has traditionally been so difficult to assess the value of any investment in risk management is that this value is typically realized only by not investing, and then only in retrospect when an event occurs that might have been prevented had the investment been made. Suppose it were possible to press the “rewind” button and replay the latter half of this decade. If the major companies had invested the necessary resources in Modern ORM and avoided losing trillions of dollars, tongues would doubtless be clucking at the waste of time and money invested in “unnecessary” risk mitigation efforts. It is only through a thorough and robust analysis of risk as outlined herein that a value can be placed on risk management investments.

The key question to ask is how much would U.S. companies (and taxpayers) have had to spend on ORM for it not to make economic sense? The return on investment from preventing or containing the global economic meltdown is incalculable. Nonetheless, even today, businesses continue to spend billions of dollars complying with regulations and reporting requirements that often fail to address the most serious risks.

If the goal of ORM is to reduce the level of risk, then one can argue that Traditional ORM approaches — including COSO, Basel II ORM and others (as currently implemented) have failed. By following a Traditional approach we continually “shrink” the volatility of the least significant risks that produce the commonly observed events. Unfortunately, nothing is done to address the truly critical risks.

One could equate the continual evolution of likelihood x impact analysis to the proverbial myopic solution: Focus on rearranging the deck chairs on the Titanic so that passengers are less likely to trip. Tripping over deck chairs is a problem, but certainly not the most critical risk passenger face. Yet, there are more deck chairs than icebergs. And the deck chairs threat is more obvious; this is regarded as addressing a more “practical” matter.

10. Conclusions and Recommendations

Historically, ORM has taken a back seat to the management of the other major risks, which are often defined as market, credit, insurance and strategic risk and sometimes include “liquidity,” “legal” and “reputation” risk. This is largely because operational risk is often confused with operations risk. This has not only caused operational risk to be underestimated, but has also obscured the underlying causes of many of the most significant financial losses.

Financial firms have spent millions of dollars on ORM, but with limited success, since they have generally taken a traditional audit-based approach (Traditional ORM). In the United States, the broad principles underlying this general approach have been incorporated into a set of enterprise risk management standards that are referred to as COSO ERM. Most major accounting firms and numerous consulting firms, rating agencies, industry bodies and independent experts advocate using this approach or a customized version thereof. A majority of national and international bank regulators have also at least tacitly endorsed this approach. Finally, a large number of corporate CFOs believe that Traditional ORM represents the standard for best practices. Consequently, virtually every organization that has implemented an ORM program has based the underlying framework on the principles of Traditional ORM.

Traditional ORM provides structure, governance standards and an intuitive approach to risk identification and assessment, but mathematically equates risk with the average, or expected loss. In reality, however, risk is more appropriately represented by the unexpected, or even “worst case” loss. This discrepancy has huge implications. Specifically, Traditional ORM relies greatly on the RCSA process, which fails to reveal the real risk and focuses instead on common threats and control weaknesses associated with routine losses — independent of risk. Therefore, the largest risks go unrecognized and institutions that follow this approach remain blissfully unaware of their most significant risks. In fact, organizations following this approach often become over-controlled in the areas where they have the least risk and remain significantly under-controlled in the areas where they have the most risk.

Traditional ORM is also designed to be implemented at a granular/process level, which makes it a hugely resource-intensive process. And, because the assessment metrics are not additive, the resulting “risk” figures cannot be aggregated. For all these reasons, most managers find it very difficult to use Traditional ORM information in a systematic way.

In contrast, Modern ORM is a top-down approach, which focuses first on the major risks — within a comprehensive and mutually exclusive risk architecture — and which drills down only in those risk areas where more granularity is required. This allows practitioners to triage the risk management process. It is significantly less resource-intensive and avoids focusing management attention and resources on immaterial risks. Modern ORM can be implemented for tactical and strategic decision analysis.

The Solvency II regulations, which are scheduled to become effective in Europe in 2012, use language that is consistent with Modern ORM. These regulations describe risk as a measure of uncertainty (specified at a 99.5% confidence level for a one-year time horizon). In addition, the Solvency II “use test” requires that internal models must play an integral role in any organization’s system of governance and risk management as well as its economic and solvency assessment/capital allocation. Solvency II will directly impact North American companies with international operations and will eventually translate into a corresponding set of U.S. and Canadian regulations. The rating agencies are also likely to evaluate insurance companies based on these or similar standards; some have already expressed an intention to do so.

