Using Bayesian Networks to model Expected and Unexpected Operational Losses

30 November 2004

Abstract

This report describes the use of Bayesian Networks (BNs) to model statistical loss distributions in financial operational risk scenarios. Its' focus is on modelling "thick" tail, or unexpected, loss events using mixtures of appropriate loss frequency and severity distributions where these mixtures are conditioned on causal variables that model the capability or effectiveness of the underlying controls process. We conclude that BNs can help combine qualitative data from experts and quantitative data from historical loss databases in a principled way and as such they go some way to meeting the requirements of the draft Basel II Accord, [Basel 2001], for an Advanced Measurement Approach (AMA).

The AgenaRisk software was used to create the example analysis contained herein.

1. Introduction

The Basel Committee on Banking Supervision, in reaction to a number of well-publicised financial disasters, has drafted a system of regulation addressing the issue of Operational Risk (OpRisk) and its assessment [Basel 2001]. Key to the regulatory process is the modelling of a business's operational risks, in terms of a variety of loss event types, in order to arrive at an appropriate regulatory capital charge. To calculate such a charge it is tempting to predict operational risk by building a statistical model based on historical data. However, from a *statistical* perspective the well-publicised financial disasters are too few in number for any meaningful inference. Moreover, until recently, banks have not historically collected loss event data on a wide and systematic basis. This general paucity of loss data means that traditional statistical approaches are unlikely to provide useful predictions of operational losses. A mixture of qualitative and quantitative methods is needed to model operational risks.

Despite the widespread belief that the OpRisk problem is peculiar to the banking community operational risk is not a new topic and is certainly not peculiar to finance. In his book James Reason argues that operational risk is faced by all organisations and he uses examples from the Financial, Rail Transport, Civil Aviation and Nuclear power sectors to support his case [Reason 1997]. Reason identifies a host of reasons why catastrophic failures occur in these safety critical industries including (but not restricted to): a failure to enforce lessons learnt from previous failures; slow degradation or collapse of safety procedures; changes in culture and management; lack of visibility and support for risk reporting and lack of attention to detail. The key conclusion from this is that accidents are not solely the result of human fallibility but are supported by organisational features that fail to defend against all-too-human mistakes, slips and (in the case of fraud) malicious acts. From this we can conclude that OpRisk prediction is inextricably entwined with good management practice and that measurement of OpRisk can only meaningfully be done if the effectiveness of risk and controls processes is regularly assessed. This contrasts sharply with the view that modelling OpRisk simply involves the investigation of statistical phenomena.

By the same arguments financial catastrophes are not a "bolt out of the blue", nor are they inexplicable. The financial scandals such as Barings, Sumitomo and the Allied Irish Bank were all the result of fraudulent activities building up over lengthy periods of time during which active management could have discovered and prevented them. Indeed, if caught early the events would not have been catastrophes at all. There is a tendency to see financial disasters as single "ultra high loss" events rather than aggregations of smaller losses accrued over a period of time. This is understandable given the fact that the losses have to be realised upon discovery, all at once. But this does not change the fact that such losses are accumulated daily and could be detected by good diligence, applied routinely. It is precisely this routine attention to good practice that, just as in safety critical industries, prevents disasters occurring. Any OpRisk scheme should therefore focus on detecting near misses and small losses on a monthly or quarterly basis before they become large losses and disasters.

In this paper we argue that Bayesian Networks (BNs) provide an attractive solution to the problems identified above. BNs enable us to combine any statistical data that is available with qualitative data and subjective judgements about the process. Hence BNs provide a method of modelling operational losses and measuring the effectiveness of a business's operational processes, as part of a self-assessment oriented "Bayesian Scorecard" approach. Using BNs we can:

- Combine pro-active loss indicators, related to the business process, with reactive outcome measures such as near miss and loss data;
- Incorporate expert judgements about the contribution qualitative estimates can make to the overall risk assessment;

- Enter incomplete evidence and still obtain predictions;
- Perform powerful "what-if?" analysis to test sensitivity of conclusions;
- Obtain a visual reasoning tool and a major documentation aid;
- Obtain output in the form of *verifiable* predictions against actual performance measures and loss event rates.

