This article originally appeared in the March 2000 issue of RISK

QUANTIFYING OPERATIONAL RISK*

Charles Smithson

Financial institutions recognize the importance of quantifying (and managing) operational risk. While losses like the $1.6 billion loss suffered by Barings in 1995 capture most of the attention, operational losses are widespread. PricewaterhouseCoopers compiled press reports indicating that financial institutions lost in excess of $7 billion in 1998 due to operational problems and that the largest financial institutions lose as much as $100 million annually.[1] Operational Risk, Inc. indicates that, since 1980, financial institutions have lost more than $200 billion due to operational risk.

Supervisors also recognize the importance of quantifying operational risk. In its June 1999 Consultative Document, the Basel Committee on Banking Supervision expressed its belief that operational risks (including reputational and legal risks) "are sufficiently important for banks to devote the necessary resources to quantify the level of such risks and to incorporate them into their assessment of their overall capital adequacy." Indeed, the Committee indicated its intention to require regulatory capital for operational risk.

The problem is how to accomplish the quantification. The Basel Committee mentioned that the options ranged from a simple benchmark to modeling techniques. The objective of this column is to provide an overview of the techniques being employed.

The Basel Committee conjectured that a simple benchmark measure of operational risk could be based on an aggregate measure of the size of the institution,

$$(\text{Operational Risk})_i = \Psi(\text{Size})_i$$

where Ψ is a parameter relating operational risk to institution size. As possible measures of the size of the institution, the Committee suggested gross revenue, fee income, operating costs, managed assets or total assets adjusted for off-balance-sheet exposures. While such a relation has some intuitive appeal and is easy to calculate, it does not capture the relation of operational risk to the nature of the institution's business. Indeed, a recent empirical examination by Shih, Samad-Khan, and Medapa (2000) suggests that little of the variability in the <u>size</u> of operational losses is explained by the size of a firm – revenue, assets, or number of employees.[2] Moreover, such an approach runs the risk of setting up perverse incentives – a financial institution that dramatically improves the management and control of its operational risk could actually be penalized by being required to hold more capital, if the improvements lead to an increase in the volume of the institution's business.

---

* Thanks is due to Dan Mudge and José V. Hernández (NetRisk), Michael Haubenstock (PricewaterhouseCoopers), and Jack King (Algorithmics) for their help with this column.

[1] *American Banker* November 18, 1999

[2] Note that this study did not deal with the frequency of operational losses.

Most of the published descriptions of operational risk modeling subdivide the models into two ***groups***: "Top Down" models estimate operational risk for the entire institution. "Bottom Up" models estimate operational risk at the individual business unit or process level. Moreover, the models could appropriately be subdivided on another dimension -- within the two groups, three ***approaches*** to modeling operational risk can be identified:

1) The approach I shall refer to as the ***Process Approach*** focuses on the individual processes that make up the financial institution's operational activities. (Because of this focus, all of the process approaches are bottom up approaches.) In the same way that an industrial engineer examines a manufacturing process, individual operational processes in the financial institution are mapped (decomposed) to highlight the components that make up the process. For example, Exhibit 1 provides a process map for a transaction settlement.

Exhibit 1

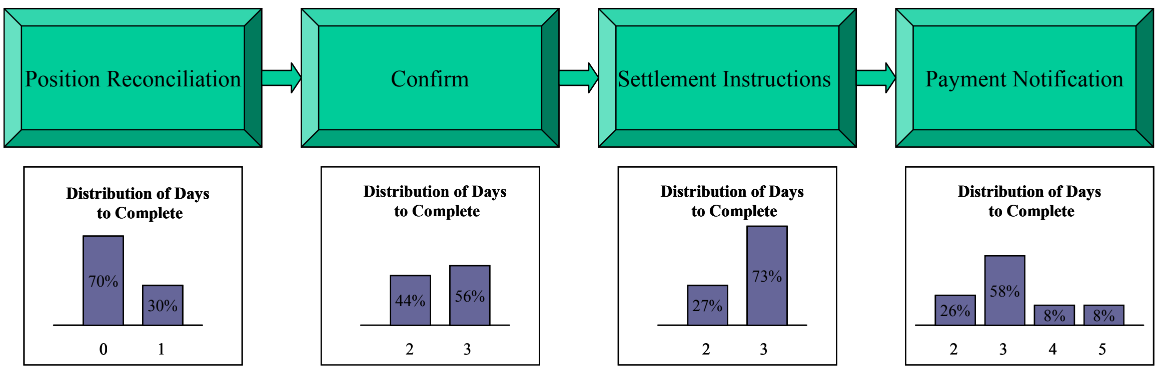

Source: Algorithmics

Each of the components of the process is examined to identify the operational risk associated with the component – e.g., in the case of Exhibit 1, the number of days necessary to complete the process. By aggregating the operational risk inherent in the individual components, the analyst can obtain a measure of operational risk in the process.