Adopting a Modern Approach will also allow companies to focus on the most important business problem: mitigating exposure to the large events — the events that have the greatest impact on financial performance and solvency. Modern ORM aims to:

- Facilitate the holistic management of all operational risks, based on a consistent definition of risk and a comprehensive risk architecture/taxonomy.
- Create a structured and transparent process for factoring risk into the business decision-making process — at both a tactical and strategic level. Specifically, provide managers, senior managers and C level executives the tools and information they need to optimize risk-reward, risk-control and risk-transfer in the context of cost-benefit analysis.
- Embed a risk culture that reflects and harmonizes the goals of key decision makers and external stakeholders.
- Reduce information asymmetries between managers and stakeholders to help confirm that managers are pursuing strategies that conform with the risk tolerance standards of the stakeholders — in other words, mitigate principal-agent risk.

For the reasons given above, the authors recommend that North American insurance companies consider developing formal ORM programs. These programs would benefit from the principles of Modern ORM. Traditional ORM has many useful aspects; some of these program features should be retained as well. However, companies that have already developed ORM programs may find it beneficial to consider scaling back several of the highly resource-intensive aspects of Traditional ORM and replacing them with their Modern ORM equivalents — which will significantly reduce cost. Transitioning from Traditional to Modern ORM, if implemented effectively, may not only improve risk management practices but may allow organizations to do so with less expense.

In order to move beyond Traditional ORM — along the path towards Modern ORM — companies will need to:

- Reorient their conceptualization of risk from a qualitative notion of unpleasant outcomes to a quantitative representation of uncertainty;

-
- Replace likelihood x impact analysis with frequency and severity analysis;
 - Adopt a more robust risk architecture/taxonomy;
 - Recognize that low-frequency, high-severity events actually pose the greatest risk; and
 - Continue to deploy Traditional methods where they are effective, but recognize that these approaches are inadequate when dealing with the most significant risks.

Regulators will find it beneficial to encourage companies to evolve from Traditional ORM to Modern ORM. Operational risk is a key risk, but “common practices” in ORM represent a series of disparate, silo-based compliance tasks that have no real business objective and produce little if any value. A move to Modern ORM would transform ORM into a purposeful exercise that legitimately adds value.

At present, many extreme losses assigned to credit, market and liquidity risk are too readily classified as outliers because of their unique nature, whereas the operational risk components of these events are very instructive. Much better information can be gleaned from evaluating the incremental impact of operational risk on these other risk categories. Therefore, instead of tossing out these losses from the experience base, the “normal” portion should be recognized with the appropriate risk classification and the incremental impact should be recognized as operational risk. Objective methods have been developed to address this problem (see Section 8.10.)

In the wake of the recent financial crisis, the general public has an increased awareness of the complexity of the global financial environment. There is now a greater perceived need for a regulatory regime that focuses on the risks that matter rather than the minutiae that consume so much collective energy and resource. While some are pointing to the need for additional regulation, most are calling out for more sensible, risk-based regulation. From the vantage point of investors and the general public, compliance with inadequate or ineffective regulation can potentially be more harmful than non-compliance with effective regulation.

Regulators and other key stakeholders (e.g., rating agencies) need to take an active role in calling for improved ORM practices. ORM should be viewed within the broader ERM context. Beyond the practical and definitional differences between Traditional and Modern ORM outlined herein, the authors believe one of the more important aspects of effective ORM relates to the concept of the payoff matrix, and in particular the problem of principal-agent risk. This particular aspect of operational risk is extremely important to understand, because the reality of today’s global economy is that virtually every business depends on the performance of agents in some capacity, and where direct day-to-day control over agents is not possible or feasible, effective risk management protocols and processes including legitimate, transparent metrics must be put in place.

As previously mentioned, developing a set of best practices for implementing ORM was not part of the scope of this project. However, because of the critical importance of operational risk the authors strongly recommend that this become the focus of a future research initiative.

The continued focus on market, credit, liquidity and legal risk — in the face of increasingly complex corporate operating models and global economic interdependencies — reminds one of the old joke about the man who was searching the ground meticulously near a street lamp one night. A passerby asked what he was looking for and the man said that he had dropped his car keys, so the passerby also began searching. After several minutes it became clear the keys were nowhere to be seen and the passerby asked the man if he remembered where he last had the keys in his possession. The man pointed down the darkened street and said that he was pretty sure he had dropped his keys near his car, which was parked about a block away. Puzzled and a little frustrated, the passerby asked the man why he was searching in this particular spot, if he thought he had dropped his keys so far from there. The man looked up, pointed down the street and said, “Because there is no light over there.” Our hope is that this research effort provides some light so the search can be moved to more productive areas.

