1

Operational risk

Anything that can go wrong will go wrong.

—Murphy

1.1 INTRODUCTION

Operational risk has only in recent years been identified as something that should be actively measured and managed by a company in order to meet its objectives for stakeholders, including shareholders, customers, and management. These objectives include future survival of the company, avoidance of downgrades by rating agencies and remaining solvent for many years to come. Operational risk is becoming a major part of corporate governance of companies, especially in the financial services industry. This industry includes both banks and insurance companies, although they have somewhat different historical cultures in most countries. More recently in other fields such as energy, where trading and hedging activity mirrors similar activity in the financial services industry, operational risk is being recognized as a vital part of a broader enterprise risk management framework.

The definition of operational risk has not yet been universally agreed upon. In very general terms, operational risk refers to "risk" associated with the "operations" of an organization. "Risk" is not defined very specifically, nor is "operations." Generally, the term "risk" refers to the possibility of things going wrong, or the chances of things going wrong, or the possible consequences of things that can go wrong. "Operations" refers to the various functions of

the organization (usually a company such as a bank or insurance company) in conducting its business. It does not refer specifically to the products or services provided by the company. In banking, operational risk does not include the risk of losing money as a result of normal banking activities such as investing, trading, or lending except to the extent that operational activities affect those normal activities. An example of such an operational risk in banking is fraudulent activity, such as unauthorized lending where a loan officer ignores rules, or rogue trading in which a trader is involved in trading activity beyond limits of authorization. The well-known classic example of a rogue trader is Nick Leeson, whose activities resulted in the failure of Barings Bank, leading to its takeover by the ING financial services conglomerate.

operational risk is generic in nature. The operational risk concept applies to organizations of all types. However, the specifics of operational risk will vary from company to company depending on the individual characteristics of the company. For example, a manufacturer will be exposed to somewhat different operational risks than a bank or an insurance company, but many are the same. The risk of shutdown of the operations of a company because of IT failure, flooding, or an earthquake exists for any company. While the principles of operational risk modeling and management apply to all types of organization, in this book we will look at operational risk from the vantage point of a financial institution, such as a bank or insurance company.

Measurement and modeling of risk associated with operations for the financial sector began in the banking industry. Operational risk is one of several categories of risk used in enterprise risk management (ERM). ERM involves all types of risk faced by a company. Operational risk is one part only.

Many financial institutions have incorporated ERM into a new governance paradigm in which risk exposure is better understood and managed. The responsibility for the risk management function in a company often falls under the title of chief risk officer (CRO), a title first held by James Lam in the 1990s [72]. The CRO is responsible for the entire ERM process of the company in all its business units. Within the ERM process are processes for each risk category. Within the operational risk category, the responsibilities include:

Developing operational risk policies and internal standards
Controlling the operational risk self-assessment in each business unit
Describing and modeling all internal processes
Testing all processes for possible weaknesses
Developing operational risk technology
Developing key risk indicators
Planning the management of major business disruptions
Evaluating the risk associated with outsourcing operations
Maintaining a database of operational risk incidents
Developing metrics for operational risk exposure
Developing metrics for effectiveness of risk controls
Modeling losses using frequency and severity

Modeling potential losses using statistical tools
Calculating economic capital required to support operational risk

This book is primarily concerned with the last three items in this list.

In the banking sector, risks are generally described to be part of market risk, credit risk, or operational risk. In carrying out normal banking activities associated with investment in bonds issued by other companies, a loss in value due to overall interest rate changes in the market place is considered market risk, a loss in value due to a downgrade or bankruptcy of the issuer is a credit risk, but a loss due to an execution error, such as an error in timing or delivery of a trade, by the bank is an operational error.

At the time of writing this book, market and credit risk are much more well developed than operational risk. One of the reasons for this is the general dearth of publicly available operational risk data. This is in direct contrast to market risk and credit risk, for which data are widely available, particularly for the shares and bonds (and the related derivative products) of publicly traded companies. In the very recent past, the situation has changed as a result of gathering and sharing of historical data on operational risk losses. At the time of writing of this book, many organizations are building historical databases on past operational events in addition to building systems for the reporting and recording of new operational risk events as they occur. One major challenge, which is addressed later in this book, is how to combine data from several companies or the industry as a whole in building a model for a single company. This problem is sometimes called "scaling" because different companies are of different sizes and are therefore subject to risks of different sizes.

