Algo Academy Notes

"Algo Academy Notes" is a feature of the *Algo Research Quarterly*. In their article, "The Actuarial Approach to Loss Distributions," **Diane Reynolds** and **Dave Syer** review the motivation behind actuarial models and their application to operational risk measurement. Their goal is to provide the reader with an overview of the principles, concepts, extensions and advantages of the approach.

In the current environment for operational risk measurement, the single overriding concern is lack of data. Because it makes efficient use of data, researchers focussing on the Loss Distribution Approach for regulatory capital (Basel Committee on Banking Supervision 2001), or one of its sisters for economic capital, tend to make extensive use of actuarial models.

Although actuaries themselves do not use the term "actuarial model," it has become common in the field of operational risk management, and so it is adopted for the purpose of this note. In the insurance industry, these models are sometimes known simply as "frequency/severity models."

The basic element of an insurance contract, from the insurer's perspective, is a loss. If the owner of the insurance contract suffers a loss, the insurer assumes the loss under the terms of the contract. To remain solvent, the insurer must estimate how much loss might be incurred annually. Unfortunately, to estimate the annual loss distribution f accurately by directly sampling annual losses z from measured data, one would need many values of z and, hence, many years of data.

Some types of insurance contracts have been issued continually for 100 years or more. However, such contracts are rare. Operational risk data typically do not span more than a few years because its importance has only recently been recognized. Also, if the business environment changes, then data from only a few years ago might no longer be relevant. Thus, efficient use of data is a priority in operational risk measurement.

In insurance, because many policies are issued, many losses can occur in the same year. The actuarial idea is simple: separate the number of losses from the amount that is lost. By so doing, the data are used more effectively. The same is true of operational risk: the actuarial approach allows the estimation of loss severity based on a larger number of data points, or events, than there are years of data.

This note focuses primarily on modelling a single operational loss process. An operational loss process is one that models a type of loss experienced by an operational unit. An operational unit can be abstract, but is often a part of the business, such as a branch, a trading desk or a department. For a more in-depth discussion of operational units and operational loss processes, see Reynolds and Syer (2002).

A mathematical outline of the actuarial model as it applies to operational loss processes is provided, followed by a discussion of the model choice for frequency and severity. Methods for combining frequency and severity models to retrieve the annual loss distribution are presented, as is a simple example of an actuarial approach to calculating a loss distribution. The paper concludes with a brief discussion of codependence between loss processes and some techniques for modelling it.

Model outline

The key element of an actuarial model is to say that the annual loss z is not a single loss caused by a single event, but the result of the aggregation of a number of loss events. This section gives a brief mathematical overview of actuarial methods. The literature (e.g., Klugman et al. 1998 and Frachot et al. 2001) provides much more detailed discussions.

Suppose, in a particular year, that there are n events that each cause a loss, and that their cash values are given by

$$x_i, \quad i = 1, \ldots n \tag{1}$$

then

$$z = \sum_{i=1}^{n} x_i \tag{2}$$

where z is the total annual loss. Understanding the composition of z is facilitated by viewing both x and n as random variables. Thus x has a distribution g such that

$$dp = g(x)\,dx \tag{3}$$

is the conditional probability of experiencing a loss with value in the range $[x, x+dx]$ given that an event has occurred; and n has a distribution h such that

$$p_n = h(n) \tag{4}$$

is the probability of experiencing n loss events in a year. Operational risk events are characteristically very rare, so often $p_0 \neq 0$.

In actuarial terms, x is the "severity" of an event, g is the "severity distribution," n is the "frequency" and h is the "frequency distribution" of the operational loss process.

The annual loss distribution can now be written as

$$f(z) = \sum_{n=0}^{\infty} h(n) g^{(n)}(z) \tag{5}$$

where $g^{(n)}$ is the distribution of annual losses, given that precisely n events occurred.