In Section 2 we provide a brief overview of BNs. In Section 3 we consider the widely accepted distinction between expected and unexpected losses in OpRisk, while in Section 4 we explain how BNs provide a unified method of predicting both types of losses. We believe this represents a significant improvement over existing approaches since the distinction is, in our view, arbitrary. We concentrate on the core problem of predicting losses using two BNs to show how they can be used to model loss event frequency, severity and heavy tailed distributions. Section 5 discusses the issues relating to prior estimation in BNs and how prior beliefs can be informed by evidence from the operational process. Finally, in Section 6 we offer some conclusions.

2. Bayesian Networks

The underlying theory of Bayesian Networks (BNs) combines Bayesian probability theory and the notion of conditional independence to represent dependencies between variables. To date BNs have proven useful in many areas of application such as medical expert systems, diagnosis of failures, pattern matching, speech recognition and, more relevantly for the OpRisk community, risk assessment of complex systems in high stakes environments.

A BN is a directed graph whose nodes represent the (discrete) uncertain variables of interest and whose edges are the causal or influential links between the variables. Associated with each node is a set of conditional probability values that model the uncertain relationship between the node and its parents.

BNs enable reasoning under uncertainty and combine the advantages of an intuitive visual representation with a sound mathematical basis in Bayesian probability. With BNs, it is possible to articulate expert beliefs about the dependencies between different variables and to propagate consistently the impact of evidence on the probabilities of uncertain outcomes. BNs allow an injection of scientific rigour when the probability distributions associated with individual nodes are simply "expert opinions".

The key to the successful design of BNs is the meaningful decomposition of a problem domain into a set of causal or conditional propositions about the domain. Rather than ask an expert for the full joint probability distribution, which is obviously a very difficult task, we can apply a "divide and conquer" approach and ask for partial specifications of the model that are themselves meaningful in the experts' domain. Consider, for example, Figure 1. Here we have a BN for a medical problem — the Chest Clinic example. Here a medical practitioner wishes to diagnose one or more diseases: Tuberculosis (*TB*), Cancer (*C*) or Bronchitis (*B*) from a set of causal or diagnostic sources of information: Smoker (*S*)? Visit to Asia (*A*)? Dyspnoea (*D*)? (shortness of breath) and an X-Ray test result (*X*). Rather than simply ask the expert to consider every possible joint event in $p(A,S,B,TB,C,D,X)$ we can instead exploit the causal assumptions that the practitioner may be prepared to make, such as that smoking causes cancer or that bronchitis causes dyspnoea. These causal assumptions (which are shown by arcs in the graph) are then represented as conditional probability statements thus: $p(C|S)$ and $p(D|TBoC,B)$. In general for our Chest Clinic example we might achieve this decomposition (this is mirrored in the structure of the BN in Figure 1):

$$p(A,S,B,TB,C,D,X) = p(B|S)p(C|S)p(TB|A)p(TBoC|TB,C)p(D|TBoC)p(S)p(A) \quad (1.1)$$