11. References

Artzner, Philippe, Freddy Delbaen, Jean-Marc Eber, David Heath and Hyejin Ku, “Coherent Multiperiod Risk Measurement” (February, 2002).

Basel Committee on Banking Supervision, “A New Capital Adequacy Framework” (June 1999).

Basel Committee on Banking Supervision (Secretariat of the Basel Committee on Banking Supervision), “The New Basel Capital Accord: an explanatory note” (January 2001).

Basel Committee on Banking Supervision, “Working Paper on the Regulatory Treatment of Operational Risk” (September 2001).

Basel Committee on Banking Supervision, “International Convergence of Capital Measurements and Capital Standards” (June 2004).

Cruz, Marcelo G., “Modeling, Measuring, and Hedging Operational Risk,” New York: John Wiley & Sons, Inc. (2002).

Dutta, Kabir and J. Perry, “A Tale of Tails: An Empirical Analysis of Loss Distribution Models for Estimating Operational Risk Capital,” Federal Reserve Bank of Boston, Working Paper (April 2007).

Frachot, Antoine, Thierry Roncalli and Eric Salomon, “The Correlation Problem in Operational Risk,” Working Paper (January 2004).

Holton, Glyn A., “Defining Risk,” *Financial Analysts Journal*, Vol. 60(6), CFA Institute (December 2004).

Klugman, Stuart A., H.H. Panjer, and G.E. Willmot, “Loss Models: From Data to Decisions,” 2nd Edition, New York: John Wiley & Sons, Inc. (2004).

Knight, Frank, “Risk, Uncertainty, and Profit,” 1st edition, Boston: Houghton Mifflin Company (1921).

New York State Insurance Department News Release: “Allstate Should Report Any Illegal or Inappropriate Use of Credit Default Swaps;” (April, 2009).

Ross, Sheldon M., “Introduction to Probability Models,” 7th Edition, New York: Academic Press (2000).

Samad-Khan, Ali, “Categorization — A Solution,” Working Paper (2002).

Samad-Khan, Ali, “Modern Operational Risk Management,” *Emphasis* (2008/2).

Samad-Khan, Ali, A. Rheinbay and S.L. Blevic, “Fundamental Issues in OpRisk Management,” *OpRisk & Compliance* (February 2006).

Samad-Khan, Ali, “Why COSO is Flawed,” *Operational Risk* (January 2005).

Shih, Jimmy, A. Samad-Khan and P. Medapa, “Is the Size of an Operational Loss Related to Firm Size,” *Operational Risk* (January 2000).

Skinner, Tara, “In Defense of AMA Methodology,” *OpRisk & Compliance* (February 2006).

Walhin, Jean-Francois and J. Paris, “On the Use of Equispaced Discrete Distributions,” *ASTIN Bulletin*, Vol. 28(2), 241 – 255 (1998).

Wikipedia: The Free Encyclopedia. Wikimedia Foundation (July 22, 2004).

Wang, Shaun S., “Aggregation of Correlated Risk Portfolios: Models & Algorithms,” *Proceedings of the Casualty Actuarial Society*, Vol. LXXXV, 848 – 939 (1998).

12. Appendix A — Principal-Agent Risk and the 2008 Global Financial Crisis

The 2008 global financial crisis wiped out countless billions of dollars of wealth and brought the global financial system precariously close to a total meltdown. It also revealed that a major corporation's most significant risk may be principal-agent risk. Principal-agent risk refers to the risk that, in circumstances where there is separation of ownership and control, an agent (who controls or acts on behalf of an organization) may pursue actions that are in his/her own interest, but are not necessarily in the best interest of the principals/stakeholders (the stockholders, bondholders, etc.). Where the institution is deemed "too big to fail," the ultimate stakeholders are the taxpayers.

Principal-agent risk was the underlying cause of two of the most widely recognized drivers of the recent financial collapse: the so-called "Sub-Prime Credit Crisis" and AIG's credit default swaps debacle.

