

Rules and Models¹

Carol Alexander² investigates the internal measurement approach for operational risk capital

There is a view that the new Basel Accord is being defined by a committee of technocrats with no real understanding of the issues, nor sympathy with banks for the difficulties that their current proposals would impose on their operations. Whilst I understand those that take this view, and value their efforts to temporize regulatory changes, I do not hold to such opinions myself.

With the publication of working paper CP2.5 on the Regulatory Treatment of Operational Risk, the Basel Committee has proposed the rules for the playing fields.³ These rules define a structure, a set of boxes, into which banks should re-organize the measurement and management of Operational Risks. The Committee has not attempted to tell us what is in these boxes, but that does not mean that they are empty.

In this article I examine the foundations and some simple extensions of the internal measurement approach (IMA), one of the advanced measurement approaches (AMA) defined in CP2.5. I show that the IMA is particularly attractive for a number of reasons. First it is surprisingly flexible, and can accommodate a number of different modelling assumptions; second, data and modelling requirements are kept to a minimum; and third, it allows a simple formula for the inclusion of insurance cover. One cannot say that it will be the best approach for all risk types – there will be more flexibility to model different risk types using the loss distribution approach (LDA) – but in my opinion, it will be possible for many banks to implement the IMA by January 2005, whereas only a few banks will have the data and technology that are necessary to implement a more advanced LDA .

The Binomial Model

For each business line/risk type

$$\text{IMA operational risk capital charge} = \text{gamma} \times \text{expected loss}$$

The total operational risk capital charge is the sum of all charges over business lines and risk types; this assumes the worst possible case, of perfect correlation between individual risks.⁴ Since the capital charge is for covering unexpected losses, the IMA assumes that unexpected loss is a multiple of expected loss. The rules proposed in CP2.5, which allow banks to calibrate their own gammas, do not require that gamma should be a constant. In fact the method by which expected loss is calculated in CP2.5 implies that it is based on the binomial model, and the logical consequence of this is that gamma will be inversely proportional to the square root of the total number of loss events. This is certainly not a constant.

There is wide acceptance in the industry that the binomial distribution is an adequate model for the frequency of loss events, at least to a first order approximation. However the IMA can

¹ 'Rules and models destroy genius and art' William Hazlitt (1839).

² Carol Alexander is Professor of Risk Management and Director of Research at the ISMA Centre at the University of Reading.

³ Basel Committee, 28th September 2001. Available from www.bis.org

⁴ In the IMA model charges are proportional to standard deviations. But the sum of standard deviations is only equal to the standard deviation of the sum when the correlation is unity.

readily be generalized to other distributions should a bank wish to do so. Let us first suppose that the IMA is based on the binomial model, as in Alexander and Pezier (2001)⁵. Then the expected total loss is $Np\mu_L$ and

$$\text{gamma} \approx k \sqrt{\{1 + (\sigma_L/\mu_L)^2\}} / \sqrt{[Np]}$$

where μ_L is the expected loss given the event, σ_L is the standard deviation of this loss and the multiplier k will be calibrated for different types of risks, as discussed below. In the binomial framework the total number of loss events is denoted Np , where N is the number of events that are susceptible to operational losses during the time horizon of the model (which is currently thought most likely to be one year) and p is the probability of a loss event.

An important point that was not stressed in Alexander and Pezier (2001) is that the introduction of separate notation for N and p is only necessary to help us understand the foundations of the model. It will not be necessary to estimate N and p separately (unless the bank wishes to use one of the more advanced parameter estimation methods mentioned below) because they always appear as a product in the formulae for gamma and the capital charge. However banks will need to estimate Np – that is, the average number of unauthorized transactions, or fraudulent deals, or claims relating to employment practices, or failed trades, or systems breakdowns and so forth. Banks will also need to estimate the average loss given that the event incurs a loss, but if they do not have an estimate of the loss standard deviation the simple formulae where $\text{gamma} \approx k/\sqrt{[Np]}$ and the capital charge is $k\mu_L\sqrt{[Np]}$ can be applied.

