Operational Risk

Is the Size of an Operational Loss Related to Firm Size?

By Jimmy Shih, Ali Samad-Khan and Pat Medapa

The European Commission recently proposed that capital charges for operational risk might be based on the size and income of an institution (*Risk Management Operations*, 13 December).

While it seems intuitive that operational risk is to some degree a function of firm size, the nature of this relationship is not straightforward. We therefore conducted a study to empirically test one aspect of this relationship: whether the magnitude of a loss experienced by a firm is correlated with the size of the firm.[1]

The study revealed that

- size only accounts for a very small portion (about 5%) of the variability in loss severity;
- the size of a firm is related to the magnitude of its loss, but the relationship between size of loss and size of firm is not linear; and
- there is clear evidence of a diminishing relationship between the size of a firm and loss magnitude.

Identifying the relevant scale variable

We began by calculating level and log correlations between operational losses and three variables associated with the size of a firm.[2] These "scale variables" were revenue, assets and number of employees.

We found that all three variables were correlated with loss size, with the revenue variable showing the strongest relationship.

We also noted that the logarithm of the scale variables showed a stronger relationship to losses than did the raw scale variables.

This suggested that the relationship between size of firm and loss magnitude was not linear. That is, a firm that is twice as large as another does not, on average, suffer a loss that is twice the size of the loss experienced by the other firm.

We subsequently examined the correlation between the logarithm of the three scale variables and found, not surprisingly, that all three variables were highly correlated.

Because multiple regression in this case would have resulted in little improvement in significance, we chose to specify our regression equation in terms of a single independent variable.

Hypothesizing the scale relationship

In view of the above, we hypothesized the following relationship between size of firm and loss magnitude:

$$L = R^{\alpha} \times F(\theta) \quad (1.1)$$

where:

L is the actual loss amount;

R the revenue size of the firm in which the loss took place;

α is the scaling factor measuring the degree of return to scale; and

θ is the vector of all the risk factors not explained by revenue R so that $F(\theta)$ is the multiplicative residual term not explained by any fluctuations in size.

In the above equation $\alpha = 1$ would represent a linear relationship between R and L, $\alpha < 1$ would represent a diminishing relationship between R and L, while $\alpha > 1$ would represent an increasing relationship.

Ordinary least squares regression

Taking the logarithm of both sides of equation (1.1) yields the following equation:

[1] The purpose of this study was to determine the relationship between the size of a firm and single-event losses. The aggregate loss that a firm suffers in one year, which is more closely related to its total operational risk, is also dependent on the number of losses that the firm experiences in one year. The relationship between the size of a firm and annual loss frequency was not studied because sufficient data do not exist to support rigorous empirical analysis.

[2] The data used for this study were obtained from the PricewaterhouseCoopers OpVaR^{SM} database, a database of publicly reported operational losses. At the time of this study, the OpVaR database contained over 4700 losses in excess of $1 million.

$$l = \alpha x r + \beta + \varepsilon \quad (1.2)$$

where $l = \ln L$,

$r = \ln R$, $\beta = E[\ln F(\theta)]$ and

$\varepsilon = \ln(F(\theta)) - \beta$

An ordinary least square (*OLS*) regression was performed on all the

A. OLS Regression Results

OLS regression results	*Coefficients*	*Standard error*	*t Statistic*	Regression statistics	
Intercept (β)	1.275950699	0.121395484	10.51069	R Square	0.054102027
Ln*R* (α)	0.151524065	0.014702495	10.30601	Adjusted R Square	0.053592658

data, based upon equation 1.2. The results are shown in Table A. Our initial regression plot showed that the residual data scatter was very symmetric about the horizontal axis, strongly indicating a linear relationship between l and r (the log of loss and the log of revenues).

The t-statistics for the coefficient also indicated a very significant linear relationship between these two variables.

B. WLS Regression Results

WLS regression results	*Coefficients*	*Standard error*	*t Statistic*	Regression statistics	
Intercept	0.232445601	0.0093497	24.86129	R Square	0.090513977
X variable 1	0.695024654	0.051166479	13.58359	Adjusted R Square	0.090023424

We noted that both the R^2 and adjusted R^2 were very small - only slightly above 5%. An R^2 of around 5% suggests that around 95% of the loss variability in the data is attributable to factors other than the independent variable.

Generalized least squares regression

We noted that the residual plot of the results exhibited a funnel shape, indicating a positive linear relationship between loss variability and firm size – a classic case of heteroskedasticity.[3] This may be because large firms suffer large losses, as well as small losses.