Ordinarily, laws and regulations monitored through legitimate, transparent metrics prevent agents from exploiting principals. However, where information asymmetries and flawed performance metrics exist, this may not hold true. Where it is difficult to differentiate between legitimate profits and phantom profits, those generating huge spurious profits will be generously rewarded, while those producing legitimate but moderate returns will be criticized, perhaps even penalized. In such situations, agents may engage in business activities that create the appearance of profitability when they are actually destroying value. Under these circumstances, even well intentioned executives, particularly those faced with a "heads I win; tails someone else loses" compensation structure, may take the path of least resistance and disregard their fiduciary responsibilities.

This raises some very troubling issues. It suggests that where information asymmetries and inaccurate performance metrics exist, irresponsible behavior at just one firm can very quickly replicate itself, eventually causing an industry trend. Thus, operational failure can be a key driver of systemic risk.

The Sub-Prime Credit Crisis

The housing boom took place partly because in the early 2000s the U.S. experienced historically low interest rates, which reduced the overall cost of homeownership. In addition, at the same time the federal government decided to broaden the mandate of Fannie Mae and Freddie Mac to encourage home ownership among sub-prime borrowers. As the mortgage market expanded to include these somewhat riskier loans, financial markets responded by promoting special instruments designed to mitigate idiosyncratic risk through diversification. Among these instruments were mortgage-backed securities, which were sold to investors in tranches. The tranches had varying degrees of risk based on type of borrower and loan characteristics. Because the pooling of large numbers of borrowers generally reduces non-systemic risk, many of these securities were given investment-grade debt ratings. Companies invested in these securities and then protected their investments with CDSs and other financial guarantees.

Potential homeowners and investors took advantage of these loans to purchase new and existing homes. This had the general effect of increasing housing prices. As home values increased, new homebuyers and investors entered the market, which further increased home prices. The resulting rapid increase in home values further convinced potential new homeowners, existing homeowners and investors that real estate was an outstanding investment, which resulted in a further increase in demand for housing and mortgage products.

Because lenders were able to pass on these mortgage products to investors, they were able to free up capital, thereby allowing them to grant additional loans. Thus, profitability for these institutions became a function of turnover efficiency. Many lenders were unable to process this large volume of loans. Therefore, to meet the surge in demand, lenders delegated a portion of their own underwriting authority to brokers. In addition, underwriting requirements in general were loosened across the industry. Brokers were authorized to offer “no-documentation” (and “low-documentation”) loans, where the borrower did not need to substantiate his/her income and other financial information (though they were charged nominally higher rates to cover the added risk of default). This naturally led to a decline in the overall credit quality of the loans.

Bubbles take place when a condition that is observed to be true in “normal” circumstances is assumed to be true in all circumstances (e.g., real estate prices never decline). This can drive micro-economic behavior (an overwhelming demand for real estate). When coupled with irresponsible lending to home buyers (no income verification loans, no down-payments), it can have macro-economic implications (prices rise because of a disequilibrium between supply and demand). When this happens, the rapid rise in prices serves to confirm the original premise (housing prices only go up) and fuels much more of the same behavior. When these bubbles eventually become unsustainable they bring about the very circumstances in which the original assumption is no longer valid (housing prices decline) — rendering the original premise a self-invalidating prophecy.

It should have been evident to any objective analyst that the housing bubble was unsustainable. But every agent in this process — from the loan originator, appraiser, rating agency, bond insurer, to the trader and CDS provider — had an incentive to ignore the risk. Even though this business carried excessive risk, because individual agents assumed no part of the risk they had no incentive to stop doing what they were doing. The payoff structure they faced was: “Heads I win, tails somebody else loses.” In such a situation it was easy for the agents to continue to believe that because these securities were backed by real estate — which could never decline in value — there was nothing to worry about.

In today’s competitive business environment, performance is benchmarked against peers. In an efficient market this has desirable consequences. However, in this situation it had the opposite effect. Because of the need to meet quarterly earnings expectations, which were benchmarked against industry averages, the CEOs of large financial institutions were “compelled” to remain in the game. Had these senior executives acted responsibly by reducing their volume of mortgage loans, they would probably have been dismissed for performing below par. To justify this strategy they turned to their credit risk models. These models were based on historical data which indicated that housing prices had never simultaneously declined across all geographies in the 72 years since such data had

been collected. Consequently, virtually all models conveniently showed that the probability of such a default was zero. Therefore, instead of reducing their exposure to the mortgage business, most lenders *took refuge* in these models and continued to service the seemingly insatiable demand.