Although operational risk was originally defined to capture all sources of risk other than market and credit risk, several more specific definitions of operational risk have become well-known. In a paper published in 1998, the Basel Committee [9] on Banking Supervision (BCBS) identified the most important aspects of operational risk as relating to breakdowns in internal control and corporate governance. Effective internal controls should result in minimizing internal errors, fraud by staff, and failures to execute obligations in a timely manner. Failure of corporate governance can lead to poor internal controls.

The British Bankers Association [18] defined risk as the "risk associated with human error, inadequate procedures and control, fraudulent criminal activities; the risks caused by technological shortcomings, system breakdowns; all risks which are not "banking" and arise from business decisions as competitive action, pricing, etc.; legal risk and risk to business relationships, failure to meet regulatory requirements or an adverse impact on the bank's reputation; "external factors" including natural disasters, terrorist attacks and fraudulent activity, etc."

This all-encompassing definition was narrowed somewhat in the definition provided by the Basel Committee. In its consultative document on a capital adequacy framework [10] and its subsequent document on operational risk

[11], the BCBS defined operational risk as "the risk of losses resulting from inadequate or failed internal processes, people and systems or from external events." It includes strategic, reputational risk and systemic risks.

In its monograph dealing with capital requirements for insurance companies, the International Actuarial Association (IAA) [60] adopted the Basel Committee definition. It further noted that the definition is intended to include legal risks but exclude strategic, reputational risk and systemic risks. There remains some controversy over these items. Is a strategic decision that is later found to be in error really an operational risk? Is a loss in reputation an operational risk or simply the result of an operational risk event?

Operational risk in the banking sector is believed to represent about 30% of the total risk assumed by a bank. This contrasts with 60% for credit risk , 5% for market risks, and 5% for remaining miscellaneous risks. It is likely that the operational risk is proportionately smaller in the insurance sector. There have been some well-known significant operational losses in the insurance sector. The "misselling" of pension annuity products in the UK in the 1990s was a direct result of a lack of controls on the way in which the products were represented to potential customers.

It should be noted that losses from both internal and external events are included in the definition of operational risk. Internal events are events that result from the failure of some process or system operated by the organization. External events are those whose occurrence cannot be controlled by the company. The company can only mitigate the impact of these external events. It cannot prevent an earthquake, but it can ensure that its main computers are in an earthquake-proof building. In contrast, the occurrence of internal events is directly under the control of the company. Its risk management strategies can address both minimizing the occurrence of the event and mitigating the impact of the event when it occurs.

1.1.1 Basel II - General

The Basel Committee (in its "Basel II" framework) has been working on developing a framework for the determination of minimum capital requirements for banks. Included in the minimum capital requirement to be implemented in 2006 or later is a capital charge for operational risk. The minimum capital requirement falls under Pillar I of a three-pillar concept. The remaining two pillars relate to the supervisory process and market conduct. In this book, we shall focus on this first pillar only by addressing the question of how to probabilistically model losses arising from operational risk events. However, it is useful to understand the entire Basel II framework.

Pillar I: Minimum capital requirements There are three fundamental elements in the minimum capital requirement for regulatory purposes: the definition of regulatory capital, risk-weighted assets, and the minimum ratio of capital to risk-weighted assets. Risks are categorized into five categories:

1. credit risk
2. market risk
3. operational risk
4. liquidity risk
5. legal risk

Explicit and separate minimum capital requirements for operational risk have been added to the Basel II framework. Specifically for operational risk, there is a range of options for determining minimum regulatory capital requirements including building internal models of the company's operational risk profile. However, such minimum capital requirements will need to be supported by a robust implementation of the second and third pillars.