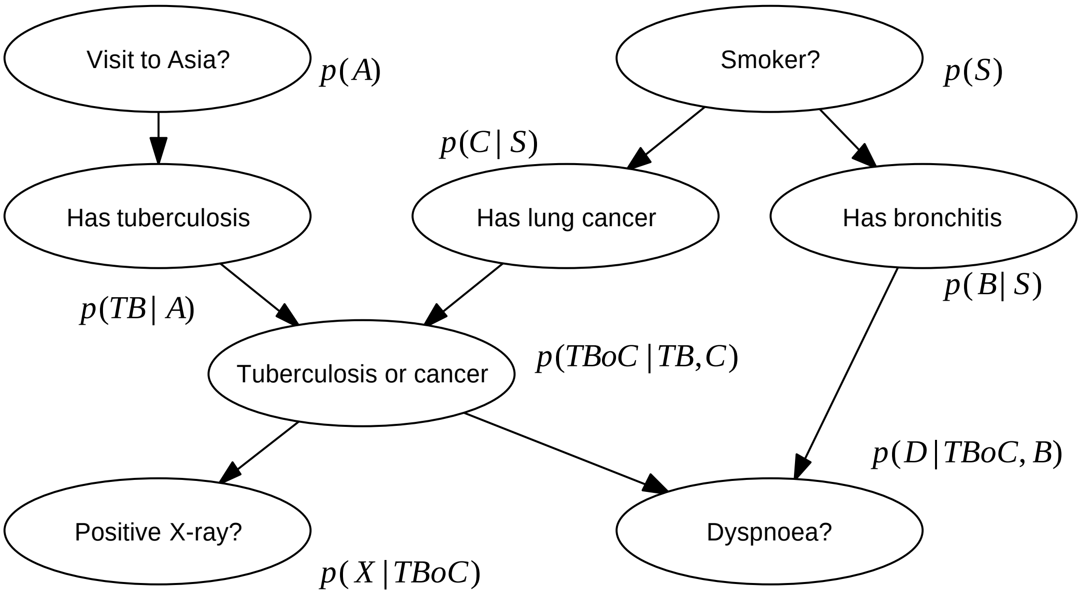

Figure 1 Example Bayesian Network for Chest Clinic Example

Once we have achieved this decomposition we have also, implicitly as a natural product of the approach, specified the covariances by virtue of the conditional probability structure.

Next, we require the expert to model the node probability tables (NPT) for each variable (node): this can either be done using historical data (with standard Bayesian parameter learning approaches or Monte Carlo simulations), or by simply asking the expert to provide a series of subjective estimates. Ideally we would expect these estimates to be based on experience and knowledge rather than blind guesswork.

We can easily embed continuous and discrete statistical distributions within the BN model, as NPTs, and generate values for these NPTs by Monte Carlo simulation methods (for continuous functions we have to discretise the model appropriately).

Once a BN is built it can be executed using an appropriate propagation algorithm, such as the Junction Tree algorithm [Jensen 1996]. This involves calculating the joint probability table for the model from the BN's conditional probability structure in a computationally efficient manner. To do this an intermediate, graph theoretic representation of the BN, called the Junction Tree (JT), is automatically derived from the BN. The JT allows localised, modular, computations to be executed using a message-passing algorithm. This is, in essence, an elaborate form of Bayes' theorem. For full details see [Jensen 1996], [Lauritzen and Speigelhalter 1988], [Pearl 1986] or [Speigelhalter and Cowell 1992]. This process is entirely automatic and can easily be hidden from the domain expert.

Once a BN has been compiled it can be executed dynamically, and exhibits the following two key features:

- The effects of observations entered into one or more nodes can be propagated throughout the BN, in any direction, and the marginal distributions of all nodes updated;
- Only *relevant* inferences can be made in the BN. The BN uses the conditional dependency structure and the current knowledge base to determine which inferences are valid.

It is worth noting that the computational weight involved in using BNs is manageable, in terms of computer memory and permanent storage space, if using an efficient implementation of the JT algorithm. However, many academic, open source and off the shelf packages do not offer implementations that are efficient enough to support large BN models, especially when combined with Monte Carlo simulation. But efficient implementations are possible for the class of BNs needed to model OpRisk problems.

3. Estimating Expected and Unexpected Losses

The Basel report [Basel 2001] classifies financial losses due to operational factors into two “types”:

- Expected losses — These are considered the “normal” losses that occur frequently, as part of everyday business, with a low severity. Examples include losses due to accidentally miscalculated foreign exchange transactions.

- Unexpected losses — These are the unusual losses that occur rarely and have a high severity. Examples include losses resulting from a major fraud activity.

Figure 2 shows the distinction between expected and unexpected losses. The demarcation line is purely arbitrary (in Figure 2 this separation is shown at total losses of $400,000). It therefore makes little sense to use fundamentally different methods for predicting expected and unexpected losses; it is better to think in terms of finding a distribution whose *tail* represents the unexpected losses.

The traditional approach to these kinds of problems is to rely purely on historical data to find the predicted distribution. Where extensive data exists traditional statistical modelling techniques work well. However, in the case of operational loss data, we have a number of special problems. Most importantly, even when a lot of loss data is available, it is unlikely that there will be enough data on the large “unexpected losses” for us to be able to estimate the tail of the distribution properly – usually we end up with tails that are too ‘thin’. Even when modelling the “expected losses” (the bulk of the distribution) the data-oriented approach suffers from the following limitations:

- Loss data will be gathered over a period of time that may represent varying levels of operational effectiveness and risk/threat level. We cannot expect that losses are generated from one single distribution with a small number of “known” parameters;

- Losses experienced are simply a sample of possible events. They may not be representative of changing operational processes. As the underling operational process degrades or improves the value of such historical data wanes;

- The reported loss data might be wrong. Under-reporting and data ambiguity can lead to significant errors in estimation;

- Any attempt to bolster loss data with data gathered from other organisations is subject to the same problems and more because very often the provenance of the data is unknown or in doubt.