2) In the approach I will refer to as the ***Factor Approach***, the analyst is attempting to identify the significant ***determinants*** of operational risk – either at the institution level or at the level of an individual business or individual process. The objective is to obtain an equation that relates the level of operational risk for institution *i* (or business *i* or process *i*) to a set of factors:

$$(\text{Operational Risk})_i = \alpha + \beta(\text{Factor 1}) + \gamma\,(\text{Factor 2}) + \ldots$$

If she/he is able to identify the appropriate factors and obtain measures of the parameters (α, β, γ, …), the analyst can estimate the level of operational risk that will exist in future periods.

3) The focus of the ***Actuarial Approach*** is on the identification of the loss distribution associated with operational risk – either at the level of the institution or at the level of

a business or process. (This contrasts to the first two approaches that both focus on identifying the sources of operational risk.) Exhibit 2 illustrates a stylized loss distribution, which combines both the frequency of the loss events and their severity.

Exhibit 2

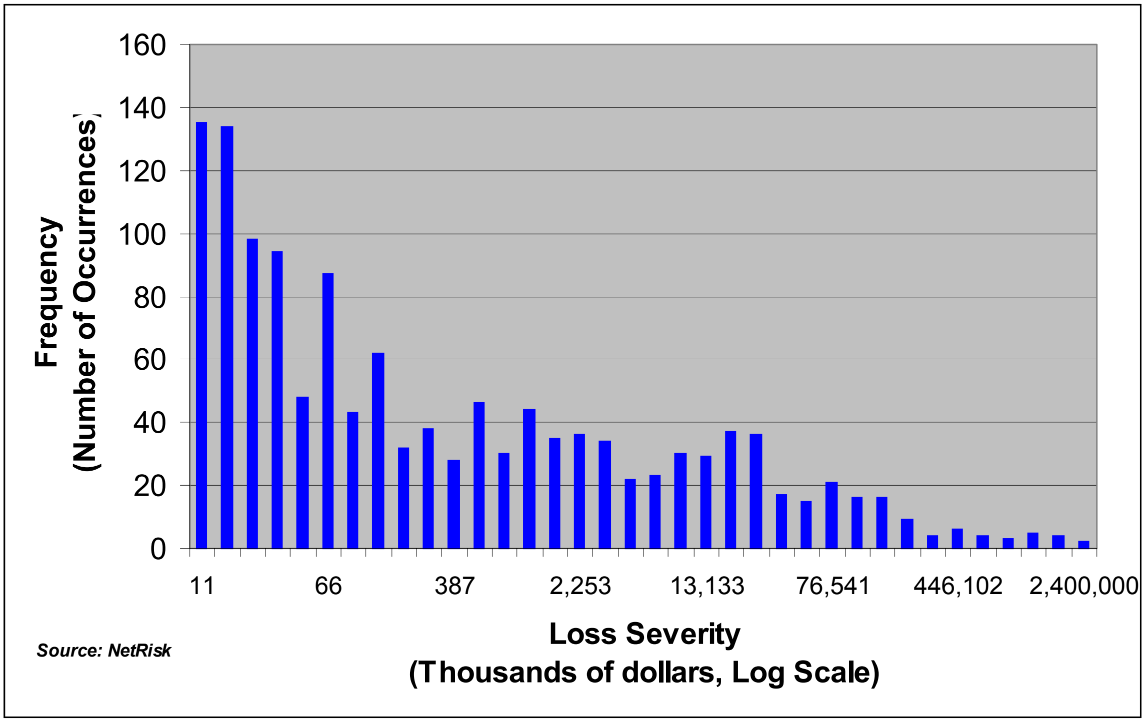

Exhibit 3 categorizes the various operational risk models that have been publicly discussed.[3]

EXHIBIT 3

	----------------------------------	*Modeling Approaches*	----------------------------------
	Process	Factor	Actuarial
Top Down Models		- Risk Indicators - "CAPM-Like" Models	- Empirical loss distributions
Bottom Up Models	- Causal Networks - Statistical quality control and reliability analysis - Connectivity	- Predictive Models	- Empirical loss distributions - Explicit distributions parameterized using historical data - Extreme value theory

[3] Much of the discussion that follows is adapted from Ceske/Hernández (1999) and O'Brien (1999).