Pillar II: Supervisory review process The second pillar focuses on the prudential supervision by regulatory authorities of banks' capital adequacy as well as the banks' internal risk management systems. There are four key principles under Pillar II:

1. Banks should have a process for assessing their overall capital adequacy in relation to their risk profile and a strategy for maintaining their capital levels. This requires: i) strong board and management oversight; ii) sound capital assessment; iii) a comprehensive system for assessment of risks; iv) ongoing monitoring and reporting; and v) internal control review.
2. Supervisors should review and evaluate banks' internal capital adequacy assessments and strategies, as well as their ability to monitor and ensure their compliance with regulatory capital ratios. Supervisors should be able to take action when they are not satisfied with the results of this process.
3. Supervisors should expect banks to operate above the minimum capital ratios and should have the ability to require banks to hold capital in excess of the minimum.
4. Supervisors should seek to intervene at an early stage to prevent capital from falling below the minimum levels required to support the risk characteristics of a particular bank and should require rapid remedial action if capital is not maintained or restored.

For operational risk, this means that banks must monitor all operational risk events and have internal control processes in place that are transparent to banking supervisors. This will assist both banks and supervisors to understand past and potential future areas of losses from operational risk events. This better understanding of operational risk should have a direct effect on the operational risk by identifying areas where the bank can reduce both the frequency and the severity of those events.

Pillar III: Market discipline The objective of Pillar III is to encourage market discipline by developing a set of disclosure requirements that will allow

market participants to assess key pieces of information on the scope of application, capital, risk exposures, risk assessment processes, and hence the capital adequacy of the institution. These are especially useful when banks are given the authority to use bank-specific internal models in assessing their own risk profiles.

1.1.2 Basel II - Operational risk

Under Basel II, banks will be allowed to chose from three approaches: the basic indicator approach, the standardized approach, and the advanced measurement approach (AMA). Banks are encouraged to move along the spectrum of methods as they develop the capabilities to do more advanced modeling. Under the basic indicator approach for operational risk, banks are required to hold a flat percentage (15%) of positive gross income over the past three years. Under the standardized approach, banks' activities are divided into eight business lines: i) corporate finance, ii) trading and sales, iii) retail banking, iv) commercial banking, v) payment and settlement, vi) agency services, vii) asset management, and viii) retail brokerage. A flat percentage, ranging from 12% to 18%, is applied to the three-year average positive gross income for each business line. The minimum capital is the sum over all business lines. Both the basic indicator approach and the standardized approach are relatively crude methods that do not in any way allow banks to take credit for doing a good job in mitigating operational risk.

Under the AMA, banks are allowed to develop sophisticated internal models of the actual risks of the company including the interactions between them and any risk mitigation strategies used by the company. However, the bank is required to make significant investment in the management of operational risk. Specifically, i) a bank's board of directors and senior management must be actively involved in the oversight of the operational risk framework, ii) its operational risk management system must be conceptually sound and must be implemented with integrity, and iii) it must devote sufficient resources to the use of the AMA in the major business lines as well as in the control and audit areas.

Before full implementation, banks will be required to demonstrate that their systems are credible and appropriate by reasonably estimating unexpected losses based on the combined use of internal and relevant external loss data, scenario analysis and bank-specific environment and internal control factors. Furthermore, the bank must have an independent operational risk management function that is responsible for designing and implementing the bank's risk operational management framework. The bank's internal operational risk management system must be closely integrated into the day-to-day risk management processes of the bank. There must be regular reporting of operational risk exposures and loss experience to business unit management, senior management, and the board of directors. The bank's operational risk

management system must be well documented and reviewed regularly by internal or external auditors.

On the quantitative requirements for using the AMA approach, the Basel Committee [12] states:

> *Given the continuing evolution of analytical approaches for operational risk, the Committee is not specifying the approach or distributional assumptions used to generate the operational risk measure for regulatory capital purposes. However, a bank must be able to demonstrate that its approach captures potentially severe "tail" loss events. Whatever approach is used, a bank must demonstrate that its operational risk measure meets a soundness standard comparable to that of the internal ratings-based approach for credit risk,(i.e. comparable to a one-year holding period and a 99.9th percentile confidence interval).*
>
> *The Committee recognises that the AMA soundness standard provides significant flexibility to banks in the development of an operational risk measurement and management system. However, in the development of these systems, banks must have and maintain rigorous procedures for operational risk model development and independent model validation. Prior to implementation, the Committee will review evolving industry practices regarding credible and consistent estimates of potential operational losses. It will also review accumulated data, and the level of capital requirements estimated by the AMA, and may refine its proposals if appropriate.*

This book will focus on probabilistic models and statistical tools that can be used for building the internal models of operational risk that can be used under the AMA by a bank or an insurance company.