AIG and Credit Default Swaps

The AIG Financial Products unit was created in 1987. It offered its clients so-called “plain vanilla” hedging instruments, such as interest rate swaps. The unit performed well. For example, it contributed about \$737 million (4.2%) to AIG’s net profits in 1999. Then the opportunity arose for AIG to begin writing credit default swaps (CDSs). CDSs are products that insure a bondholder against default by a bond issuer in exchange for a premium payment. Selling these instruments to holders of highly rated blue-chip company debt proved enormously profitable because of the extremely low default rate of the corporate bonds and mortgage-backed securities AIG was insuring. In addition, because these instruments were not classified as insurance products, AIG Financial Products was not required to hold any additional regulatory capital to protect against the risk of insolvency. As a result, the AIG Financial Products business was able to generate abnormally high rates of return on capital.

From 1987 to 2004, the AIG Financial Products unit contributed more than \$5 billion to AIG’s pre-tax corporate earnings. The managers and employees of this unit also did very well; the average annual compensation per person exceeded \$1 million from 2001 through 2005.

By aggressively pursuing the CDS business the AIG Financial Products unit was also able to generate abnormally high returns on capital, but only by taking risk that was in excess of tolerance standards of the stakeholders. Sadly, the assumptions underlying their risk models turned out to be too optimistic. As the underlying securities began to experience defaults, AIG was forced to post collateral, ultimately leading to a liquidity crisis at the mammoth institution and a bailout by the U.S. government.

Had AIG been required to reserve capital to support its credit default swap writings, returns on equity for this business would have been lower and the deal-makers within Financial Products would have had to pay much closer attention to the risk-reward trade-off. Alternatively, had the true risk of the AIG credit default swap portfolio been recognized and managers’ compensation had been based on risk-adjusted profits, the Financial Products unit would likely have been less aggressive in writing this business.

Summary and Conclusions

The principal-agent problem has always existed, but it takes extraordinary circumstances for this problem to cause a financial meltdown. The 2008 financial crisis was devastating because the sub-prime lending business and the credit default swap business complemented each other and were allowed to grow unchecked for years, even expanding across international borders.

So one may naturally ask, if the underlying cause of principal-agent risk is human nature — i.e., people will do what they need to do to gain/survive, even if this comes at the expense of others — how can future crises be

prevented? There has been a great deal of discussion on this topic. Much of the discussion has focused on compensation and incentives. Many have suggested reducing compensation outright or, alternatively, changing the compensation mix for all executives to include a large portion of deferred stock compensation. This solution is infeasible. Currently, most senior officers face a “heads I win, tails someone else loses” payoff structure. Introducing deferred stock compensation will not change their incentive structure. It will only change their payoff structure to “heads I win big, tails I win small.” This is comparable to allowing an executive to keep his \$5 mm salary but having him give up a \$10 mm bonus. This is clearly not enough of a disincentive.

In order to mitigate principal-agent risk, the payoff structure must be changed to “heads I win, tails I lose.” In other words, people who knowingly engage in value destroying activities must lose in proportion to the level of damage they cause. Specifically, there must be negative consequences for persons who knowingly violate their fiduciary responsibilities.

While compensation and incentives are important they are only a part of solution. The more important issue has to do with the way risk is measured and managed. The fundamental problem is that currently many organizations use risk models that systematically underestimate the level of risk because they fail to accurately incorporate the impact of the rare events (see section 8.6.5). Where risk-adjusted performance measures are based on the results of flawed models “risk-reward arbitrage” opportunities will exist. Where decision-makers recognize these opportunities, and where information asymmetries also exist, some will act irresponsibly and take excessive risks in order to maximize their personal rewards. This will then create its own set of dynamics, because in order to keep up with their peers, other firms will follow suit — causing an industry trend.

Preventing a repeat of the 2008 global financial crisis will require a paradigm shift in risk management practices. Risk models can provide valuable insight into complex problems, but the quality of these models and critical assumptions need to be validated by objective and independent experts. In particular, these models must be able to incorporate both empirical data and expert opinion in a credible, transparent and theoretically valid manner. Principal-agent risk needs to be explicitly recognized in both the quantification and mitigation of risk, particularly where the pressure to match peer performance encourages a “follow the herd mentality” and creates systemic risk. Boards of directors and senior management must also be held more visibly accountable when they place their own interests above those of shareholders. In order for any of these changes to work in practice, however, the understanding of risk and risk management must improve dramatically - senior executives and board members should have more than a layperson’s knowledge of risk as a prerequisite to assuming their roles. The ability to manage risk has to keep pace with other business innovations.