This formulation has some nice properties that can be exploited to derive the statistical properties of f from the properties of g and h. The expected value and variance of g and h, respectively, can be written as

$$\begin{aligned} E_g(x) &= \mu_x, & \mathrm{Var}_g(x) &= \sigma_x^2, \\ E_h &= \mu_n, & \mathrm{Var}_h(n) &= \sigma_n^2. \end{aligned} \tag{6}$$

The expected value of x over $g^{(n)}$ is $n\mu_x$, and the variance is $n\sigma_x^2$. Thus, the expected value of z is

$$E_f(z) = \sum_n h(n) n \mu_x = \mu_n \mu_x \tag{7}$$

and the second moment of z is

$$E_f(z^2) = \sum_n h(n)(n\sigma_x^2 + n^2\mu_x^2), \tag{8}$$

so the variance of z is

$$\mathrm{Var}_f(z) = \mu_n \sigma_x^2 + \sigma_n^2 \mu_x^2. \tag{9}$$

Using a simulation approach, no particular class of model is required for either frequency or severity distributions. Simulation provides the flexibility to specify the precise forms and calibration methods of both distributions, and most particularly the severity distribution. Some of the common choices are discussed briefly in the next sections.

Frequency

The actuarial approach brings most significant dividends if n is large, however, this is not the case for many operational risk applications. The estimation of h is affected by the small number of years data, but actuaries are generally ready to accept assumptions about h. Principal among these is that independent events have a Poisson distribution:

$$h(n) = \frac{\lambda^n e^{-\lambda}}{n!}, \tag{10}$$

which has a single parameter λ, the average number of events per year. A homogeneous Poisson distribution (one with a fixed average frequency) will often be used where events in a operational loss process are thought to be independent.

Correlations between events in a single operational loss process lead to frequency distributions with characteristically fatter tails (more likelihood of a larger number of events per year). An example is the negative binomial distribution

$$h(n) = \binom{\alpha + n - 1}{n} \left(\frac{1}{1+\beta}\right)^{\alpha} \left(\frac{\beta}{1+\beta}\right)^{n} \tag{11}$$

with $\alpha > 0$, $\beta > 0$. It is interesting to note that the negative binomial distribution can be derived as a mixture of Poisson distributions with different frequencies λ. The distribution of Equation 11 is obtained when λ has a gamma distribution. (See Klugman et al. 1998.) Equation 11 represents a process where there are expected to be $\bar{\lambda} = \alpha\beta$ loss events per year, but there can be more (or less) than this—the standard deviation of event frequency is $\beta\sqrt{\alpha}$.

Other potential frequency models do exist, but are often less pleasant to work with analytically. Fortunately, the scope of available frequency models broadens considerably if simulation is employed to determine the loss distribution. In fact, almost any discrete distribution may be used.

Severity

The severity distribution attempts to model the size of a single loss. It is crucial that this distribution accurately represent the size of possible future losses. Typically, the distribution (parametric or non-parametric) is estimated from past loss experience. However, other data such as scenarios raised in a risk self-assessment or near misses can play an important role in calibrating the tails of the severity distribution. Ultimately, risk is a tail-measure of the loss process. Its tails and, hence, the risk measures, will better represent the true risk if the severity distribution is calibrated using all available data.

Candidate models for severity distributions include a variety of parametric and non-parametric distributions. On the parametric side there are: normal, lognormal, tail-adjusted lognormal, beta, Weibull, Pareto or other standard continuous distributions. Non-parametric choices include various histogram or bucketed representations, or simple resampling of the input data. The list of severity models includes a number of Extreme Value Theory models. These models, particularly those formed from limit processes, are often more appropriately used to analyse the distribution of the annual losses.

Loss processes

Typically, one assumes that loss events are conditionally independent, given the value of n. The advantage of assuming independence is that there are efficient analytic and numerical techniques for evaluating $g^{(n)}$. Given this assumption, $g^{(n)}$ is equal to g convolved with itself n times. This can be written iteratively as

$$\begin{aligned} g^{(n)}(x) &= \int_{-\infty}^{\infty} g^{(n-1)}(y-x)g(y)dy \\ g^{(0)}(x) &= \delta(x), \end{aligned} \tag{12}$$

where δ is Dirac's delta function.

Intuitively, the independence assumption is hard to believe. Consider an operational unit that experiences only loss events of five types: critical, very high, high, moderate and low. If a critical loss occurs, the operational unit will cease operations. This means that after a critical loss, the probability of further losses is zero. If a very high loss occurs, the manager might take out insurance, or enforce a policy change, thus affecting the probability of future losses at each level, or the number of future losses, respectively. The consequence of relaxing the assumption of conditional independence is that $g^{(n)}(z)$ has no explicit form, even as an integral.