In the same document, the Basel Committee goes on to state:

> *A bank's risk measurement system must be sufficiently "granular" to capture the major drivers of operational risk affecting the shape of the tail of the loss estimates.*

This means that any model acceptable for the AMA must be very detailed and be sensitive to the possibility of extreme events. The shape of the tail of a loss distribution determines the likelihood of large losses. These need to be well understood because a single large loss can have a significant impact on a company. The issue of different tails of distributions is addressed throughout this book.

Continuing in the same document, the Committee states:

> *Risk measures for different operational risk estimates must be added for purposes of calculating the regulatory minimum capital requirement. However, the bank may be permitted to use internally determined correlations in operational risk losses across individual operational risk estimates, provided it can demonstrate to the satisfaction of the national supervisor that its systems for determining correlations are sound, implemented with integrity, and take into account the uncertainty surrounding any such correlation estimates (particularly in periods of stress). The*

bank must validate its correlation assumptions using appropriate quantitative and qualitative techniques.

This means that it is important to understand that there may be a possibility of diversification between operational risks. However, it is recognized that this may not be possible "in periods of stress," that is, in periods where everything seems to be going wrong. This idea can be captured through tail correlation, which is covered later in this book.

The Basel Committee document goes on to discuss data requirements for an internal risk measurement system. Internal loss data are crucial for the credible modeling of an organization's operational risk profile.

Banks must track internal loss data according to the criteria set out in this section. The tracking of internal loss event data is an essential prerequisite to the development and functioning of a credible operational risk measurement system. Internal loss data is crucial for tying a bank's risk estimates to its actual loss experience. This can be achieved in a number of ways, including using internal loss data as the foundation of empirical risk estimates, as a means of validating the inputs and outputs of the bank's risk measurement system, or as the link between loss experience and risk management and control decisions.

Internal loss data is most relevant when it is clearly linked to a bank's current business activities, technological processes and risk management procedures. Therefore, a bank must have documented procedures for assessing the on-going relevance of historical loss data, including those situations in which judgement overrides, scaling, or other adjustments may be used, to what extent they may be used and who is authorised to make such decisions.

Internally generated operational risk measures used for regulatory capital purposes must be based on a minimum five-year observation period of internal loss data, whether the internal loss data is used directly to build the loss measure or to validate it. When the bank first moves to the AMA, a three-year historical data window is acceptable ....

Thus building a loss data history is imperative to moving to an AMA for modeling risk capital. A bank's internal loss collection processes must meet the certain standards established by the Committee. The Committee is also very explicit about what data should be collected:

A bank's internal loss data must be comprehensive in that it captures all material activities and exposures from all appropriate subsystems and geographic locations. A bank must be able to justify that any excluded activities or exposures, both individually and in combination, would not have a material impact on the overall risk estimates. A bank must have an appropriate de minimis gross loss threshold for internal loss data collection, for example 10,000 Euros. The appropriate threshold may vary somewhat between banks, and within a bank across business lines and/or event types. However, particular thresholds should be broadly consistent with those used by peer banks.

The concept of a threshold becomes very important in the statistical analysis of operational risk losses. In statistical terms, ignoring small losses is called truncation of the data, in particular left truncation. It is important to know the truncation threshold for each recorded loss, because the threshold could be changed over time, or it could be different for different types of losses. It is particularly important when combining data from different banks into a single industry database or when combining external data with a bank's own data. The statistical issue of truncation will be dealt with thoroughly in this book. External data can be combined with bank data in a rigorous systematic way.

> *A bank's operational risk measurement system must use relevant external data (either public data and/or pooled industry data), especially when there is reason to believe that the bank is exposed to infrequent, yet potentially severe, losses. These external data should include data on actual loss amounts, information on the scale of business operations where the event occurred, information on the causes and circumstances of the loss events, or other information that would help in assessing the relevance of the loss event for other banks. A bank must have a systematic process for determining the situations for which external data must be used and the methodologies used to incorporate the data (e.g. scaling, qualitative adjustments, or informing the development of improved scenario analysis). The conditions and practices for external data use must be regularly reviewed, documented, and subject to periodic independent review.*