The presence of heteroskedasticity in the error term meant that any OLS estimator for α would be inefficient, ie a reduction in variance of the estimator is no longer guaranteed asymptotically. We therefore ran a weighted least square (*WLS*) regression, which is a special form of a generalized least square regression.

This was accomplished by dividing both sides of equation 1.2 by r and then regressing the following relationship:

$$y = \alpha + \beta x x + \phi \quad (1.3)$$

where $y = l/r$, $x = 1/r$ are the new dependent and independent variables, respectively, and ϕ is the new error term. (Note that the regression coefficient for l in the *OLS* equation reappears as the intercept in the *WLS* equation.)

The GLS regression results in Table B show that both the R^2 and adjusted R^2 improved. The *WLS* estimate of the scaling coefficient was determined to be $\alpha = 0.2324$ instead of the original 0.1515.

Remarks

If the size of a firm is so weakly related to its size of loss, what are the causes of loss variability? We suspect that that the vast majority of the variability is caused by factors such as inherent differences in risk (based on the type of business conducted), the competence of management, and the quality of the internal control environment.

Note that the results of this study could be applied to the statistical modelling of operational value at risk.

The correlation coefficient could be used to scale-adjust external loss data to the size of the firm that is being modelled.

The authors are consultants in the Financial Risk Management practice of PricewaterhouseCoopers in New York.

Operational Risk Capital – Why Actuarial Approaches Work

By Michael Haubenstock

"Operational, legal and reputation risks are sufficiently important for banks to ... quantify the level of such risks and to incorporate them into their assessment of their overall capital adequacy," says the June 1999 consultative paper from the Basel Committee.[1]

But how? What type of model can meet the expectations of the regulators, recognize the unique attributes of each institution and reinforce the right behaviour in bank employees?

The Basel Committee does not yet have an answer, though it may be feeling its way towards one (pages 10 – 12).

The June paper supposes that there could be a measure based on revenues, costs or other indicators, but, in the end, it looks to the industry for help in solving the problem.

On one point, everyone is in agreement: the measurement methodologies used by banks today are inadequate for assessing regulatory capital.

The trend in the industry is clearly towards actuarial internal models.[2]

Half-way measures

As of today, banks have used a whole variety of approaches to measure their operational risk (*see Box 1, page 16*).

These include peer comparisons, earnings-at-risk models, approaches based on the capital asset pricing model, scenario analysis and various types of top-down allocations.

Often, analysts then modify the results by factoring in internal assessments of one kind or another.

These might include audit results; operational risk self-assessment programmes; adding data such as fixed assets or long-term contingent liabilities; or including risk indicators or transaction and volume flows. The combinations are endless.

Almost all these formulas generate a capital number that can be used for internal capital allocation and for risk-adjusted return on capital calculations.

But the number does not support operational risk management as a process. Since the results are not specifically related to each type of risk, they cannot be analysed to determine the cause of problems.

Likewise, they are of no use to managers trying to quantify specific problems or for those trying to establish priorities for corrective action.

Moreover, the underlying finance theory is questionable. The theory of risk capital rests on the principle that capital is present to absorb volatility when actual losses are in excess of the expected annual loss.

The approaches presented in Box 1 attempt to quantify the unexpected loss without having bothered to establish the expected loss. The end result is difficult to defend – and is unlikely to be accepted by regulators.

Statistical/actuarial approaches are better

From what we've seen in the market thus far, there is really only one way to meet the combined needs of internal risk managers, managers of capital and the growing expectations of regulators.

It is to employ an empirically based "bottom-up" statistical/actuarial model.

This technique is characterized by the collection of actual operational loss event data, both from within an institution and from external sources.

Analysts can then use this data to build statistical distributions of both the loss severity associated with an event, and the frequency of its occurrence.

Next, a Monte Carlo simulation can be used to calculate a value-at-risk number for each type of operational risk and for each business line.

Is the approach practical?

This is the approach adopted by PricewaterhouseCoopers. Working with various financial institutions, we have created distributions for each of their businesses.

It has not proved straightforward. Perhaps the largest issue is the limited amount of data available. Even with a reasonable history of internal loss data, there is usually not enough to build good frequency and severity distributions for each risk and business line. The quality of internal loss data necessitates the use of external data as well.

But combining the internal data and external data creates great statistical challenges. The external data comes from institutions of different sizes, with different levels of controls, and suffers from biases related to the level of data capture.

