Quantifying Event Risk: The Next Convergence

ROBERT CESKE, JOSÉ V. HERNÁNDEZ, AND LUIS M. SÁNCHEZ

ROBERT CESKE is managing director and head of operational risk management at NetRisk in Greenwich, Connecticut.

JOSÉ V. HERNÁNDEZ is vice president of the methodology group at NetRisk.

LUIS M. SÁNCHEZ is vice president of the advisory group at NetRisk.

While the existence of event risk is not a new phenomenon for financial services companies, its measurement, as part of integrated risk management programs, has been the subject of recent focus. Property and casualty insurers have measured components of this risk class as part of the pricing decision-making process. However, few financial services firms are measuring this risk class on an integrated basis (i.e., including insurance liabilities with other risks, such as unauthorized trading/investment activities, improper sales practices, fraud, etc., which are included in the broader classification of "operational risk"[1]).

While certain types of financial services firms view event/operational risk as a revenue center, through underwriting risk transfer products, all financial services firms are exposed to the direct and indirect (especially reputational) costs of operational risk events. Whether for revenue generation or risk management purposes, the first step for any firm that wishes to manage its operational risk is to measure that risk. This paper will explore ways in which this may be done in practice.

One of the first decisions you need to make is which of many modeling approaches to use. As with any relatively new discipline, there are a number of methodologies to choose from, each with its advantages and target applications. The ultimate choice of the methodology/methodologies to use in your institution depends on a number of factors, including:

- Time sensitivity for analysis;
- Resources desired and/or available for the task;
- Approaches used for other risk measures;
- Expected use of results (e.g., allocating capital to business units, prioritizing control improvement projects, satisfying regulators that your institution is measuring risk, providing an incentive for better management of operational risk, etc.);
- Senior management understanding and commitment; and
- Existing complementary processes, such as self-assessment.

Institutions that have already been measuring operational risk on an enterprise-wide basis typically started with simpler, top-down approaches and migrated (or are migrating) to the more sophisticated bottom-up and hybrid actuarial approaches. The top-down approaches are generally less costly and time-consuming to use than the bottom-up and hybrid approaches. However, measuring enterprise-wide operational risk — even with relatively simpler approaches, such as Key Control Indicators, CAPM-based Modeling and Top-down Actuarial Models — can provide valuable insight for risk managers and other members of senior management.

Exhibit 1 provides an overview of some of the alternative methodologies, their applications and practical issues in using these methods.

Exhibit 1
Operational Risk Methodology Alternatives

Approach	Description	Application	Practical Issues
Top-down approaches	Top-down approaches estimate operational risk for the institution based on firm-wide or industry-wide data. Results from top-down approaches are then typically allocated to business units, based on pre-defined criteria.	Top-down approaches are particularly useful as a starting point to quantifying operational risk. These approaches are relatively quick and easy to use and generally require limited data. Often top-down approaches are combined: one approach (such as earnings volatility or CAPM-based modeling) is used to estimate firm-wide risk figures, which are then allocated to business units using another approach (such as key control indicators or activity-based).	• Main focus at the corporate level. • Almost no information at the business process level. • Limited value for risk management.
"Residual" approach	The simplest operational risk measure, the residual approach assigns operational risk, based on the following equation: Operational = Total Capital – Market Risk – Credit Risk	Although many institutions use relatively more sophisticated methods for allocating operational risk to business units, this "measure" is quite common. The residual approach is useful if no "true" measures for operational risk are in place, but the institution wishes to allocate operational risk capital to business units.	• Difficult to manage a residual. • Operational risk measurement is based on the firm's capital, not its risk. • Provides no real incentive for the firm to improve its control environment or other risk factors.
Activity-based	Activity-based measures estimate operational risk based on objective measures of the level of business (e.g., number of employees, operating expenses, trade volumes, etc.). Also used as an allocation method for other top-down approaches.	Rumored to be a leading candidate for regulatory operational risk capital charges. Activity-based measures are simple to determine and objective. They are often used within an institution for allocation of other top-down risk measures.	• Assumes existing capital is sufficient. • Not necessarily a direct measure of risk, but rather a measure of operations resources. • If based on expense, creates a perverse incentive, whereby firms and/or business units can lower risk score by cutting controls/expenses. • Assumes operational risk grows with the size of the institution – which is not necessarily true, particularly for catastrophic operational risk events.
Key risk indicators	These models identify factors and apply them as an early warning tool or to estimate expected losses, through the application of regression techniques or trend analysis. Key risk indicators include factors, such as audit ratings, employee turnover, transaction volume or recent operational loss rates.	In application, this methodology is quick to implement, easy to quantify and capable of business unit analysis. On an on-going basis, this approach can provide early warning signs and become a good relative indicator of various business units. It is typically used as a technique for allocating operational risk capital, determined by other means.	• Does not calculate risk. • Difficult to ensure key risk indicators are predictive. • Diminishing marginal effectiveness. • Could lead to misallocation of management focus on measured controls versus non-measured controls.