1.2 OPERATIONAL RISK IN INSURANCE

On the insurance side of the financial services industry, the development of capital requirements for operational risk has significantly lagged the developments in the banking sector. Insurers deal with risk and the management of risk on a day-to-day basis. However, this risk is primarily the risk inherent in the insurance contracts assumed by the insurer. In the jargon of risk management this type of risk is "business risk." As a business that is less transaction-oriented and less trading-oriented than banks, insurance companies have paid less attention to operational risk. But this is changing. At the global level, the International Association of Insurance Supervisors (IAIS) is in the process of developing a parallel but somewhat similar framework for the overall regulation of insurance. Its early work suggests three blocks of issues and a set of eight principles or "cornerstones" that will result in guidance to insurance companies. The three blocks of issues roughly parallel the three pillars of Basel II. The fifth principle dealing with absorption of losses states, "Capital requirements are needed to absorb losses that can occur from technical and other risks." The discussion of this principle refers directly to operational risk. The International Actuarial Association (IAA) book [60] reflects early work conducted by the IAA as a contribution to the IAIS effort in developing the regulatory framework.

Within Europe, the European Commission has initiated a "Solvency II" project for insurance regulation that also parallels Basel II but is applicable to European insurers. What we refer to as operational risks are somewhat covered by the term "risks that are difficult to quantify or to measure a priori." These include failings of management, major business decision risk, and failings in underwriting and claims handling. This list is rather short and misses some other key operational risks associated with failings in other operational areas such as sales. Furthermore, risk of external events has not been considered. At the time of writing this book, because of the dearth of available data and other difficulties in definition and measurement, operational risk is to be treated within the second pillar (governance process and controls) under Solvency II. However, as databases are developed, it is expected that Pillar I-type measurement and modeling will become the norm. Some insurance-related organizations are building data bases that make use of data coming directly from insured losses covering events that might be considered operational risks. The IAA book [60] recommends that operational risk should ultimately be handled with a Basel II approach under the first pillar. However, it would be reasonable to use a second pillar approach until insurance regulators, the industry, and the actuarial profession develop definitions and methods of measurement necessary for a first pillar approach.

Some recent external operational risk events in the US have pointed North American insurers in the direction of more active risk management (second pillar). The concentration of insurance brokerage employees in the World Trade Center on September 11, 2001 identified a personnel concentration risk to insurers. Extensive power blackouts in 2003 tested companies' computer systems and business continuity plans. The SARS epidemic in 2004 tested the abilities of banks and insurers in Hong Kong to continue operations as the movement of employees was severely restricted.

In this book, we will not try to define the various types of operational risk events that must be considered. This needs to be done at some point by every company, by industry groups, and by regulators. However those events are defined, in this book we will focus on modeling the chances that the event will occur and the consequences of the occurrence of the event.

1.3 THE ANALYSIS OF OPERATIONAL RISK

Various definitions of operational risk refer to events. The only events that are interesting to us from the point of view of operational risk are those that result in a loss. Inconsequential events are of no interest and as such are not treated as events for the purpose of the analysis of risk. As will be pointed out later in this and later chapters, the definition of an event is critical to any analysis because the definition affects how we count the events.

In order to measure the impact of operational risk, it seems natural to consider both how many events might occur and the potential impact of each

of those events. This approach to analysis is called a frequency/severity approach. This approach is commonly used in studies of losses in the insurance industry. The approach requires the risk analyst to separate the "count" or frequency of losses from the "impact" or severity of the losses. This is especially natural when the severity (per loss) does not depend on the number of losses, as is commonly assumed in modeling most risks. Consider, for example, errors made by automatic banking machines in dispensing money. An operational error can occur if the machine dispenses two bills that stick together as one bill. The number of errors increases as the number of machines increase, but the loss per loss event is unaffected.

Also for many types of losses, the severity or size of individual losses may be expected to increase over time as a result of normal inflationary growth. Similarly, expected frequency also increases as the number of exposure units (customers, employees, transactions, etc.) increases.