Alternative Models

The basic framework of the IMA assumes that

$$\text{capital charge} = k \times \text{standard deviation} = \text{gamma} \times \text{expected loss}$$

In CP2.5 the expected loss is defined as an exposure indicator (N) \times expected loss given event (μ_L) \times probability of a loss event (p) so it has been natural to think in terms of the binomial model. However, another loss frequency distribution that can be used with the IMA is the Poisson, with parameter λ which corresponds to the expected number of loss events in the time horizon.⁶ The expected total loss is $\lambda\mu_L$ and the variance is $\lambda\mu_L^2$, so if the Poisson model is used to fit loss frequency instead of the binomial, the gamma will be equal to $k \sqrt{\{1 + (\sigma_L/\mu_L)^2\}} / \sqrt{\lambda}$ and the capital charge will be given by the formula

$$k \mu_L \sqrt{[(1 + (\sigma_L/\mu_L)^2) \lambda]}$$

A single parameter family probably offers insufficient scope to fit loss frequency distributions for all the different risk types and business lines encompassed by the bank's activities. In that case the bank may consider using a more flexible distribution such as the gamma distribution, which has two parameters α and β and the density function

$$f(x) = x^{\alpha-1} \exp(-x/\beta) / \beta^\alpha \Gamma(\alpha) \quad x > 0.$$

⁵ *Binomial Gammas* by Mr. and Mrs. Pezier of Purley, Surrey in *Operational Risk* 2 Issue 3, pp8-10.

⁶ When the probability of loss is small the Poisson is a close approximation to the binomial distribution.

The mean and variance of the gamma distribution are $\beta\alpha$ and $\beta^2\alpha$ respectively. Therefore if the loss frequency is gamma distributed, $\text{gamma} = k \sqrt{1 + (\sigma_L/\mu_L)^2} / \sqrt{\alpha}$ and the capital charge will be given by the formula

$$k \mu_L \sqrt{[(1 + (\sigma_L/\mu_L)^2) \beta^2 \alpha]}$$

The reader may easily construct other examples for themselves. The point to note is that the conceptual foundation of the results holds true whatever the distribution assumed for the frequency of loss events: the capital charge will always be directly proportional to the square root of the average number of loss events; therefore it will increase like the *square root* of the size of the banks operations, and like the *square root* of the time over which losses are measured. There is no 'linearity' in the IMA, and I think this is already recognized in CP2.5.

The Multiplier

Leaving aside, for the moment, the important issues of data collection and parameter estimation, let us examine how the multiple k should be calibrated. Since capital charges are to cover unexpected losses, k is the ratio of the unexpected loss to the standard deviation. For example, in the standard normal distribution and for the 99.9% confidence level that is recommended in CP2.5 for the LDA, $k = 3.10$, as can be found from standard normal tables. For the binomial distribution with $N = 20$ and $p = 0.05$ (so the expected number of loss events is 1) the standard deviation is 0.9747 and the 99.9% percentile is 5.6818, so $k = (5.6818 - 1)/0.9747 = 4.80$. For the Poisson distribution with expected number of loss events equal to 1, the standard deviation is 1 and the 99.9% percentile is 5.84, so $k = (5.84 - 1)/1 = 4.84$; but for higher frequency risks where the expected number of loss events is, say, 20, the Poisson distribution has standard deviation $\sqrt{20}$ and 99.9% percentile 35.714, so $k = (35.714 - 20)/\sqrt{20} = 3.51$.

In general, the value of the multiplier k depends more on the type of risk than the type of distribution that is assumed for loss frequency. High frequency risks, such as those associated with transactions processing, will have lower multipliers than low frequency risks, such as a fraud. This is because for high frequency risks the expected number of loss events is high, relative to their standard deviation, and the calculation of k should take expected loss into account, as we did above. That is, unexpected loss is defined to be the difference between the upper percentile loss and the expected loss. Normally, accountants should make special provisions in the balance sheet to cover expected losses, so they do not need to be taken into risk capital charges. Some banks do not take unexpected loss to be the difference between the upper percentile and the expected loss, and this will increase capital charges for low impact high frequency risks in particular.

Regulators might use their approval process to introduce a 'fudge factor' to the multiplier, as they have done with internal models for market risk. They may wish to set the multiplier by calibrating the operational risk capital obtained from this "bottom-up" IMA approach to that determined from their "top-down" approach. This is what they are attempting to do with the multipliers (alpha and beta) for the Basic Indicator method and the Standardized Approach to operational risk capital measurement.

Insurance

CP2.5 makes the rather contentious statement that banks will only be permitted to reduce capital charges by allowing for insurance cover if they use one of the advance measurement approaches, such as the IMA. The justification is that "this reflects the quality of risk

identification, measurement, monitoring and control inherent in the AMA and the difficulties in establishing a rigorous mechanism for recognizing insurance where banks use a simpler regulatory capital calculation technique". Banks that mitigate certain operational risks through insurance will thus be given the incentive to invest in the data and technology required by an AMA. They will also need to develop an appropriate formula for recognition of insurance that is risk-sensitive but not excessively complex.⁷

A simple formula for including insurance cover in the operational risk charge can be deduced using the binomial model. Insurance reduces the loss amount when the event occurs (an expected amount μ_R is recovered) but introduces a premium C to be paid even if the event does not occur. In the binomial model an expected amount $\mu_L - \mu_R$ is lost with probability p and C is lost with probability 1 , so the expected total loss is now $N[p(\mu_L - \mu_R) + C]$.