Exhibit 1
Continued

Approach	Description	Application	Practical Issues
Operating leverage	This approach seeks to understand the relationship between cost structures (the mix of fixed and variable expenses) and revenues. Risk is potential deviation between expected and actual expenses. The expected operating expense is modeled using historical time series of revenues and expenses and regression techniques.	Transparent method for assessing high frequency, low probability events.	Model uses historical earnings and therefore is unlikely to capture low frequency/high severity events. Relies on historical earnings times series: • Statistically significant time series may not exist for all businesses. • Relationship between costs and revenues may not be stable through time. • Risk assessment evolves through time, but changes in risk measure are likely to lag changes in actual risk profile.
Earnings volatility	Earnings volatility is a macro-level approach that defines the residual income or (loss) of total earnings volatility, less market volatility and credit volatility, as the earnings volatility attributable to operational risk.	Where the focus of senior management is primarily on periodic returns and net income, the residual approach provides a quick, ex-post analysis of net income to determine the source of such income/loss. As the sources of incomes and losses are determined, the risks associated with these results are identified and linked to the respective uses of capital. This approach is often used for reporting to market analysts and shareholders and is consistent with EVA (Economic Value Added) measures.	• Difficult to manage a residual. • Does not delineate below macro-level to business units. • Assumes existing capital is sufficient. • Restricted to ex-post analysis.
CAPM-based modeling	The CAPM method for measuring operational risk is based on the Capital Asset Pricing Model used in corporate finance models. CAPM-based models attempt to attribute all components of a firm's valuation (not just the volatility of earnings or leverage components, as in the prior models) to risk classes, including operational risk. These models maintain the fundamental assumptions of CAPM, (i.e. efficient markets theorem, zero arbitrage, etc.).	Useful for measuring operational risk at a firm-wide level. CAPM-based measures are typically used in conjunction with activity-based or key risk indicator measures for allocation of top-down operational risk measures to business units. CAPM models are quick to implement, yet work well in conjunction with EVA measures.	• Difficult to allocate to lines of business, products. • Difficult to link management action to capital implications. • Risk-reward calculations are distorted below the total bank level. • Textbook theorem – does not account for "real world" irregularities.

EXHIBIT 1
Continued

Approach	Description	Application	Practical Issues
Top-down actuarial models	Top-down actuarial modeling uses actuarial methods to model the probability/frequency of operational risk events and the resulting severity. The frequency and severity curves combine to produce a profile (aggregate loss distribution) of operational Capital-at-Risk at every confidence level. Top-down measures typically rely on external loss databases and management scenarios for generating frequency and severity curves.	Where operational risk loss data is available, but in limited quantities – an actuarial top-down approach can be used to estimate firm-wide Capital-at-Risk for operational risk. The data used in actuarial models is quite useful in identifying key areas for management to focus on improving controls (i.e., learning from everyone else's mistakes). Top-down actuarial model outputs are typically allocated to business units, using activity-based or key risk indicator methods.	• The limited availability of industry-wide data in the high frequency/low severity case, in particular. • Estimating frequencies with limited internal data is particularly difficult.
Bottom-up approaches	Bottom-up approaches model the relationship between causes and its corresponding effects (losses) at individual business unit or process level. The results are then aggregated to determine the risk profile of the institution.	Of all approaches to operational risk measurement and management, bottom-up approaches represent the highest level of specialization. With strong managerial support and adequate resources, a bottom-up approach will generate high-quality results. Bottom-up approaches capture the idiosyncrasies of specific business units or processes and the quality of their associated control environments. These methods should be used when significant resources are available with strong senior management commitment.	• Often require costly development of detailed models. • With causal network and related methods (statistical quality control and connectivity), these methods tend to reveal only operational risks related to narrow processes (such as back-office processing), and fail to capture extreme events.
Causal networks	Relate causes and effects through conditional probabilities and the application of Bayes theorem. The links between causes and effects are displayed using graphical models, which are based on detailed analysis of the workflows of the process.	Where product lines and and services are highly interrelated, the application of causal network methodologies should prove beneficial. The process of applying various probabilities and/or correlations will help identify potential operational risk loss events and illustrate more clearly the possibility of extreme losses in high stress environments.	• Dependent upon the maintenance of adequate historical records identifying previously incurred losses, their effects and contributory causal factors. • Not appropriate for new business lines with no previous activity. • Restricted to ex-post analysis. • Institutions are continuously changing in structure, thereby creating difficulties when analyzing historical events.
Statistical quality control and reliability analysis	These approaches are closely related to causal networks. Manufacturing companies, to minimize errors with the subsequent savings in operating costs, have commonly used them.	When customer satisfaction is one of the key factors to product and/or services success. An example of these methods is the Six Sigma process used by GE.	• Extremely difficult to achieve such high levels of accuracy. • Not efficient in heterogeneous product offerings.