Risk managers use a variety of tools to assess and manage risk exposure. Frequency and severity are usually separately addressed in risk management strategies. Process control of internal processes can be used to minimize the frequency of risks associated with internal procedures. The development of internal policy and procedure manuals assists in defining what is acceptable activity and, more importantly, what is not acceptable activity. Process control systems, such as "six sigma" methodologies, can be employed to study and improve the performance of high-frequency transactions. Risk managers can also employ methods to control the severity of operational loss. For example, most organizations purchase directors and officers liability insurance coverage to protect against actions taken against directors and officers. Similarly, the company may purchase business interruption insurance to protect it against loss as a result of external events such as power grid failure (as occurred in the US and Canada in 2003), terrorist attack (as occurred in the US in 2001), or a hurricane (as occurs frequently in the US). Insurance usually carries with it a deductible so that a portion of the risk is still retained (or self-insured) by the company.

Risk managers will measure all risks consistently but may add special procedures for very large risks, often called "jumbo" risks. This reflects the different approaches to risks with different frequency/severity profiles. Risks can be classified according to whether the frequency is high or low and whether the severity is high or low. Here the terms "high" and "low" are used in a purely relative sense, that is, relative to other risks or, perhaps, relative to the size of the company. In general, we shall refer to the spectrum running from high-frequency-low-severity (HFLS) to low-frequency-high-severity (LFHS). It is not necessary to discuss high-frequency-high-severity risk because a history of this type of risk will certainly put a company out of business! Similarly, low-frequency-low-severity risk has little potential impact.

Model-based frequency/severity modeling is a main focus of this book. Because senior management (and regulators, rating agencies, and shareholders) are interested in the potential overall impact, frequency and severity modeling

are combined in the development of "aggregate" loss models. Frequency and severity modeling is done the same way for both LFHS and HFLS situations, at least in principle. However, LFHS will often attract additional analysis, that is, there will be serious analysis of the single possible big events that can bring down (or at least impair) the company. The next section discusses the model-based approach to operation risk management. Later, Chapter 7 will deal with possible extreme losses.

Model-based approaches to operational risk require significant amounts of data in order to calibrate the models to provide realistic outcomes. In the banking sector, a number of databases have been developed to help understand the frequency and severity of various types of operational risk. On the insurance side, at the time of preparation of this book, some organizations have begun to build databases. However it will be some time before their databases are broadly usable for calibrating models.

The remainder of this book is premised on the assumption that data will be available for the risk analyst. The tools in this book come mainly from the insurance industry, where actuaries have been involved in modeling the risk assumed by insurers in the insurance contracts that they sell.

1.4 THE MODEL-BASED APPROACH

The model-based approach involves the building of a mathematical model that can be used to describe, forecast, or predict operational loss costs or to determine the amount of capital necessary to absorb operational losses with a high probability. The results of the model can be used to better understand the company's exposure to operational risk, and the potential impact on the company of various possible mitigation and management strategies.

A model is a simplified mathematical description that is constructed based on the knowledge and experience of the risk analyst combined with data from the past. The data guide the analyst in selecting the form of the model as well as in calibrating the parameters in the model.

Any model provides a balance between simplicity and conformity to the available data. Simplicity is measured in terms of such things as the number of unknown parameters; the fewer the simpler. Conformity to data (or "fit") is measured in terms of the discrepancy between the data and the model or, equivalently, how well the model fits the data.

There are many models and many models with the same level of complexity; for example, the same number of parameters. Model selection requires consideration of both the mathematical form of the model and the number of parameters in the model. Model selection is based on an appropriate balance between the two criteria, namely, fit and simplicity. Appropriateness may depend on the specific purpose of the model.

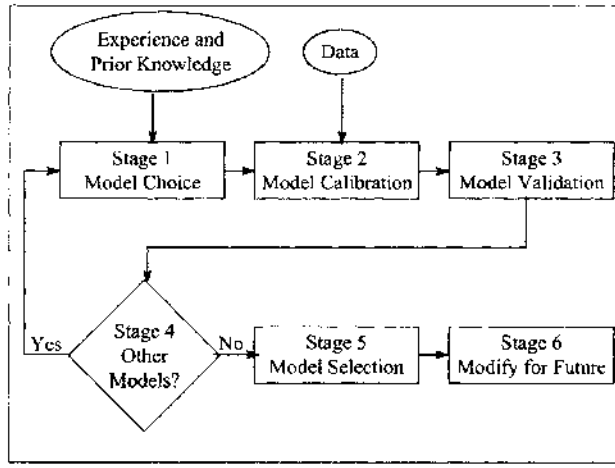

Fig. 1.1 The modeling process

1.4.1 The modeling process

The modeling process is illustrated in Figure 1.1, which describes six stages.