So simply combining the external and internal data is theoretically invalid. An objective aggregation process is necessary if the results are to be meaningful.

There is another problem, too.

Monte Carlo simulation can be used to generate sets of "hard" results – statistically derived from actual loss data. But while historical data takes time to accumulate, organizations change daily.

That's why it's important to include, in any approach, elements of qualitative assessment.

Qualitative assessments might include the level of controls (as compared to the historical data), product complexity, or the legal environment in any one country, for example. Or factors such as new products (e.g. an e-business initiative), rising volumes due to merger, or product growth.

Self-assessment and other management processes, integrated with the quantification model, allow the bank to alter its cap-

1. "A New Capital Adequacy Framework", Basel Committee on Banking Supervision, June 1999.

2. This is supported by the BBA/ISDA/RMA industry survey, Operational Risk – The Next Frontier, autumn 1999.

Box 1. Which other approaches are in use today?

Peer comparison Here a company looks to other financial institutions or to other industries, for risk characteristics that are similar to those in each of its business lines. Then it attempts to estimate its own capital levels based on those of other companies. Firms using this methodology find it a struggle not only to identify good comparisons, but also to satisfy themselves that the two units are similar enough to justify using the data.

Revenue/expense volatility The theory here is that risk causes volatility, and that after eliminating the effects of market and credit risk, the remaining volatility can be attributed to operational risk - defined in its broadest sense to include business and strategic risks. There are two major problems with this approach. One is obtaining clean data. The second is that the approach assumes that every operational risk exposure has been experienced by the company during the period studied.

CAPM The capital asset pricing model can be used to estimate the total amount of operational risk capital expected by equity investors. This total figure is then allocated through some means. While interesting from a capital management perspective, how can external investors really understand the level of risk taken internal to a company, when even its own managers struggle to understand it?

Scenario analysis Applied in different forms, this approach attempts to ask and answer the question, "What if this [disastrous event] happened?" The analyst tries to quantify the effect of various operational risk events and, consequently, to derive a range of distributions. A variant of the method involves analysing each risk in each business and then assigning a high, medium or low frequency and severity to the risk; the analyst then calculates capital requirements on the basis of the assumed frequencies and costs per event. This is a risk-based approach, but it remains subjective and is labour intensive.

Top-down allocations Often companies already have a rule of thumb figure for firm-wide operational risk capital, but face the problem of allocating this among various business lines. Most commonly, the allocation is based on non-interest expenses. The logic here is that the primary expenses are people and technology, and the more a business has of these, the higher its operational risks. This is not a sound assumption, particularly from a capital perspective. High expense businesses (e.g. retail banking) may have the lowest volatility when it comes to operational risk, while lower cost businesses, e.g. trading, may have much greater exposure to operational risk. Worse still, it fails to address the critical question - what is the right level of total operational risk capital?

ital allocation as the organization changes shape.

In this way, the amounts of capital charged to various business lines create the right behavioural re-enforcement – low levels of loss and good assessments of risk result in lower capital charges.

Statistical models are not created equal

There are various approaches to building statistical models. Many rely on simplistic assumptions.

We suggest weeding out the weaker approaches, using the following criteria:

• Models must be based primarily on internal loss data in order to capture the unique attributes and experience of each institution.

• External data needs to be incorporated in a statistically valid manner.

• The type of distribution and its parameters should be fitted to the data, as opposed to pre-selecting a certain type of distribution.

• The approach should be bottom-up – business line value-at-risk should be derived from data from that business line, not through some allocation process.

• Insurance coverage should be explicitly considered.

• The model must allow validation through back-testing.

Missing piece of jigsaw

Why are we so sure this is the right approach?

Statistical/actuarial models are the only approaches that are:

- based on actual results;

- use mathematical models consistent with market and credit risk;

- can be calibrated and back-tested for each institution;

- evolve as the institution changes; and

- can take account of any insurance programmes in place.

In addition, the ability to calculate exposure to each specific risk within each business line supports the operational risk management process by greatly improving the analysis of root causes.

These approaches are feasible today using internal data, supplemented by external events gleaned from press articles.

They will become even more accurate when the next piece of the jigsaw becomes available: the financial industry must start to share the data it collects on internal loss events.

Only rich collaborative databases will allow detailed modelling, and help make the statistical results defensible in the eyes of business line leaders and regulators.

Michael Haubenstock is a Partner of PricewaterhouseCoopers