13. Appendix B — Modeling ORM: Key Concepts

The Poisson Distribution

The Poisson probability distribution can be mathematically represented by the following function for non-negative integers:

$$P(X = x | \lambda) = \frac{\lambda^x}{x!} e^{-\lambda}, x = 0, 1, 2, \dots \text{ and } \lambda \geq 0$$

where λ is the parameter of the distribution.

An important characteristic of the Poisson distribution is that the sum of two independently distributed Poisson random variables is still a Poisson random variable with mean equal to the sum of the means of the two component distributions.

Since the mean and variance are equal, the Poisson is commonly used where the mean and variance are expected to be close. Example: number of wins in a season for a major league baseball team.

The Negative Binomial Distribution

The negative binomial distribution is given by the following probabilities on the set of non-negative integers:

$$P(X = x | r, p) = \frac{\Gamma(r + x)}{x! \Gamma(r)} p^r (1-p)^x, x = 0, 1, 2, \dots \text{ and } r > 0, 0 < p < 1$$

where $\Gamma(r) = (r-1)!$, where r is an integer and r and p are the two parameters of the distribution. Because it has two degrees of freedom, the negative binomial distribution is more versatile than the Poisson distribution.

And the mean and the variance of the distribution are given by:

$$\text{Mean}(X) = r(1-p)/p, \text{ Variance}(X) = r(1-p)/p^2$$

Since the variance of the negative binomial always exceeds the mean, this distribution is commonly used to model event frequencies subject to a high degree of variability. Example: the number of earthquake-driven tsunamis in a specific region of the world in a given time period.

The Binomial Distribution

The binomial probability distribution is given by the following function for non-negative integers:

$$P(X = x | p, n) = \frac{n!}{x!(n-x)!} p^x (1-p)^{n-x}, \quad x = 0, 1, 2, \dots, n \text{ and } 0 \leq p \leq 1$$

where n and p are the two parameters of such a distribution.

Using its probability generating function, it is easy to show that for a binomial distribution of X ,

$$\text{Mean}(X) = np, \text{ Variance}(X) = np(1-p)$$

It is easy to demonstrate that the binomial variance is always less than the mean. Consequently, the binomial distribution is generally used to model the frequency of events where there is relatively little variability. Example: number of “heads” resulting from ten tosses of a balanced coin.

Using Modified MLE for Fitting Truncated Data

The standard MLE likelihood function is the density function, but where loss data are truncated the likelihood function must be modified to describe the conditional likelihood, the likelihood of loss in excess of the reporting threshold. For left truncated data, one can achieve this by taking the original density function and dividing it by the probability of loss above the threshold, as shown below:

$$P(x_1, x_2, \dots, x_n | \theta) = \frac{\prod_{i=1}^n pdf(x_i | \theta)}{[1 - cdf(T | \theta)]^n}$$

where θ is the parameter vector

the x_i refer to the actual empirical data, and

T is the threshold value above which the data are collected.

Using the method above, one can use modified MLE to fit truncated data.

The Functional Form of Relevant Severity Distributions

Lognormal-Gamma:

The PDF of the LNG distribution is:

$$p(x; \mu, \sigma, \kappa) = \int_0^\infty \gamma(y; \kappa) \phi(x; \mu, \sigma^2 \times y) dy$$

where $\gamma(y; \kappa)$ is the density function for the gamma distribution and $\phi(x; \mu, \sigma^2)$ is the density function for the normal distribution with mean μ and variance σ^2 . Note in this representation the gamma distribution $G(\alpha, \beta)$ is restricted such that the mean is equal to 1, which effectively

transforms the gamma distribution into a one parameter distribution. In addition, since $\alpha = 1/\beta$, the variance $v = \alpha\beta^2 = \beta$; therefore β is linearly related to the kurtosis (κ) of the LNG distribution.

The CDF of the LNG distribution is:

$$CDF(x, \mu, \sigma, \kappa) = \int_{-\infty}^x p(t; \mu, \sigma, \kappa) dt$$

for all $x > -\infty$, where $p(x; \mu, \sigma, \kappa)$ as in the LNG density function given above.