Stage 1 One or more models are selected based on the risk analyst's prior knowledge and experience and possibly on the nature and form of available data. In studies of the size of operational risk losses, a set of statistical distributions, such as lognormal, gamma, or Weibull, may be chosen.

Stage 2 The model is calibrated based on available data. In studies of operational losses, the data may be information about each of a set of actual losses. The model is calibrated by estimation of parameters based on the available data.

Stage 3 The calibrated model is validated to determine whether it conforms adequately to the data. Various diagnostic tests can be used. These may be well-known statistical tests, such as the chi-square goodness-of-fit test or the Kolmogorov–Smirnov test, or may be more qualitative in nature. The choice of test may relate directly to the ultimate purpose of the modeling exercise.

Stage 4 This stage is particularly useful if Stage 3 revealed that all models are inadequate. It is also possible that more than one valid model will be under consideration at this stage.

Stage 5 All valid models considered in Stages 1–4 are compared, using some criteria to select between them. This may be done by using the test

results previously obtained or may be done by using other criteria. Once the best model is selected, the others may be retained for later model sensitivity analysis.

Stage 6 Finally, the model is adapted for application to the future if the data were from the past and the model is to be used for the future. This could involve adjustment of parameters to reflect anticipated inflation or change in exposure from the time the data were collected to the period of time to which the model will be applied.

As new data are collected or the environment changes, the six stages will need to be repeated to improve the model. In practice, this should be a continuous process.

1.5 ORGANIZATION OF THIS BOOK

This book takes the reader through the tools used in modeling process beginning with organization of the remainder of this book is as follows:

1. Review of probability—Almost by definition, uncertain events imply probability models. Chapter 2 reviews random variables and some of the basic calculations that may be done with such models.

2. Probabilistic measurement of risk—Probability models provide a probabilistic description of risk. Risk measures are functions of probability models. They summarize in one number (or very few numbers) the degree of risk exposure. Chapter 3 provides a technical description of the state of the art in risk measurement analytics.

3. Understanding probability distributions—In order to select a probability model, the risk analyst should possess a reasonably large collection of such models. In addition, to make a good a priori model choice, characteristics of these models should be available. In Chapters 4 and 5, a variety of distributional models are introduced and their characteristics explored. This includes both continuous and discrete distributions. The range of distributions in these chapters is much greater than in most standard books on statistical methods.

4. Aggregate losses—To this point the models are either for the amount of a single loss or for the number of payments. What is of primary interest to the decision-maker, when modeling operational losses, is the total possible amount of losses. A model that combines the probabilities concerning the possible number of losses and the possible amounts of each loss is called an aggregate loss model. Calculations for such models are covered in Chapter 6.

5. Extreme value theory—In studying operational risk, special attention must be paid to high-impact extreme, but rare, events. This is the subject of Chapter 7.

6. Copula methods—Dependencies among risks must be understood so that appropriate credit can be given for diversification when risks may exhibit correlation of some type. Chapter 8 introduces many relevant copula models.

7. Review of mathematical statistics—Techniques of mathematical statistics are needed to calibrate models and make formal choices among models based on available data. While Chapter 9 is not a replacement for a thorough treatment of mathematical statistics, it reviews the essential items needed later in this book. The reader with a good background can skim this chapter quickly.

8. Calibrating parametric models—Chapters 10 and 11 provide methods for parameter estimation for the continuous and discrete models introduced earlier. Model selection is covered in Chapter 12.

9. Chapter 13 applies special statistical methods for the study of very large possible losses, the jumbo risks that require deeper individual study.

10. Finally, in Chapter 14, we consider estimation methods for multivariate models, in particular the estimation and selection of copulas.

This book provides many tools necessary for carrying out the modeling of operational risk for an organization. However, we do not attempt to discuss building an operational risk management program for an organization, a program that would include process controls and other aspects of risk management. As such, our scope is relatively narrow. Within this narrow scope, the treatment of topics is quite comprehensive and from a practical perspective. We have not incorporated some topics that are, at this stage, more interesting to the theoretician than the practicing risk analyst.