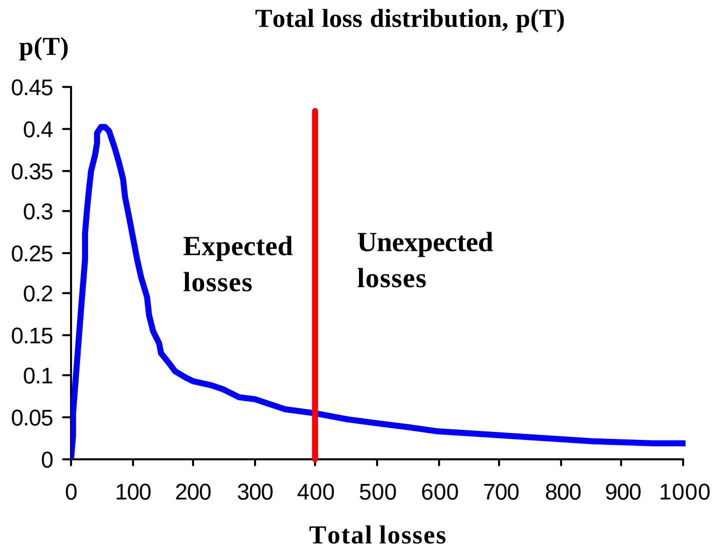

Figure 2 Expected Vs Unexpected Total losses

4. Using BNs to predict losses

Given the serious limitations of approaches based purely on historical loss data it is inevitable that we will have to use methods that enable us to incorporate other types of evidence. Where different types of evidence need to be combined classical statistical methods do not work. Bayesian methods, and in particular BNs, do provide a way forward since they can offer the following specific benefits in OpRisk:

- Explicit approach for combining objective and subjective data by explicitly examining the connection between the operational environment and the loss event process;

- Can model "thick" tail distributions for the unexpected loss component of the total loss distribution;

- A method for eliciting subjective components of risk forecast from experts by explicitly modelling scenarios involving different operational processes or threats to the business, with likely outcomes;

- A verifiable means for dealing with expertise such that models can then be used independently of the expert in the same way a medical expert system would support less qualified practitioners by making "expert" advice available.

In this section we describe two BN models for predicting operational losses. The first predicts total losses from event frequency and severity assuming that these are independent. The second assumes dependence between frequency and severity. These models are by necessity simple and are presented mainly to dissuade readers of some of the misconceptions about how BNs might be used in OpRisk and give a glimpse of the potential of BNs in this area. It should be noted that the calculations here are presented from first principles – in practice all such calculations are performed by special purpose BN software tools and so their complexity is completely hidden.

4.1 Predicting Total losses from Event Frequency and Severity

We can predict total losses from loss frequency and severity where total losses, T, is determined by multiplying the frequency of loss, F, by the severity of loss should it occur, S: $T = F \times S$.

Loss event frequency and severity are random variables, each with an appropriate probability density function (pdf). For a given loss event frequency, *E*, and severity distribution, *S*, we wish to predict the total loss distribution, *T*. The joint probability distribution $p(F,S,T)$ is:

$$p(F,S,T) = p(T|F,S)p(S)p(F) \tag{1.2}$$

$p(F,S,T)$ can be depicted graphically by a BN as shown in Figure 3.