Burr (type XII):

The PDF of the Burr Distribution is:

$$f(x, A, B, C) = \frac{A \cdot C \cdot (x/B)^C}{x \cdot (1 + (x/B)^C)^{A+1}}$$

for all $x \geq 0$, where $A > 0$, $B > 0$ and $C > 0$ are constant parameters of the distribution.

The CDF of the Burr distribution is:

$$F(x, A, B, C) = 1 - \frac{1}{(1 + (x/B)^C)^A}$$

for all $x > 0$, where A , B and C are as described above.

14. Appendix C — Glossary of Operational Risk Terminology

Basel II: the 2004 Basel Capital Accord, which includes a set of recommendations on banking laws and regulations issued by the Basel Committee on Banking Supervision. Formally referred to as the “International Convergence of Capital Measurements and Capital Standards — A Revised Framework. Published in June 2004.”

Burr Distribution: a three-parameter probability distribution used to model loss severity.

Business Strategic Risk: the risk of loss from an unanticipated change in an exogenous factor, such as the macro-economic environment, where no operational failure was present.

Censored Data: data that is censored above or below a specified value. Censoring occurs when the event is known to have happened but the value is known only to be below or above a specified amount. Data can be left censored or right censored (see also Truncated Data).

Coherent Risk Measure: a risk measure that satisfies properties of monotonicity (preserves rank order), sub-additivity (risk of A+B is less than or equal to the risk of A plus the risk of B), homogeneity (if all values in a portfolio are increased by a factor of X, the risk measure increases by a factor of X) and translational invariance (adding a constant value to a portfolio of risks, such as adding cash to an investment portfolio, reduces the risk of that portfolio by the same amount).

Conditional Expected Loss: the amount of loss that one expects given the occurrence of one specified event or presence of a specific condition. This can be calculated as the probability weighted mean of the severity distribution.

Conditional Risk: a measure of adverse deviation from the conditional expected loss.

Conditional Tail Expectation: the expected value of outcomes exceeding a stated threshold (also known as tail VaR (TVaR)).

Contributory Factor: a risk factor or controllable factor that contributes to loss frequency or loss severity.

Controllable Factor: an endogenous contributory factor that contributes to loss frequency or loss severity (such as inadequate training, unsafe working conditions or lack of supervision).

Cost of Risk: a financial measure of risk defined as the expected loss, plus the capital multiplied by the cost of capital.

Credit Default Swap: an agreement between two parties whereby one party pays the other a fixed periodic coupon for the specified life of the agreement; the other party makes no payments unless a specified credit event occurs.

Credit Risk: the risk of loss due to uncertainty in a counterparty's ability to meet its obligations.

Event: an occurrence; something that has happened (e.g., a loss).

Expected Loss: the average loss; this can be calculated as the probability weighted mean of the severity distribution. The aggregated expected loss, with a one-year time horizon, is the average amount of money that is expected to be lost in one year, on average. This represents the probability weighted mean of the aggregate loss distribution and can be calculated as the product of mean frequency and mean severity.

Expected Losses: nontechnical term used to describe losses one commonly observes — typically the smaller losses. This is NOT to be confused with “Expected Loss,” which refers to the statistical mean of a loss distribution.

External Loss Data: loss data that comes from sources outside an individual organization and that reflects the loss experience of other institutions.

Frequency: the number of events occurring during a specified time period.

Frequency Distribution: a statistical distribution of the number of events with associated probabilities. Common frequency distributions include Poisson, Binomial and Negative Binomial.

Generalized Pareto: this is a special class of statistical distribution function that is excess over threshold (i.e., data fits only over a threshold).

I.I.D.: independent and identically distributed. A requirement for modeling that stipulates that the individual data points must be independent (uncorrelated) with one another and must all exhibit the same distributional characteristics.

Impact: the aftermath of what happens after a negative event occurs; e.g., loss of reputation, write-down of an asset, etc.

Incident: synonymous with event; term often used by personnel outside the risk management profession.

Internal Loss Data: loss data that is collected by an individual institution and reflects its own loss experience.

Insurance Risk: the risk of loss above the expected from insured claims in the underwriting portfolio.

Legal Risk: term used to describe exposure to lawsuits. However, this is a misnomer in that “Legal Risk” is *not* a risk (event). Instead it is an impact/effect. (See also, Reputation Risk.)