Exhibit 1
Continued

Approach	Description	Application	Practical Issues
Connectivity	This approach consists on analyzing the connections between individual exposures within a process or business unit and combining them (using a connectivity matrix) to estimate potential losses. The results are aggregated to obtain total potential losses for the institution. This approach is intended to capture dependencies (failure effects propagate across processes or business units), which may be used as the basis for simulation.	Similar to causal networks, connectivity analysis is most valuable where the inter-correlation of product lines and services are high. This methodology will assist management in determining the probability of extreme loss events as well as the mid-severity, more frequent loss effects.	• Dependent upon subjective input from the respective analyst – that is the levels of correlation applied, the scope of correlations simulated.
Bottom-up actuarial models	As with top-down actuarial modeling, bottom-up modeling uses actuarial methods to determine frequency and severity curves for calculating operational Capital-at-Risk profiles. However, with bottom-up methods, risk profiles are determined at the business unit level and then aggregated to determine firm-wide profiles.	Bottom-up actuarial modeling is performed by institutions with sophisticated risk measurement capabilities. Calculating operational risk using this method requires substantial data (or management estimates) from each business unit.	• The time and resource commitment required to collect sufficient data is high. • Estimating frequencies with limited internal data is particularly difficult.
Predictive modeling	This approach attempts to identify predictive factors in order to determine the probability and severity of future loss events. Such practices have been utilized in the credit card business where financial institutions attempt to determine the probability of credit losses based upon a pre-determined set of predictive factors. It uses statistical methods by which historical data is used to predict the likelihood of future events. This model uses regression or factor analysis to determine correlations between easily quantifiable statistics and the likelihood of an event occurrence. These correlations are then utilized to highlight potential risk areas.	The application of predictive modelling is often seen as the end goal of most bottom-up approaches. The ability to identify those factors which may prove predictive is the result of a trial and error process, as well as the development and analysis of historical cause and effect relationships. The business environment should be conducive to a strong statistical relationship between easily quantifiable events and future losses. Also, such models require an environment in which subjectivity is not relied upon.	• Significant resources and efforts required to obtain this level of analysis. • Relies upon the assumption that history is a good predictor of the future. • Difficult to predict ahead of time, which factors may work well in this model.
Hybrid approaches	Hybrid approaches combine top-down and bottom-up approaches to take advantage of their complementary aspects. Typically they use external databases of industry-wide historical operational losses, internal data and business line-specific scenarios to determine aggregate loss distributions at several institutional levels. External losses are modified to reflect the state of the institution and its business lines.	A combined methodology is commonly used where some business units have sophisticated operational risk measurement and management tools and techniques in comparison to other business units with minimal operational risk management attention. These hybrid approaches represent good interim management tools as more sophisticated and comprehensive methodologies and processes are being developed and implemented.	• Approach may not provide a comprehensive view of all operational risk exposures. • Conflicting results and/or duplication of coverage may require further analysis to determine the appropriate factor or result.

Once you have selected the preferred modeling approach to use, two major questions remain:

1. Are there any models available from third parties, or do I need to build them myself?
2. What will you do for the inputs to the models?

Although these questions are relatively easy for market risk and getting easier for credit risk, they are not as straightforward for operational risk. In fact, because of the particular difficulty of Question #2, risk managers often find themselves faced with another question:

3. What methodological changes are necessary in light of limited data?

In this article, we will help you answer these questions, by giving practical insights and constructive alternatives for working through these issues. In particular, we will focus on sources of inputs to operational risk models and on methods for using available data (or data that require relatively less work to collect) in place of data from "ideal" sources. We will then discuss methodological refinements and techniques for bridging the gap between measuring operational risk in theory, which requires large populations of loss data or in-depth models of complicated processes, and measuring it in practice, where more data are not available.

MODEL AVAILABILITY

One of the characteristics of a new market and/or risk management discipline is that relatively few service providers have models and/or systems available to that market. This is the case for operational risk measurement models. While there are many providers of control (self)-assessment services and software (most notably Frasin, Horizon, and MethodWare), few companies offer operational risk models.

If you are interested in using bottom-up or hybrid actuarial approaches, there are a few choices. NetRisk offers the web-based RiskOps™ system, which has a loss database and sophisticated modeling capabilities. PricewaterhouseCoopers (PWC) also offers a loss database and models, dubbed OpVaR, in conjunction with its consulting services, while Algorithmics' WatchDog system is being developed using causal network approaches.

If you choose to use a top-down approach to measuring your operational risk (except with a top-down actuarial model, which may be accomplished with hybrid models) you will need to develop your own modeling capabilities. However, since top-down approaches are significantly less complicated than other approaches (see Exhibit 2), the effort required to build such models is not significant.

MODEL INPUTS

The answer to the second question depends to a great extent on the modeling approach you choose to follow. Data requirements for top-down approaches are significantly less than for bottom-up and hybrid actuarial approaches — one reason for their relative ease of use. Although most institutions begin with less sophisticated top-down approaches to measuring operational risk, the leading institutions in risk management have focused efforts on the newer frontiers of operational risk measurement, namely bottom-up and hybrid/actuarial approaches.