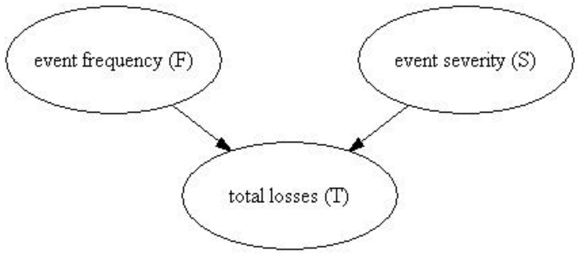

Figure 3 BN for $p(F,S,T)$

Given BNs accommodate the use of Monte Carlo methods to generate the probability tables we do not need to restrict our model to any given family of conjugate probability distributions. For simplicity the prior pdf for event frequency might be represented best by single parameter distributions. Here we use a Poisson distribution, with rate parameter, l, and an Exponential distribution, with parameter, q, as the prior distribution for severity, *S*, thus:

$$p(F) = \frac{e^{-lf} l^{f}}{f!} \tag{1.3}$$

$$p(S) \approx f(S) = q e^{-qs} \tag{1.4}$$

We can generate $p(T|F,S)$ by sampling from S and F using Monte Carlo methods and calculating $T = F \times S$ for each combination of *F* and *S* sampled.

Once the BN has been specified and the NPTs generated we can calculate the marginal probability of any node in the BN by invoking the propagation algorithm. For the marginal distribution of total losses, $p(T)$, this is calculated by marginalising out *F* and *S* from $p(F,S,T)$:

$$p(T) = \sum_{F,S} p(T|F,S)p(S)p(F) \tag{1.5}$$

Given our assumptions we simply need single parameter estimates for the pair (l,q) to generate $p(F,S,T)$. Let us assume that, from past experience elicited from discussion with an expert, the mean loss event rate per year is approximately ten, $l = 9.7$, and that the mean loss severity is forty thousand dollars per loss event, $q = 40$. Then the resulting distribution for $p(T)$, as calculated by the BN is as shown in Figure 4:

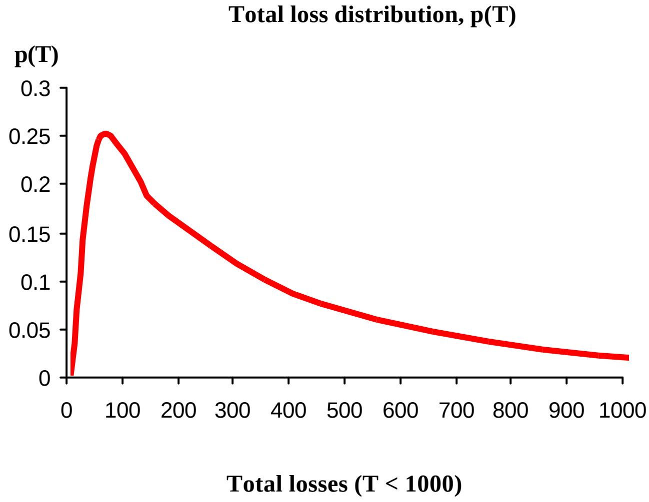

Figure 4 Marginal distribution for $p(T)$ where T < 1000

4.2 Modelling Dependence between Event Frequency and Severity

The model above assumes that loss event frequency and severity are independent of one another. This is optimistic — in reality we can expect them to co-vary. We can easily model covariance between severity, S, and frequency, F, by introducing a common cause, which we will name *process effectiveness*, E, into the BN model.

The new joint probability distribution $p(F,S,T,E)$ is:

$$p(F,S,T,E) = p(T \mid F, S)\, p(S \mid E)\, p(F \mid E)\, p(E) \tag{1.6}$$

This is shown graphically in Figure 5.

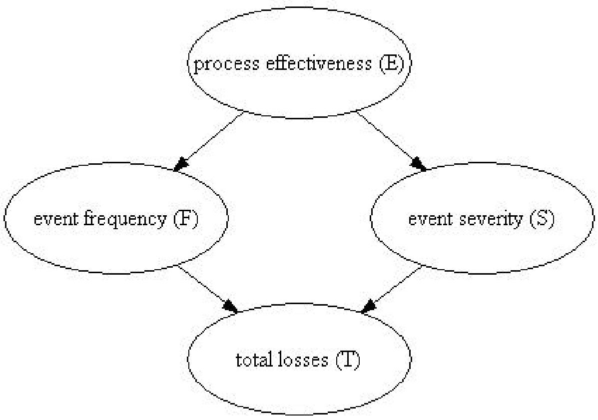

Figure 5 BN for $p(F,S,T,E)$ where *F* and *S* are conditioned on *E*

We can compare this "dependence" model, conditioned on E, with the independent version discussed previously by directly comparing $\sum_{F,S} p(T \mid F, S)\, p(S)\, p(F)$ and $\sum_{F,S,E} p(T \mid F, S)\, p(S \mid E)\, p(F \mid E)\, p(E)$.