If we assume that the premium is fairly priced then the introduction of insurance will not affect the expected loss significantly, only the standard deviation of loss will be reduced. Thus the expected loss will be approximately $Np\mu_L$ as it was before insurance, and the premium will be set approximately equal to the expected pay-out, that is $C \approx p\mu_R$. However if p is small, the standard deviation is now approximately $(\mu_L - \mu_R) \sqrt{[Np]}$ and so if we denote the expected recovery rate μ_L/μ_R by r , $\gamma \approx k(1 - r) / \sqrt{[Np]}$ and the capital charge will be

$$k \mu_L (1 - r) \sqrt{[(1 + (\sigma_L/\mu_L)^2) Np]}$$

As before, this can be generalized to other types of distributions for loss frequency. The general result is the same in each case: If risks are insured and the expected recovery rate per claim is r , the capital charge should be reduced by a factor of $(1 - r)$. Of course, insurance is more complex than this because contracts will not cover individual events except perhaps for very large potential losses. However, as stated in CP2.5, a simple formula such as this will be necessary for banks that wish to allow for insurance cover when calculating capital charges.

Data

We have seen that low frequency high impact risks will have the largest effect on the bank's total capital charge. But for these risks, data are very difficult to obtain: by definition, internal data are likely to be sparse and unreliable. Even for high frequency risks where there are normally plenty of data available there will be problems following a merger, acquisition or sale of assets. When a bank's operations undergo a significant change in size, it may not be sufficient to simply re-scale the capital charge by the square root of the size of its current operations. The internal systems, processes and people are likely to have changed considerably and in this case the historic loss event data would no longer have the same relevance today.

In general, there is trade-off between relevance and availability of data. The bank will be left with no other option than to use 'soft' data that is available, but not necessarily as relevant as they would like. The bank may consider using subjective data in the form of opinions from industry experts, or data from an external consortium. In CP2.5 there is no mention of the use of expert opinions, but it is recognized that banks may supplement their internal loss data with the external industry loss data that are being collected in large data consortia such as www.moreexchange.org. The working paper states: "Member banks wishing to use these data in

⁷ However the total capital charge from the AMA, with or without allowance for insurance, will not be less than 75% of the capital charge under the Standardized Approach and it is not clear whether this reduction will be sufficient incentive for banks to develop AMA.

their advanced measurement models must establish proper procedures for the use of external data as a supplement to its internal loss data."

Bayesian Methods

How should 'soft' data, i.e. data from external sources, or expert opinions, or reflecting pre-merger internal practices, be used in conjunction with 'hard' data from internal, current operational processes? Classical methods, such as maximum likelihood estimation, treat all data as the same. But, for the significant risks that banks face, no two cases are the same. The risk assessment must take into account the specific elements in each case, and be subjective. Parameter estimation methods that allow subjective beliefs to play a role are called 'Bayesian' methods. Bayesian methods combine two different types of information: (a) 'prior beliefs' which may be based on the subjective opinions of industry experts or the less subjective data from an external consortium, and (b) 'sample likelihoods' which are based on 'hard' data. Often the internal 'hard' data are very sparse, but when combined with prior beliefs about model parameters the bank can obtain Bayesian parameter estimates.

Prior beliefs and sample likelihoods are expressed in terms of probability densities which are multiplied to give a posterior density on the model parameter. From this posterior density, a point parameter estimate – called the Bayesian estimate – may be obtained as the mean, mode or median of the posterior density, depending on the loss function of the decision maker. Often we assume that the decision maker has a quadratic loss function, in which case our point Bayesian estimate of the parameter will be the mean of the posterior density.

Loss Amounts

If both 'hard' internal data and 'soft' data are available on the distribution of losses, then Bayesian methods can be used to estimate μ_L and σ_L . To illustrate the method, suppose that in the 'hard' internal data the expected loss given a loss event is 5m\$ and the standard deviation of this loss is 2m\$; suppose that the 'soft' data, being obtained from an external consortium, shows an expected loss of 8m\$ and a loss standard deviation of 3m\$.