This evolution toward more sophisticated measures of risk is consistent with what we have seen with prior risk classes in the risk management industry. Leading institutions, with the assistance of vendors and other service providers, work to develop improved methodologies and techniques to meet the practical issues in using more sophisticated risk measures. The remainder of this paper

EXHIBIT 2
Frequency and Severity of Loss Events

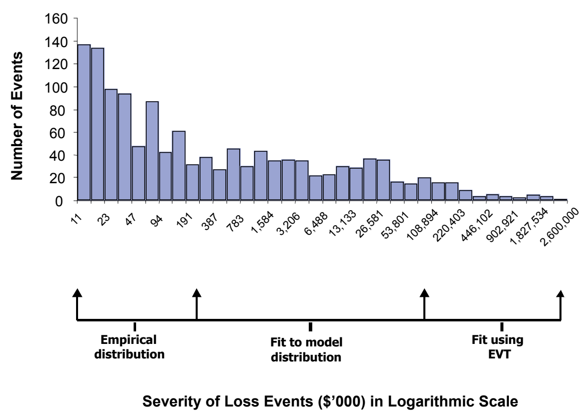

Different sections of the severity distribution can be modeled using different methods, depending on data availability. The vertical axis represents the number of operational losses and the horizontal axis represents their corresponding severities. Severities are displayed in a logarithmic scale to facilitate the visualization of the range of severity values.
Source: RiskOps™.

will focus on identifying techniques and approaches to help institutions measure operational risk with the more sophisticated bottom-up and hybrid actuarial approaches.

Finding inputs for operational risk models is significantly more difficult than for market risk. Operational risk loss events or failures in processes (two primary inputs to models) do not happen with anywhere near the frequency of changes in market values of portfolios. In addition, data on credit-related events have already been collected for many years, making inputs to credit models more easily available. As a result, operational risk managers are faced with the following choices for obtaining model inputs:

- Collect data
- Use others' data as a proxy for yours
- Use educated opinions to estimate inputs
- Estimate/extrapolate data based on limited samples

Obtaining inputs based on the first three data sources involves time and effort, while the last is accomplished through methodological refinements, which is discussed in the next section. Generally, measuring operational risk will require the collection and use of internal data, external data and management scenarios.

Internal Data

Ideally, most data used in operational risk modeling would come from your institution's internal sources, because they most closely reflect your institutions' control environment, business mix, etc. However, few institutions have collected large volumes of internal data, and establishing robust collection mechanisms and gathering such data is time-consuming. Additionally, internal data tend to be biased toward high frequency, low severity events, missing the low probability, high severity operational loss events that can take the institution into a state of financial distress. Furthermore, because of the changing nature of institutions, even internal data has to be scaled, based on activity levels, and filtered.

Internal data is particularly important in determining frequency distributions, as these distributions are particularly sensitive to the internal institution characteristics. In addition, internal data helps determine the shape of the severity distribution and is used to relate external data to the specifics of the particular institution. However, internal data generally lacks information on catastrophic events, because they occur with a low probability.

In addition to using your institution's internal data for measuring operational risk, other institutions' internal data can provide valuable insights into your relative operational risk profile, as well as provide statistically significant data for modeling purposes. However, others' internal data need to be related to the characteristics of your institution — namely the business unit mix, activity level, geography, control environment, etc. By collecting data on business/activity levels (also referred to as "scaling factors") and characteristics that influence the frequency and severity of operational risk losses ("predictive factors"), operational risk measurement can positively influence management decision-making.

By collecting information on the different characteristics of the institutions and business units within these institutions that have sustained operational losses, one can begin to determine the causal relationships around operational risk events. For example, collecting scaling factors can relate the impact of operational loss events on different institutions. Factor analyses on predictive factors can help relate characteristics of institutions and/or business units (such as the staff experience, relative age of processing systems, etc.) and the probability/severity of operational loss events. Such "predictive modeling" gives management important information for making strategic decisions, such as "is it better to cut expenses through hiring less experienced staff or spending less on technology"?

However, gathering others' internal data will require collaboration among financial institutions — by means of a consortium managed by an independent party. Unfortunately, there is no consortium yet in existence, but progress is likely to be made soon (see sidebar).

External Data

External data, or data that are available from public sources, contains information on both low and high severity events. However, high severity loss events are more likely to make the financial press and other public sources. As with others' internal data, external data may not be directly relevant to your institution. Therefore, it is necessary to filter external data, based on their relevancy, and to scale them to reflect institutional characteristics.

Severity distributions are scaled using variables such as size, number of employees, trading volume, etc. Frequency distributions are scaled using the results of bottom-up approaches or management assessments. External data, along with management scenarios, are valuable in modeling the tail of severity distributions —

which, given the kurtosis of operational risk distributions, is an extremely important region in the analysis of this risk class.