To best illustrate the differences between the models we can construct a model with mean values for *p(S)* and *p(F)* that are very close to the original independence model. The example mixture of NPTs chosen are shown in Table 1.

E	*p(E)*	*p(F | E)*	*p(S | E)*
1	0.05	Poisson(0.5)	Exp(5)
2	0.11	Poisson(2)	Exp(10)
3	0.22	Poisson(5)	Exp(20)
4	0.43	Poisson(10)	Exp(50)
5	0.11	Poisson(15)	Exp(60)
6	0.05	Poisson(25)	Exp(70)
7	0.03	Poisson(40)	Exp(80)

Table 1 NPTs for *p(E)*, *p(F | E)* and *p(S | E)*

The expectations for event frequency, *F*, and severity, *S*, are easily derived from the BN. These are almost identical to those used when *F* and *S* are independent: $E(F)=10$ and $E(S)=40$. However the key issue relates to the differences in the tail of *p(T)*, or to put it as a question: Are the "unexpected" losses *larger* when *F* and *S* co vary?

Figure 6 shows the tail probability density functions for values of $T > 1000$ and Table 2 shows a comparison of the mean and 95th percentile (we might assign the 95th percentile as the Value at Risk (VaR) measure) statistics, when *F* and *S* are independent and dependent respectively. We can see that when *F* and *S* are dependent we get a "thicker" tail than when they are independent. This difference in tails is further revealed in the difference in the 95th percentile values for *T*: under dependence the value is 1990 and under independence it is 679. Therefore under independence the VaR measure will be optimistic.

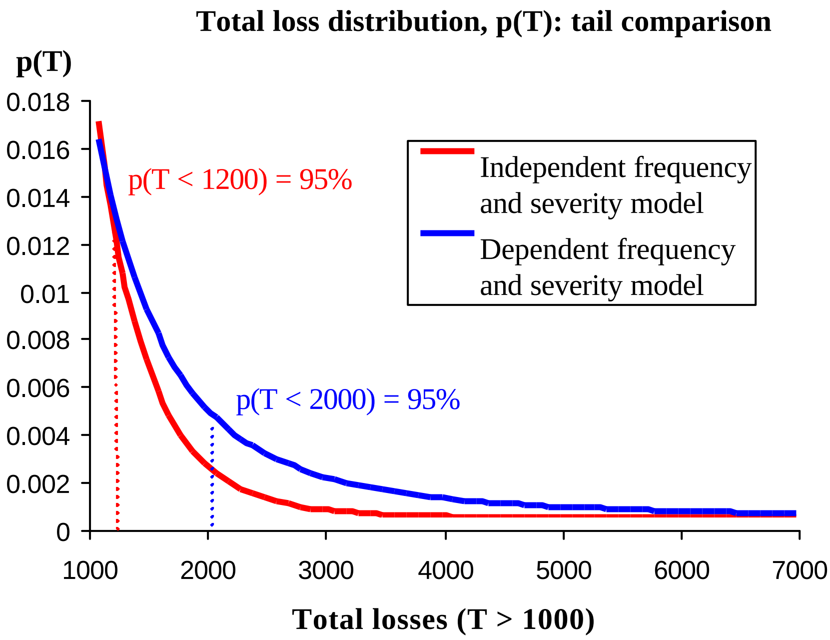

Figure 6 Tail distribution comparisons

If the "unexpected" losses differ so markedly under the different assumptions, how do the "expected" losses differ? To help answer this question we can compare the probability distribution of total losses, *p(T)*, under each assumption as shown in Figure 7, where we plot p(T) in the "expected loss" range (i.e. T < 1000). Here we can see that, under independence, the total loss density is larger in the mid-range between 200 and 800 than it is under the dependence assumption. This is because under the independence assumption small frequency losses can have high severity consequences, whereas under dependence this is not the case. Under independence the estimates of "expected" losses are therefore inflated.