Convolution techniques

Several techniques may be employed to obtain the distribution of annual loss. In general, the more analytic the solution, the more restrictive the distributional assumptions that must be made, particularly concerning independence of losses.

Analytic. For certain combinations of frequency and severity models, a full analytic solution is available. Typically, analytic solutions require a Poisson-type frequency, and a severity model with specific properties. For a further discussion of specific analytic solutions see, for example, Panjer and Wilmot (1992).

Simulation. A simulation framework, by contrast, provides a much more practical solution. It is straightforward to implement via programming, and provides a great deal of flexibility. For example, simulation can accommodate some special forms of dependence between frequency and severity through the process of constructing $g^{(n)}$. For instance, there could be a rule of "self-correcting behaviour" such that the severity distribution of the first event in a year is different to subsequent events. Thus, the business learning from its mistakes can be simulated—a large loss will often lead to a lesson being learnt that prevents the same magnitude of loss recurring. In the extreme case that a business unit is closed down after a very large loss, subsequent events would be impossible. This simulation can be accomplished by assigning zero severity to subsequent events.

Semi-analytic. The middle ground is a semi-analytic approach. Using this technique, the losses are assumed to be independent, and a convolution is used to obtain a distribution of losses, conditional on their frequency. A variety of analytic or semi-analytic (e.g., Fourier transform) techniques are employed in the convolution. After this step, the frequency of loss is simulated. The significant benefits here are that fewer scenarios are required to obtain a result of given accuracy, as compared to simulation, and the required assumptions concerning independence and model choice are less rigid than those for the full analytic solution.

Given the same inputs, enough resources (i.e., a large enough simulation), and within the limitations of the least flexible approach (analytic), the results of the three techniques will be the same. The necessary right balance between resource constraints, mathematical beauty and modelling flexibility determines which technique is most appropriate in any given circumstance.

Example

To illustrate the use of an actuarial model in operational risk measurement, a simple example has been devised. The data used are purely hypothetical. This example is intended to demonstrate an application of the model, rather than to provide insight into operational risk itself.

The example shows the calculation of enterprise-wide annual capital for operational risk. There is a single time horizon of one year, and a single operational loss process representing the entire firm. Note that annual capital is only one of the risk measures typically calculated, but it is used here because it is the focus of current regulatory attention. This capital figure could be used, for example, as the starting point for a scorecard allocation of capital among business units.

A list of internal losses, including dates and amounts, is presumed to be available. This data will be used to calculate capital. It consists of six years of losses totalling 293 observed losses, and is summarized in Table 1.

Suppose a Poisson distribution for the frequency of loss and an empirical distribution for the severity of the loss are selected. This requires an estimation of the parameters of the distributions, and an assessment of the appropriateness of the assumptions based on the input data.

The frequency, whose single parameter is its mean, λ, can be calibrated using the average number of events per year over the six years. The result is $\lambda = 48.83$. The Poisson distribution is likely to be appropriate because of the clustering of the number of losses about the mean. This indicates a low vari-

Year	Number	Total Loss	Expected Loss	Standard Deviation
1996	43	4.27 M	99.4 K	79.7 K
1997	44	4.71 M	107.0 K	102.3 K
1998	56	7.26 M	129.7 K	120.6 K
1999	54	5.71 M	105.7 K	98.8 K
2000	46	4.40 M	95.8 K	93.9 K
2001	50	4.87 M	97.4 K	90.5 K
Overall	**293**	**31.22 M**	**106.6 K**	

Table 1: Summary of (illustrative) input data for firm-wide capital calculation (USD)

ance, much in line with the effects of a Poisson distribution. A more detailed backtesting analysis would have to be carried out to formally determine the appropriateness of the model.

To construct the severity distribution, use the given 293 individual loss events, with severity x_i, $i = 1, 2, \ldots, 293$. Their broad statistical properties can be deduced from the data in the table, namely, mean $\mu = 106\text{K USD}$ and standard deviation $\sigma = 99.3\text{K USD}$. Assuming that all previous losses are equally likely to reappear, sampling with replacement can be done directly from the vector of past loss severities. In more formal terms, the implied assumption is that the loss events are conditionally independent, given *n*.

Having determined and calibrated the frequency and severity distributions, simulation may be used to determine the distribution of annual loss for the enterprise. In pure simulation, a sample is drawn first from the frequency distribution, then many samples are drawn from the severity distribution and totalled. This produces a single scenario for total annual loss. Repeating the process thousands of times results in a discrete approximation of the distribution of annual loss.