Assuming normality of loss amounts, the prior density that is based on external data is $N(8, 9)$ and the sample likelihood that is based on internal data is $N(5, 4)$. The posterior density will also be normal, with mean μ_L that is a weighted average of the prior expectation and the internal sample mean. The weights will be the reciprocals of the variances of the respective distributions. In fact the Bayesian estimate for the expected loss will be $\mu_L = [(5/4) + (8/9)] / [(1/4) + (1/9)] = 5.92\text{m\$}$ and the Bayesian estimate of the loss variance will be $[4 \times 9] / (4 + 9)$, so that the standard deviation of the posterior is $\sigma_L = 1.66\text{m\$}$.

The Bayesian estimate of the expected loss is nearer the expected loss in the internal data (5m\$) than that of the external data (8m\$) because the internal data has less variability than the external data. Given the heterogeneity of members in the data consortia, it is likely that the uncertainty in the internal estimates will be less than that of the external estimates, so the Bayesian estimate of the expected loss will, in general, be nearer the internal mean than the external mean.

Loss Probability

It is easier to obtain data on the average number of loss events rather than separate data on the number of events N and the probability of a loss, p . Consider the effect of separating these parameters. What does one mean by the number of 'events'? According to the line of business/loss distribution categorization outlined in CP2.5, the number of 'events' will be the

number of transactions, or employees, or new deals, or trades, or computers, or software systems and so forth that are expected over the time horizon of the model. This is going to be extremely difficult to quantify.

However, the value for N used to calculate the capital charge should really represent the *forecast* over the risk horizon (one year, in CP2.5) because the operational risk capital charge is supposed to be forward looking. Thus we should really use a target or projected value for N – assuming this can be defined by the management – and this target could be quite different from its historical value.

Using internal data alone to forecast a loss probability, p is going to be difficult for low frequency events because, by definition, very little internal data will be available. Bayesian estimates for p can use prior densities that are based on external data, or subjective opinions from industry experts, or 'soft' internal data. Bayesian estimation of a probability are often based on beta densities of the form

$$f(p) \propto p^a (1 - p)^b \quad 0 < p < 1.$$

For example, if in a sample of 100 events there are two loss events, the beta density that represents the probability of a loss event is $p^2(1 - p)^{98}$. Beta densities for probabilities are often used because the product of two beta densities is another beta density; so if the prior and likelihood are both beta densities, the posterior density will also be a beta density.

If we assume decision makers have quadratic loss functions, the Bayesian estimate of a proportion will be the mean of the posterior density. It is easy to show that a beta density $f(p) \propto p^a(1 - p)^b$ has mean $(a + 1)/(a + b + 2)$. This gives the Bayesian estimate that takes account of both data sources, with a and b being the parameters of the posterior density. For example, if internal data indicate that 2 out of 100 new deals have incurred a loss due to unauthorized or fraudulent activity, the sample likelihood will $\propto p^2(1 - p)^{98}$; and if in an external database there were 10 unauthorized or fraudulent deals in the 1000 deals recorded, then the prior density will be $\propto p^{10}(1 - p)^{990}$; so the posterior will be $\propto p^{12}(1 - p)^{1088}$ and the Bayesian estimate of p will be $13/1102 = 0.0118$.

It is clear that there will be great potential to massage operational risk capital charge calculations using targets for N and Bayesian estimates for p , μ_L and σ_L . One hopes that the internal models groups in the regulators have already gained considerable experience with this problem, since internal 'Value-at-Risk' models for market risk also depend on parameters that are notoriously difficult to forecast, e.g. correlations and volatilities.

Conclusion

The Basel Committee working paper CP2.5 on the Regulatory Treatment of Operational Risk has proposed various rules for the pillar 1 capital charge. Banks are currently in the process of making constructive suggestions for the modification of these rules in time for the next consultative paper, due early next year. I do not wish to pass judgment on whether the rules should be accepted. As a mathematician I take the rules as given. Without rules there can be no assumptions and no logical deduction leading to any conclusion.

This article aims to shed light on the models that could be developed within the internal measurement approach. The simplest IMA models assume a binomial or Poisson loss frequency distribution, and these models have relatively few data requirements: it is only necessary to

estimate the expected loss given event and the average number of loss events for the time horizon of the model. However banks that have more data and modelling resources may still chose the IMA because it turns out to be surprisingly flexible. Alternative loss frequency distributions and more sophisticated parameter estimation techniques have been described, but the basic result that capital charges will increase like the *square root* of the size of the banks operations, or the time over which losses are measured, still holds. The IMA also gives a simple formula for recognition of insurance that is risk-sensitive but not excessively complex.

Acknowledgements. This paper has benefited from some very useful comments by an anonymous referee and the usual constructive criticism by my husband, Dr. Jacques Pezier.