		F and S	
		independent	dependent
Frequency, F	**mean**	9.7	10.8
	95th percentile	14.9	28.0
Severity, S	**mean**	40.0	43.0
	95th percentile	51.0	75.0
Total losses, T	**mean**	391.0	621.0
	95th percentile	679.0	1990.0

Values in 000's $

Table 2 Mean and 95th percentile statistics for *F*, *S* and *T* under dependence and independence assumptions

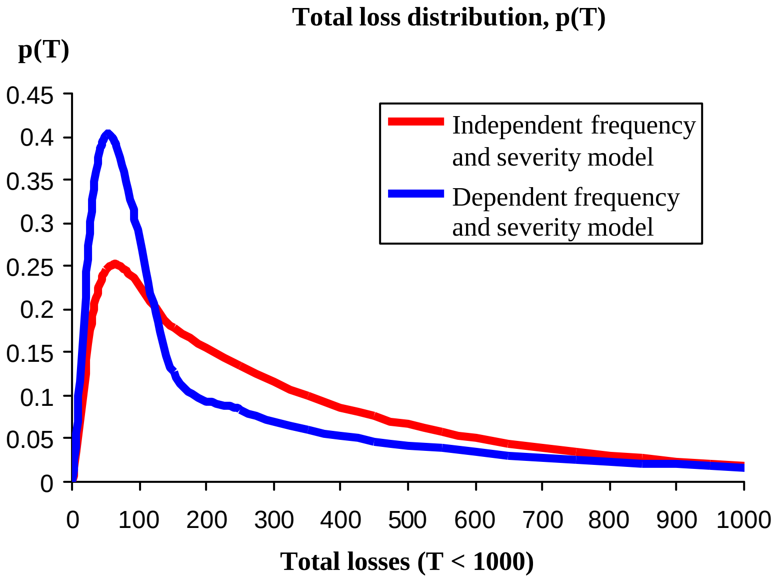

Figure 7 pdf for *p(T)* comparisons

5. Modelling Process Effectiveness

Estimating the prior distribution for process effectiveness, *E*, is very important if we are to arrive at sensible estimates for the total loss distribution, *p(T)*. The reason it makes sense to estimate *p(E)* rather than simply estimate *p(T)* directly is because *p(T)* is a compound measure of event frequency and severity; experts find it hard to perform the mental calculations to combine these directly. They find it much easier to break it down into separate assessments of severity and frequency. Moreover, experts have direct experience of the process; they are involved in it, they have an intimate understanding of the controls, procedures, staff and threats the business may face. We can exploit this direct experience to elicit information about the process separately from the outcomes of the process (i.e. event severity and frequency).

Thus, to populate the model we need to assess $p(S|E)$, $p(F|E)$ and $p(E)$. We can interpret *E* as the level of maturity, or effectiveness, in preventing undesirable events. The measurement scale for *E* given in Table 1 assigns "one" for the most effective process and "seven" for the least effective. We could, of course, label these differently and give detailed descriptions of the mix of operating procedures, technology, staff skills etc assigned to each level and doing so would make process effectiveness an observable quality.

Now we have some model for E, we need to consider what the prior probability distribution actually means in Table 1. If we examine $p(E=1)=0.05$ we can interpret the subjective beliefs (probabilities) of the expert in a number of ways:

- $p(E)$ reflects uncertainty about the current process — The chance of *this* process at this moment in time being a level one process is 1 in 20.
- $p(E)$ reflects the frequency of occurrence of a changing process — Over 20 years we would expect our *changing* process to be equivalent to level one (the best), on average, once.

The first interpretation echoes the Bayesian definition of probability, as a degree of belief, and the second echoes the Frequentist definition. Both of these interpretations are mathematically equivalent but have subtly different semantics. The first implies that we never really know how "good" the process actually is (we are in doubt) and the second suggests that the process changes or evolves over time in some way i.e. there is some long-run frequency with which the process occurs at a particular level. Both interpretations are correct but an expert may find it easier to think in subjective rather than frequency terms or vice versa.