For comparison, since a large number of events per year are expected, the results of an analytic approach are also provided. In the analytic approach, the Central Limit Theorem is applied to the annual loss distribution. With this approximation, the frequencies are simulated as before, and the firm-wide annual loss distribution can be efficiently calculated to a high degree of accuracy.

The result is shown in Figure 1 as the curve, and the corresponding risk measures are listed in Table 2.

In Figure 1, the firm-wide annual loss distribution is calculated using the input data as summarized in Table 1. The histogram uses a simulation approach with resampling. For comparison, the smooth line uses Central Limit Theorem (CLT). Note that the distribution is only slightly heavy tailed, and is skewed to large losses. The expected loss is $8.11M USD, and the standard deviation is, from Equation 9, $1.37M USD. The differences in the two results of risk statistics are largely due to sampling errors owing to the small number of scenarios used.

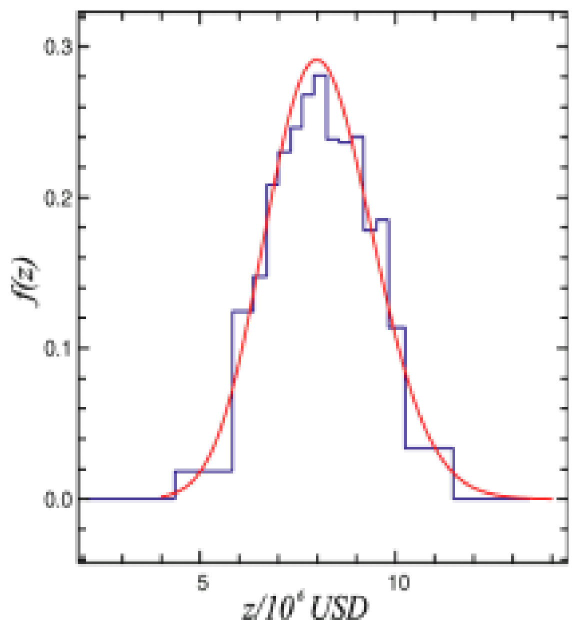

Figure 1: Firm-wide annual loss distributions

Table 2 shows firm-wide risk measures using both resampling and CLT. The expected loss μ is given along with the standard deviation σ. VaR(*p*) is defined as the difference between the p-*th* percentile and the expected loss.

Method	μ	σ	VaR (95%)	VaR (99%)	VaR (99.9%)
Resampling	8.14 M	1.40 K	1.82 M	2.41 M	2.79 M
CLT	5.20 M	1.02 M	1.67 M	2.36 M	3.15 M

Table 2: Firm-wide risk measures using a non-parametric simulation approach (Resampling) compared to those using CLT as an approximation (all quantities are in USD)

Codependence

Suppose an operational unit is modelled by two operational loss processes. The loss distributions of the two processes might be specified actuarially, as above, or in another manner. The total loss distribution for the unit could be determined analytically using a convolution, provided the two processes are independent.

However, it is possible that the loss processes are related in some fashion. For example, both might incur losses due to the same technology failure. If operational risk processes are likely to be correlated, information containing the relations between the two loss processes is also required. The relations (codependences) can be expressed within the actuarial framework in the form of the joint frequency distribution of all *m* loss processes. Suppose loss process *k* has n_k events per year, then the required distribution is

$$h^{(m)}(n_1,\, n_2,\ldots,n_m)$$

with

$$p = h^{(m)}(n_1,\, n_2,\ldots,n_m) \tag{13}$$

equal to the probability of n_1 events at operational unit 1, n_2 events at operational unit 2, and so on. The marginal distributions (Equation 4) are

$$h_k(n_k) = \sum_{n_j \neq k} h^{(m)}(n_1,\, n_2,\ldots,n_m). \tag{14}$$

Because the joint distribution must be specified, the frequency and codependence models are intrinsically linked.

The following sections present some models that contribute to achieving codependent frequency models, such as arrival time models, copulas and latent variable models.

Arrival time models

The frequency distribution *h(n)* can also be expressed as a distribution of arrival times of an event. For example, the simple Poisson case (see Equation 10) can be written in terms of the arrival time *t* as

$$q(t) = \lambda e^{-\lambda t} \tag{15}$$

where $q(t)dt$ is the probability of the next event arriving after *t* years. The arrival time formulation is particularly convenient for some kinds of problems, and can help with the specification of the codependence between different operational loss processes.