Management Scenarios

In many cases data may not be available at all (e.g., Y2K problem). In those cases it is necessary to complement internal data and data from external sources with management-generated loss scenarios. These scenarios must be developed in close collaboration with the business lines, or can be based on internal self-assessments already performed by most financial institutions. However, management scenarios must include frequency and severity estimates that are quantitative (i.e., high, medium, low is not sufficient).

Loss scenarios can be described by:

- Separating the losses due to a particular event into several potential loss sizes and assigning them their relative frequencies.
- Specifying parametric models (and their parameters) for both frequency and severity distributions. Severity distribution models utilized in practice include lognormal, gamma, and Pareto distributions. Frequency distribution models include Poisson, binomial, and negative binomial.

Once the data to be used in measuring operational risk have been collected, you must then reconcile overlaps from different sources and fill in holes with additional scenarios. This process requires a solid understanding of your institutions' business units, activities, and an appreciation for the diverse types of operational risk.

ARE OPERATIONAL RISK DATA CONSORTIUMS THE ANSWER TO THE DATA PROBLEM?

One alternative to the operational risk data problem that has been gaining favor over the past year is sharing public and disguised non-public operational risk information by means of a data consortium. Although such an initiative seemed unlikely one year ago, several initiatives are moving quickly.

The vision of an operational risk consortium is to create one common definition for operational risk and one comprehensive database for all industry practitioners to supply with information and benefit from the respective lessons learned. The data consortium would need to be organized with a third party:

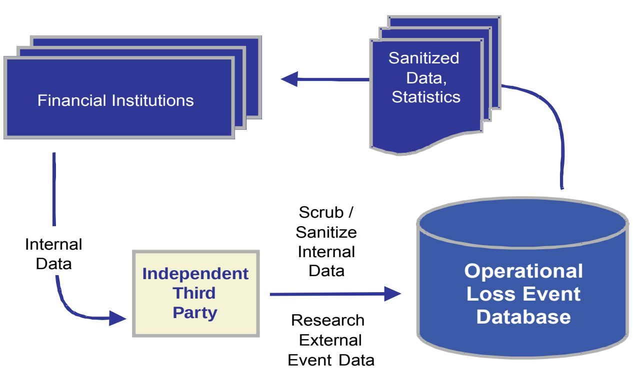

- Implement agreed-upon uniform data standards.
- Scrubbing internal data (conceal identity of contributor).
- Researching external loss events.
- Managing internal and external data.

Such a consortium might be organized as follows:

If open, the consortium database could then be used by banks to improve risk measurement and management, by insurance companies to develop and price risk transfer products to be used to alter risk profiles, by software vendors to develop risk management systems, and by consultants for specific analyses, such as benchmarking studies and trend analyses. Alternatively, data could only be made available to contributing institutions.

The database would contain non-public internal operational loss events to be shared, with the sources of such losses concealed. Such a forum would facilitate the sharing of lessons learned at various institutions, assist in the development of identifying predictive factors and provide an opportunity to decrease the systemic risk of the financial industry. The development of a consortium to share non-public operational risk data will also likely influence regulatory initiatives and meet the needs of industry practitioners.

Other industries share potentially sensitive information on confidential bases. Examples include insurance companies' sharing of claims information and energy companies' sharing of information on catastro-

METHODOLOGICAL CHANGES

As a practical matter, the sparsity of operational risk data will not be overcome overnight (even with pooled efforts). It is therefore necessary to explore changes to modeling techniques and/or refined methods for measuring operational risk with limited data. This section describes several alternatives for answering our third major question: *What methodological changes are necessary in light of limited data?*

In order to measure operational risk, using bottom-up and hybrid approaches, it is necessary to determine the probability distribution of aggregate operational losses over a period of time relevant to the institution (generally, one year). Depending upon the application, aggregate losses can be calculated at several levels, including firm wide, business line, loss type, loss severity, data source, etc.

Aggregate loss distributions are skewed (asymmetric) and have fat-tails. This is due to the fact that low severity events have a high frequency, while high severity events occur with a low frequency. Aggregate losses over a given period are the result of the total number of losses (governed by the frequency distribution) and the severity of those losses (described by the severity distribution). Aggregate losses are obtained by compounding frequency and severity distributions. This can be achieved through Monte Carlo simulation or through the application of analytical techniques. Even though aggregate losses can be modeled directly, keeping frequency and severity separate provides flexibility and insights into the measurement processes.

Frequency distributions are discrete ($p_k = PR\{N = k\}$, $k = 0, 1, 2, \ldots$, where k is the number of operational loss events). When some, but not a significant

phes. Can the financial services industry follow these examples to collect the data necessary for modeling operational risk?

Such a consortium has a number of benefits to all participants. It would:

- Minimize data handling costs for individual institutions;
- Create flexibility for accessing loss information, while maintaining confidentiality;
- Build upon existing methodologies and best practices;
- Assist in the identification of predictive factors; and
- Help reduce the likelihood of regulatory capital directives that are activity-based, not risk-based.