Process Approaches

Causal Networks – The process illustrated in Exhibit 1 is a causal network. The analysis begins with a graphical map of the components in the process, with linkages between the components visible. Historical data are used to produce statistics for the behavior of the components and the process in the past. (This permits the analyst to identify problem areas.) Scenarios or simulations can be employed to predict how process will behave in the future.

Statistical quality control and reliability analysis – Similar to causal networks, this technique is widely used in manufacturing processes.

Connectivity – The focus is on the connections between the components in a process. The analyst creates a "connectivity matrix" that is used to estimate potential losses for that process. If the processes are aggregated, a "failure" in one component will propagate across the process and through the institution

Factor Approaches

Risk Indicators – The analyst identifies the significant factors using regression techniques. (In addition to volume, factors can include audit ratings, employee turnover, employee training, age of the system, and investment in new technology.) The analyst can use the resulting equation to estimate expected losses.

"CAPM-Like" Models – In contrast to focusing on the frequency and/or severity of operational losses, this approach would relate the volatility in share returns (and earnings and other components of the institution's valuation] to operational risk factors.

Predictive Models – Extending the risk indicator techniques described above, the analyst uses discriminant analysis and similar techniques to identify factors that "lead" operational losses. The objective is to estimate the probability and severity of future losses. (Such techniques have been used successfully for predicting the probability of credit losses in credit card businesses.)

Actuarial Approaches

Empirical loss distributions – The objective of the actuarial approach is to provide an estimate of the loss distribution associated with operational risk. The simplest way to accomplish that task is to collect data on losses and arrange the data in a histogram like the one illustrated in Exhibit 2. Since individual financial institutions have data on "high-frequency, low-severity" losses (e.g., interest lost as a result of delayed settlements) but do not have many observations of their own on the "low frequency, high-severity" losses (e.g., losses due to "rogue" traders), the histogram will likely be constructed using both internal data and [properly scaled] external data. In this process, individual institutions could benefit by pooling their individual observations to increase the size of the data set. Several industry initiatives are underway to facilitate such a data pooling exercise – the Multinational Operational Risk Exchange (MORE) project of the Global Association of

Risk Professionals (GARP) managed by NetRisk, a project at PricewaterhouseCoopers, and a BBA project.

Explicit distributions parameterized using historical data – Even after making efforts to pool data, an empirical histogram will likely suffer from limited data points, especially in the tail of the distribution. A way of smoothing the histogram is to specify an explicit distributional form. A number of analysts have concluded that, rather than specifying a distributional form for the loss distribution itself, better results are obtained by specifying a distribution for the frequency of occurrence of losses and a different distribution for the severity of the losses.[4] In the case of frequency, it appears that most analysts are using the Poisson distribution. In the case of severity; analysts are using a range of distributions, including a lognormal distribution and the Weibull distribution. Once the two distributions have been parameterized using the historical data, the analyst can combine the two distributions (using a process called "convolution") to obtain a loss distribution.

Extreme value theory – Because large operational losses are rare, an empirical loss distribution will be sparsely populated (i.e., will have few data points) in the high severity region. Extreme value theory – an area of statistics concerned with modeling the limiting behavior of sample extremes -- can help the analyst to obtain a smooth distribution for this important segment of the loss distribution. Specifically, extreme value theory indicates that, for a large class of distributions, losses in excess of a high enough threshold all follow the same distribution (a generalized Pareto distribution).

References

Basel Committee on Banking Supervision, 1999. Consultative Paper: "A New Capital Adequacy Framework." Basel: Bank for International Settlements, June.

Ceske, Robert and José V. Hernández, 1999. "Measuring Operational Risk – Where Theory Meets Practice," NetRisk, October. An abridged version of this paper appeared in *Operational Risk*, a Risk Special Report, November 1999.

O'Brien, Niall, 1999. "The Case for Quantification," *Operational Risk*, a Risk Special Report, July.

Shih, Jimmy, Ali Samad-Khan, and Pat Medapa, 2000. "Is the Size of an Operational Loss Related to Firm Size?," *Operational Risk*, January.

[4] The proponents of this approach point to two advantages: (1) It provides more flexibility and more control. (2) It increases the number of useable data points.