To write the joint frequency distribution in terms of arrival times requires $q^m(t_1,t_2,\ldots,t_m)$ with

$$dp = q^{(m)}(t_1,t_2,\ldots,t_m)dt_1dt_2\ldots dt_m \tag{16}$$

equal to the probability of the arrival times being in the infinitesimal neighbourhood of $(t_1,t_2,\ldots,t_m)$. The marginal distribution q_k is given by

$$q_k(t_k) = \int_{t_{j \neq k}} q^{(m)}(t_1,t_2,\ldots,t_m)dt_1dt_2\ldots dt_m\,. \tag{17}$$

Extending the concepts above, arrival time modelling can make it easier to include more complicated ideas in a simulation framework. For instance, instead of $q^{(m)}$ being constant, it could be dependent on the most recent event.

Copulas

An important and convenient mechanism for specifying joint distributions is through the use of copulas (Frey and McNeil 2001, Embrechts et al. 2001). Copulas are a special form of joint distribution of continuous variables, so in this context they would be used to specify $q^{(m)}$.

Latent variable models

Another example of a method to specify the codependence structure is based on latent variables, in the same way that the Merton model operates in portfolio credit risk modelling (Merton 1974, Bucay and Rosen 2000, and Kreinin 2000). This is an important special case of a frequency-codependence structure that is equivalent to specifying the joint frequency distribution. This is not obviously so, since it is expressed in a very different way—as a latent variable model, based on covariate normal risk factors, and event frequency determined by a threshold model.

There is a set of m risk indexes y_k at the operational loss processes, which are random variables with a covariate normal joint distribution. An event at operational loss process k is deemed to have occurred if y_k crosses a threshold η_k. The marginal distribution of frequencies of each loss process is a Bernoulli distribution: possible values of n are zero or one, with probability

$$p = \int_{\eta_k}^{\infty} N(1,0)(x)\, dx. \quad (18)$$

If identical, uncorrelated, operational loss processes with probability p are grouped together, a binomial marginal frequency distribution is obtained for the group. When the group has v members, the maximum frequency is v, and the probability of a single event is p^v. In the limit, p is very small, but $pv \equiv \lambda$ remains finite, so the distribution tends to a Poisson distribution with intensity λ.

Generalizations of the covariate normal approach are possible. Typically, they involve rank correlations and marginal distributions of y_k, which are not normal.

Conclusions

Actuarial models, mainly because of their simplicity, flexibility and efficient use of data, are commonly used in operational risk measurement. This brief overview is meant only as a review and introduction to these models and their rich potential.

References

Basel Committee on Banking Supervision, 2001, "Working Paper on Regulatory Treatment of Operational Risk," September, http://www.bis.org/publ/bcbsca.htm.

Bucay, N. and D. Rosen, 2000, "Applying portfolio credit risk models to retail portfolios," *Algo Research Quarterly*, 3(1): 45–73.

Embrechts, P., A.J., McNeil and D. Straumann, 2001, in *Risk Management: Value at Risk and Beyond*, ed. M. Dempster and H. Moffat, Cambridge: CUP.

Frachot, A., P. Georges and T. Roncalli, 2001, "Loss Distribution Approach for Operational Risk," Crédit Lyonnais: Paris, http://gro.creditlyonnais.fr.

Frey, R. and A.J. McNeil, 2001, "Modelling Dependent Defaults," preprint, ETH, Zurich, http://www.math.ethz.ch/mcneil.

Klugman, S.A., H.H. Panjer and G.E. Willmot, 1998, *Loss Models*, New York, NY: Wiley.

Kreinin, A., 2000, "Generating function method for portfolio credit risk," Research Paper Series 00-04, Toronto, ON: Algorithmics Publications.

Merton, R., 1974, "On the pricing of corporate debt: The risk structure of interest rates," *Journal of Finance*, 29: 449–470.

Panjer, H. and G. Wilmot, 1992, "Insurance Risk Models," USA: Society of Actuaries.

Reynolds, D. and D. Syer, 2002, "Mark to Future and operational risk," *Algo Research Quarterly*, 5(2): 39–51. (This issue)