One of the most promising initiatives is a consortium sponsored by The Global Association of Risk Professionals (GARP) with NetRisk as the managing agent. The Multinational Operational Risk Exchange (MORE) is having great acceptance among integrated financial services companies. MORE is targeted to financial institutions and owned by members who provide operational loss data to the consortium. Participants in the consortium maintain ownership of the data, and its respective uses.

MORE has received early commitment from several leading financial services companies, and has attracted interest from over forty institutions globally.

MORE plans to 1) establish a secure framework for financial institutions to share non-public operational loss data, 2) provide the necessary elements to quantify operational capital at risk, 3) collect information on predictive factors that will allow managers to focus operational risk mitigation efforts in areas with the highest expected return on these efforts, and 4) influence regulatory decisions on operational risk capital requirements.

Nevertheless, potential members of the MORE consortiums have expressed their concerns over the following issues:

- Integrity and confidentiality of data;
- Disparity in number of data points supplied;
- Disparity in quality of data;
- Commitment before full disclosure of methodology;
- Commitment before reaching "critical mass"; and
- Cost/benefit analysis and estimated budgeting.

The largest concern is the issue of confidentiality. Can internal loss events be sanitized to allow the benefit of sharing while maintaining the integrity of individual institutions? Could regulators demand access to such a database? Could contents affect the outcome of litigation?

While there are certainly risks, it seems that the benefits of a data consortium are beginning to outweigh the risks in the minds of many institutions. This is evidenced by the growing interest in the MORE consortium and commitment from early leaders.

amount of, data are available, the frequency distribution can be modeled using the empirical distribution or by estimating the parameters of a model distribution. The latter is achieved by applying fitting techniques such as maximum likelihood estimation (MLE). Model distributions often used in practice include the Poisson and negative binomial distributions.

Similarly, severity distributions ($F(x) = \Pr\{X \leq x\}$) can be modeled by their empirical distribution if enough data are available, or by smoothing the distribution when there are insufficient data. Smoothing can be achieved by fitting the empirical distribution to a model distribution (parametric methods) or by using the data available without any assumptions regarding the shape of the underlying distribution (nonparametric methods). In parametric approaches, the empirical distribution is smoothed by fitting a model distribution such as lognormal, exponential, Weibull, Pareto, etc. MLE is commonly applied to estimate the parameters of the model distribution. For those cases when enough data is available, nonparametric smoothing methods, such as kernel density estimation can be applied. Kernel density estimation replaces individual sample points by distributions ("kernels") centered on those points and then aggregates the kernels to obtain the distribution. This technique creates a smoother distribution that is entirely based on the given data points.

In some cases, different amounts of data are available at different severity levels. If this is the case, then the severity distribution may be modeled by applying different distributions to the different severity regions. For example, as in Exhibit 2, the low severity region can be modeled by its empirical distribution, the medium severity region may be fitted to a lognormal distribution and the high severity region can be approximated through the application of extreme value theory (EVT).

Finally, loss distribution for each case should be combined to the desired aggregation level, which may be firm wide or by business unit or type of loss. Aggregate loss distributions can be combined through convolution (under the assumption of independence) or through simulation (if information on correlation is available).

Extreme Value Theory

Extreme value theory provides a useful framework for the application of parametric smoothing methods to fit the tail of loss distributions beyond a certain level.[2] That is, EVT helps you to estimate the shape of the distribution deep into the tail, where relatively little data are available. For example, operational losses of the order of a few million dollars occur relatively frequently and their distribution may be modeled using their empirical distribution. On the other hand, loss events with direct losses exceeding $1 billion, occur infrequently, but when they occur they have dramatic consequences and thus the severity distribution in this region requires careful modeling. This is the domain of application of EVT.

The main result of EVT that is relevant for operational risk measurement is the outcome that losses beyond a high enough threshold follow a member of the class of generalized Pareto distributions (GPD). The GPD $G_{\xi,\mu,\beta}(x)$ is given by

$$G_{\xi,\mu,\beta}(x) = \begin{cases} 1-\left(1+\xi\dfrac{x-\mu}{\beta}\right)^{-1/\xi} & \text{for } \xi \neq 0 \\ 1-\exp\left(-\dfrac{x-\mu}{\beta}\right) & \text{for } \xi = 0 \end{cases}$$

where $\beta > 0$, $x - \mu \geq 0$ when $\xi \geq 0$ and $0 \leq x - \mu \leq \beta/\xi$ when $\xi < 0$. The parameters ξ, β and μ are called shape, scale and position parameters, respectively. Distributions for the $\xi > 0$ case are Pareto distributions, for the $\xi = 0$ are exponential distributions and for the $\xi < 0$ case are of the Pareto II type.