The next possible concern about E is that the experts may not have experienced the whole range of E directly, and will thus be unable to make any confident statement about $p(S|E)$, $p(F|E)$ and $p(E)$ over this range. However, experts are very good at anticipating and then recommending or taking action to avoid unpleasant outcomes (this is why we employ humans after all!). Indeed, the whole thrust of regulation is to force businesses to anticipate and act in this way. With this in mind we can generate, with the help of the expert, hypothetical scenarios involving E and asking:

- What is the chance of this level of degradation or improvement occurring in the process, $p(E)$?
- Given the process effectiveness is at a particular level, what would the event severity, $p(S|E)$, and frequency, $p(F|E)$, distributions look like?

In our experience we have found that experts are willing and able to answer these two questions if they are asked in a structured fashion *and* they can quickly see the results of their assertions in the BN and then refine it.

Answering these questions may be difficult if they are do not appear grounded enough. Under these circumstances we could extend the BN model to identify causal factors related to E, such as that shown in Figure 8, where we have identified another layer of causes in the form of three causal factors: operational procedure quality, O, technology fidelity, TF, and staff quality, Q. Collecting priors on these might then be easier for the expert.

Figure 8 Refinement of BN model to include another layer of causes

From a statistical modelling perspective the probability distribution for process effectiveness, *E*, simply acts as a mixture parameter that mixes a set of statistical or empirical frequency and severity distribution models that will result in a "thicker" tail and a more realistic *VaR* estimate [Venkataraman 1997]. We obviously favour a causal interpretation of *E* rather than a purely statistical one in order to focus expert attention on the underlying generative process and hopefully to generate a more sensible Bayesian model.

6. Concluding Remarks

BNs can help combine qualitative data from experts and quantitative data from historical loss databases in a principled way and as such they go some way to meeting the requirements of the draft Basel Accord, [Basel 2001], for an Advanced Measurement Approach (AMA). Adopting a BN based approach should therefore lead to better operational risk governance and a reduced regulatory capital charge. Relying purely on historical loss data and traditional statistical analysis techniques will neither provide good predictions of future operational risk losses, nor a mechanism for controlling and monitoring such losses.

We have shown how BNs can be used to model operational risk via two small examples in which total losses are based on event frequency and severity. In the second of the models we took account of the possible dependence between frequency and severity by introducing a common cause *process effectiveness*, *E*, and we showed that we could use this BN to model "heavy tailed" distributions in a way that would exploit the expertise available within an organisation. This helps produce a realistic and reproducible Value at Risk (VaR) estimate. We call the method used to construct the total loss distribution the "Bayesian Scorecard" because of the obvious similarities to the less sophisticated Scorecard measurement process.

BNs focus on assessing the effectiveness of the underlying business process and we propose using them as a form of self-assessment. This would involve monitoring the underlying business process on a frequent basis (such as quarterly or monthly) and translating these self-assessment scores into total loss predictions via the BN.

We could go further and automatically *learn* from loss data, as it is observed, and self-assessment data *together* as part of a dynamic approach. Likewise entering hypothetical self-assessments and notional loss data into the BN model would then easily support "What-if" and sensitivity analysis. Such analysis would then help assess the accuracy of and help recalibrate their expertise over time. These topics, and others, will be covered in subsequent papers.

References

[Basel 2001] Basel Committee on Banking Supervision. Working paper on the Regulatory Treatment of Operational Risks. http://www.bis.org, 2001.

[Jensen 1996] F. V. Jensen. *An Introduction to Bayesian Networks*. UCL Press, London, 1996.

[Lauritzen and Speigelhalter 1988] S.L. Lauritzen and D.J. Spiegelhalter, "Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems (with discussion)". *Journal of the Royal Statistical Society Series B*, Vol. 50, No 2, pp.157-224, 1988.

[Pearl 1986] Pearl J., "Fusion, propagation, and structuring in belief networks", *Artificial Intelligence*, Vol. 29, 1986.

[Reason 1997] Reason J. *Managing the Risks of Organisational Accidents*. Ashgate Publishing Limited, 1997.

[Speigelhalter and Cowell 1992] D. J. Spegelhalter and R.G. Cowell. "Learning in Probabilistic Expert Systems", *Bayesian Statistics*, 4, pp. 447-465. Oxford University Press, 1992.

[Venkataraman 1997] S. Subu Venkataraman. Value at risk for a mixture of normal distributions: the use of Quasi- Bayesian estimation techniques. Economic Perspectives, Federal Reserve Bank of Chicago, 1997.