Likelihood: commonly used synonymously with probability, used in likelihood x impact analysis; technically, likelihood is the probability of an outcome being randomly generated from a specific probability distribution, expressed in terms of the parameters of that distribution (this latter context underlies “Maximum Likelihood Estimation” and related statistical methods).

Liquidity Risk: term used to describe potential for loss from a liquidity squeeze. However, this is a misnomer in that “Liquidity” is *not* a risk (event) because one does not measure the impact in terms of loss. Instead one measures the loss in a reduction in the value of marketable securities (market risk). Where the liquidity squeeze is driven by exogenous factors it is a risk factor — something that exacerbates risk (e.g., market risk or credit risk). Where it is driven by endogenous factors (poor liquidity management), it represents a controllable factor (an operational failure).

Lognormal Distribution: a statistical distribution of a random variable whose logarithm is normally distributed.

Loss: an adverse financial outcome resulting from an event.

Loss Data Sharing Consortium: a group that collects loss data from various participating institutions and organizes the information in a standardized database available to members of the consortium.

Market Risk: the risk of loss in the market value of an investment portfolio (provided there was no operational failure).

Mean: the expected value (or probability weighted average) of a random variable X .

Median: the value in a data sample below which 50% of the observed data points lie when the entire sample is sorted in ascending order by size; also referred to as the 50th percentile of a sample distribution.

Mode: the value that occurs most frequently in a data set or probability distribution.

Monte Carlo Simulation: the use of generating random numbers for computing risk figures; expected loss and unexpected loss (for ORM).

Near Miss: either a non-event, which nearly became an event (two airplanes nearly crash), or an event that did not result in any significant injury, illness or damage but had the potential to do so.

Normal Distribution: the standard Gaussian distribution, commonly described as a “bell curve.”

Pareto Distribution: this is a single parameter statistical distribution function, which is also an excess over threshold class.

Poisson Distribution: a statistical distribution commonly used to model loss frequency. The Poisson distribution is parameterized by mean and variance where the mean and variance are equal. Two similar distributions which are also parameterized by mean and variance are the Binomial and the Negative Binomial. In the Binomial case the mean is less than the variance, in the Negative Binomial it is greater.

Reputation Risk: represents the risk of a loss in franchise or brand value. Reputation risk is not a risk (event); it is an effect or impact and measures the impact an event may have on future income through decreased revenues or increased expenses.

Risk: there are two definitions of risk. Formally, risk is a metric — a measure of the level of adverse deviation from the expectation. Informally, the term risk is used to describe the underlying event or incident (e.g., a fraud or system failure) which could create the possibility of an adverse outcome or a class of such events (e.g., operational risk or market risk).

Risk Capital: from a risk management perspective, this is the amount of capital required to cover adverse deviations from expected results and is the basis for the capital charge component of the cost of risk.

Risk Factor: a characteristic of an entity or sub-entity that impacts the relative frequency and/or severity of loss events (e.g., being a smoker increases the likelihood of death from heart and lung disease).

Risk Management: the discipline by which exposures to loss are identified, quantified, mitigated and either financed, hedged or transferred to another party.

Severity: the monetary (direct or indirect) value of a loss.

Severity Distribution: the statistical distribution of the value of individual or aggregate losses, with associated probabilities (this does NOT include loss frequency).

Solvency II: a set of regulatory requirements for insurance firms that operate in the European Union, which represent an updated version of the original Solvency regulations. The Solvency II regulations become effective in 2012.

Standard Deviation: a statistical measure of the variability of a population or probability distribution; the square root of the variance.

Truncated Data: data that is truncated at a specified amount. Truncation occurs when events below or above a specified amount are never recorded (see also Censored Data).

Unexpected Loss: the level of loss at a stated confidence level minus the expected loss; this represents the amount of adverse deviation beyond the expected loss at a stated confidence level.

Unexpected Losses: nontechnical term used to describe losses one does NOT commonly observe — typically the large losses. This is NOT to be confused with the term “Unexpected Loss,” which is a technical term and has a precise mathematical definition.

Value at Risk (VaR): two definitions: (1) the level of loss at a stated level of confidence or (2) the level of loss at a stated confidence level minus the expected loss.

Variance: a statistical measure of dispersion equal to the probability-weighted average of the squared distance of all possible values from the mean of the distribution.