The most important case for operational risk is $\xi > 0$. The reason for this is that, just like operational loss distributions, Pareto distributions exhibit fat tails. In practice, to estimate the distribution in the region of extreme losses, fit the existing data from the tail of the severity distribution to a Pareto distribution. (See Exhibits 3–5.)

Combining Operational Losses from Different Sources

In operational risk it is often required to combine losses arising from different sources. For example, it may be necessary to combine people risk and relationship risk for a business line. Combining operational losses requires determining their joint (multivariate) distribution. If the risks are independent, then their joint distribution can be obtained through convolution.

Suppose that X and Y represent operational losses in FX and precious metals, respectively, and that you are interested in calculating the distribution of total opera-

EXHIBIT 3

Generalized Pareto Distributions for Different Values of the Shape Parameter: Cumulative Distribution Function

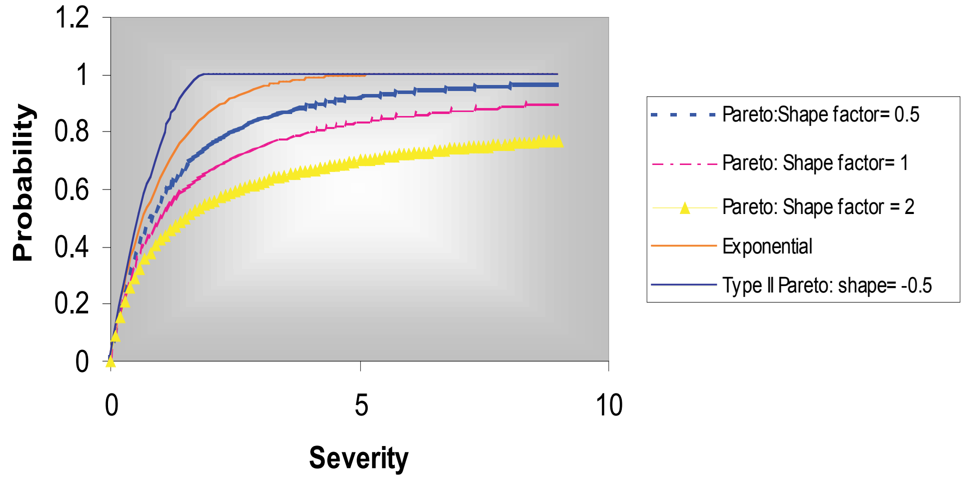

Cumulative distributions for the GPD class. The shape parameters for the Pareto distributions are 0.5, 1, and 2. The shape parameter for the Pareto II is -0.5. Note that for the Pareto distribution the larger the value of the shape parameter, the fatter the tail.

tional losses for both units. That is, the distribution of S, where S = X + Y. If X and Y are independent, then this distribution is given by

$$F_S(s) = \int F_X(s - y) f_Y(y) dy$$

where the Fs denote distribution functions and $f_Y(y)$ is the probability density function for the random variable Y.

If the risks are dependent and their joint distribution is a multivariate Gaussian, then their correlation matrix ρ, means, and standard deviations completely determine their combined loss distribution. However,

EXHIBIT 4

Generalized Pareto Probability Density Functions

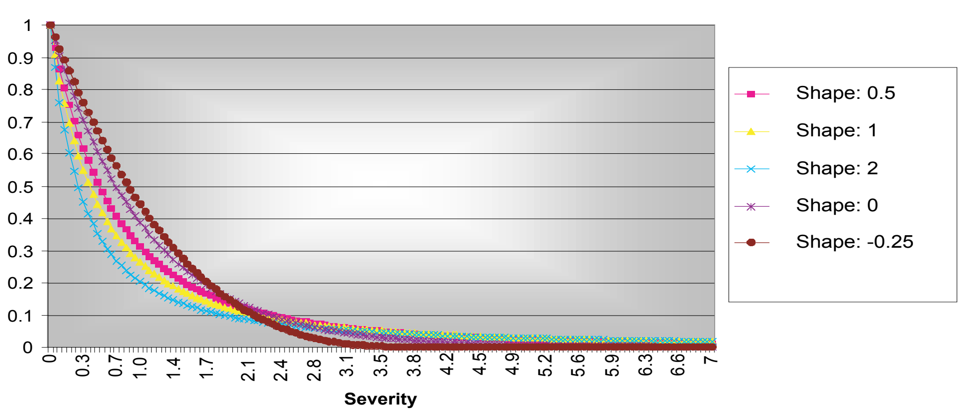

Generalized Pareto probability density functions for different values of the shape parameter. In all cases, the position (μ) and scale (σ) parameters are 0 and 1, respectively. The shape parameter varies from -0.25 to 2. Note that the larger the shape parameter, the fatter the tail.

EXHIBIT 5

Generalized Pareto Probability Density Functions: Detail

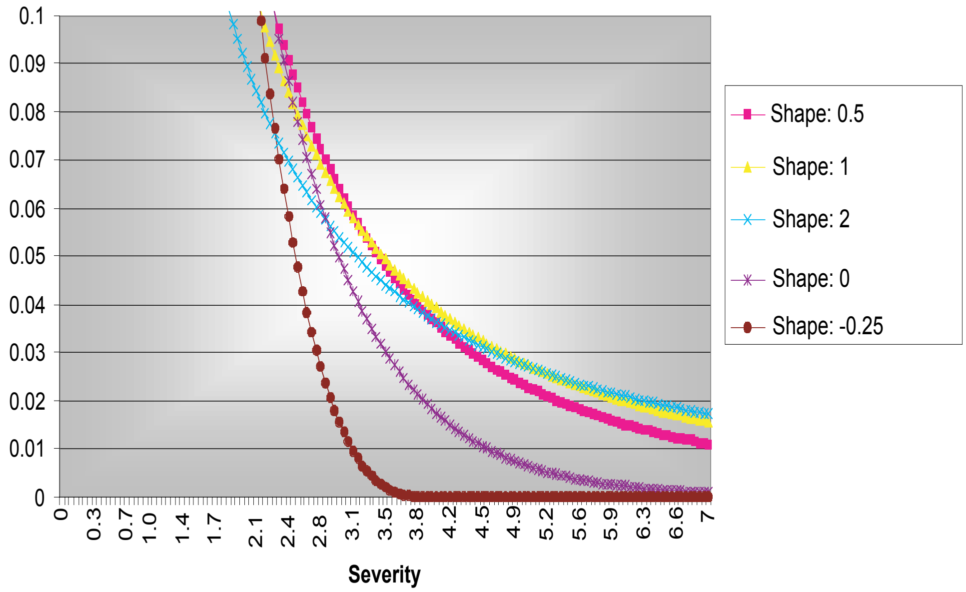

This chart illustrates that the larger the shape parameter, the fatter the tail. The parameters are the same as in Exhibit 4.

loss distributions in operational risk are not Gaussian (they are kurtotic and skewed). Knowledge of the individual loss distributions from specific business units, data sources, etc. (i.e., the marginal distributions), $F_1(x_1), \ldots, F_n(x_n)$, corresponding to the n risks and their correlation matrix are not enough to generate the corresponding multivariate loss distribution. For operational loss distributions it is necessary to specify the dependency structure in addition to the marginal distributions and their correlation matrix.

Copulas are a tool for combining correlated risks.[3] A copula C is the joint distribution function of n uniform ~(0, 1) random variables

$$C(u_1, \ldots, u_n) = \Pr\{U_1 \leq u_1, \ldots, U_n \leq u_n\}$$

For any set of marginal distributions the expression

$$F(x_1, \ldots, x_n) = C[F_1(x_1), \ldots, F_n(x_n)]$$

defines a joint distribution function with marginals $F_1(x_1), \ldots, F_n(x_n)$. The joint distribution F is specified by its dependence structure (the copula) and the marginal loss distributions. This separation provides great flexibility in modeling. For example, it is possible to combine losses from different business lines utilizing their marginal loss distributions and correlation matrix and assuming that the dependence structure is given by a Gaussian copula.[4]

Copulas can be used in conjunction with Monte Carlo simulations to aggregate correlated losses. The selected copula should reflect the empirical properties and properties of the risk measurement problem.

CONCLUSION

Although deciding to devote resources to measure operational risk is the first step toward managing and mitigating this risk at your institution, there are numerous practical issues to overcome. Although initiatives, such as multinational data consortiums, and the continued development of third-party models provide encouragement, operational risk measurement is far behind the measurement of other risk classes.

However, devoting resources to measuring this potentially catastrophic risk class is certainly called for. Improved firm-wide management of operational risk begins with understanding the sources of risk and measuring their potential impact on the capital structure and shareholder value of your firm. With improved measures, markets for financing and transferring unwanted operational risk will be more available, as will management's ability to prioritize the allocation of resources to improve your firm's risk-adjusted returns.

ENDNOTES

[1]Throughout this article, the terms "operational risk" and "event risk" will be used interchangeably. For a detailed discussion of the types of events included in the definition operational risk, see NetRisk [2000]. Available publicly through www.netrisk.com, www.morexchange.com, and www.watersinfo.com.

[2]For a detailed discussion, see Embrechts, Kluppelberg, and Milkosch [1997].

[3]For example, see Embrechts, McNeil, and Straumann [1999].

[4]A Gaussian copula is the copula that when applied to Gaussian marginals with the given correlation matrix leads to the corresponding multivariate Gaussian distribution.

REFERENCES

Embrechts, P., C. Kluppelberg, and T. Mikosch. *Modelling Extremal Events for Insurance and Finance*. Berlin: Springer, 1997.

Embrechts, P., A. McNeil, and D. Strauman. "Correlation: Pitfalls and Alternatives." *Risk*, May 1999.

Joe, H. *Multivariate Models and Dependence Concepts*. London: Chapman & Hall, 1997.

"Methodology for the Classification of Operational Losses." NetRisk, Inc., 2000.