

MODELLING OPERATIONAL RISK IN FINANCIAL INSTITUTIONS USING BAYESIAN NETWORKS

By

YEW KHUEN YOON

1 May 2003

A dissertation submitted for the award of the degree of Master of Science

In

Actuarial Management



Faculty of Actuarial Science and Statistics

Abstract

Recent failures and losses of financial institutions have prompted the regulatory authorities such as the Financial Services Authority in the UK and the Bank for International Settlements to highlight the importance of adequate systems and controls to deal with operational risk (OR). An important measure taken is requiring these organisations to allocate OR capital. This has given the industry an impetus to develop methods for measuring and modelling OR.

Although most firms have only begun to collect data and initiate modifications to their systems to deal with quantification of OR, there has not been a shortage of methods proposed. Broadly, these can be divided into (i) top-down methods that starts with enterprise-wide risk, which is then allocated to the business units and (ii) bottom-up methods that analyses risk at the business unit level and then builds up to form a firm-wide evaluation of OR. The consensus appears to be that bottom-up methods will be the favoured approach in the long run as data is gathered and as more sophisticated methods develop. The preferred approach is one that deals with the cause and effects of OR events.

Amongst the more popular methods are linear models such as time series models, econometric models and empirical actuarial models. Since OR losses tend to have ‘fat-tailed’ distributions, extreme value theory has been proposed as a solution. These models tend to rely too much on data availability, a problem due to the non-standard and infrequent nature of operational loss events. Furthermore, the complex interaction between OR variables often render linear methods mathematically intractable. Various non-linear methods have been proposed to overcome these problems, one of which is Bayesian networks (BNs).

BNs are a combination of graph theory and Bayesian statistical theory to produce networks of variables (called nodes) that are linked via directed edges. These edges imply a causal relationship between nodes, hence BNs are also classified as a type of *Directed Acyclic Graph*. Expert input and past data are combined in the process of specifying the structure and the underlying probability distributions at each node. Rules on conditional independence of the nodes allow alternative causes to be ‘explained away’ when events occur – providing the foundation for causal analysis.

The BN can be structured in such a way to have the various OR factors cascade into an overall loss distribution, thus allowing risk capital to be specified. The causal dependencies inherent in the structure make BNs useful for scenario analyses in identifying and measuring the impact of operational loss events. This would meet the main requirements of a model for OR in financial institutions.

Developments in an algorithm known as *triangulation* allow directed graphs to be expressed as a network of undirected sub-graphs. This allows computations to be performed locally and therefore more efficiently. Bayesian statistical theory is applied in a straightforward and intuitively appealing way to update the probabilities in the model as new data arrives – this process is called *learning*. Model assessment is also relatively easy to do via monitors that quantify in real-time how well the model prediction compares to actual data.

An illustration of these aspects is provided by way of an online insurance business belonging to a fictitious company. A BN is set up and used to model the risk of failure in the system network supporting this business. The BN is used to calculate risk capital, scenario testing, causal analysis and simulation of future scenarios. The learning process is also shown and monitors used to discriminate between models with and without learning.

Contents

<i>Item</i>	<i>Page</i>
List of Abbreviations.....	iv
Chapter One: Developments in Regulatory Capital Requirements Driver for Operational Risk Modelling.....	1
Chapter Two: A Brief Survey of Existing Methods for Modelling Operational Risk	4
Chapter Three: Bayesian Networks	12
Chapter Four: Applying Bayesian Networks to Modelling Operational Risk	39
Chapter Five: Discussion	60
Post-script.....	63
Acknowledgements.....	64
References	65
Appendix I: Network Risk – Prior Distributions.....	67
Appendix II: Network Risk – Propagation Example.....	69
Appendix III: Network Risk – Simulation Results	75
Appendix IV: Network Risk – Calculations for Parent-Child Monitor for NF Node	77

List of Abbreviations Used

AMA	Advanced Measurement Approach
BN	Bayesian network
BCBS	Basel Committee for Banking Supervision
DAG	Directed acyclic graph
EVT	Extreme value theory
FI	Financial institution
LS	Logarithmic score
NR	Network risk
OR	Operational risk
VaR	Value-at-risk

Chapter One

Developments in Regulatory Capital Requirements as a Driver for Operational Risk Modelling

In the past decade, the financial world has been hit by catastrophic failures and losses – the causes of which are largely non-financial in nature. Names such as Barings, Daiwa Bank, Orange County, Metallgesellschaft, and Long Term Capital Management have entered the vocabulary in the history of operational risk (OR). OR has been increasingly recognized as an important source of business uncertainty, especially in financial institutions (FI), thus being given much attention by the financial community, the public and not least of all, the industry regulators.

It is in such an environment that the new framework for capital adequacy proposed by the Basel Committee for Banking Supervision (BCBS) of the Bank for International Settlements has highlighted the importance of OR in determining the adequacy of regulatory capital¹. OR has been defined by BCBS as:

“The risk of loss resulting from inadequate or failed internal processes, people and systems or from external events.” (p2, BCBS (2001a))

A regulatory standard for capital adequacy calls for quantification of risks that are to be supported by the economic capital of FIs. Naturally, there has been increasing interest in methods of measuring and modelling OR in response to these developments.

¹ See ¶82 – 87 of BCBS (1999).

In the United Kingdom, the Financial Services Authority (FSA) has set a 2004 target date for its policy on OR management systems and controls (scheduled to be published in 2003) to take effect, with consultation currently well under way². With the release of the Integrated Prudential Sourcebook by the FSA, FIs will be required to demonstrate adequate systems and controls in managing OR. Invariably this would involve some level of scrutiny of the company's framework for measuring and modelling OR.

BCBS is proposing a continuum of approaches to the allocation of capital to OR, starting from the Basic Indicator Approach where a fixed percentage is applied across the board, to the Standardised Approach where the percentage to be applied for each business line varies as defined by the supervisory authority and finally, the Advanced Measurement Approach (AMA) where the bank may use an internal mechanism. The policies proposed encourage banks to develop more sophisticated methods by allowing those using the AMA a possibility of lower capital charges (up to 25% lower than the Standardised Approach)³. However, only banks that meet a set of qualitative and quantitative criteria set by BCBS may use the AMA.

The interesting question would be "How prepared are banks and FIs in general to operate an internal system of OR management and capital allocation?" Or more importantly, "What sort of OR management system would be sufficient to meet the criteria set by BCBS?" Authors on the subject reveal that most FIs do not have sophisticated treatments of OR. King (1999), one of the earlier reference texts on the subject, summarized the current framework used by organizations into the following broad categories:

² See Consultation Paper 142: Operational risk systems and control and the feedback on the FSA's website: www.fsa.gov.uk

³ For more details, refer to BCBS (2001a) and chapter 2 of van den Brink (2002).

Framework	Description
<i>Control self-assessment</i>	This involves questionnaires and workshops to solicit qualitative opinion of employees on areas of risk in the organization. The subjective nature of this method restricts its suitability for capital allocation and management performance review.
<i>Process Analysis</i>	Operational processes are analysed in detail at task level to identify and classify causation of loss events. The problem of subjectivity remains and the documentation is difficult to maintain.
<i>Loss Categorisations</i>	Losses are entered into a database and attributed to loss categories with sufficiently detailed classification for risk measures. Double-counting and identification of causes are challenges to overcome.
<i>Performance Analysis</i>	Performance measures (such as Value Chain and Balanced Scorecard) are used to develop associated risk measures, ensuring consistency between performance and risk. However, the focus is more on gains and not losses.

More recently, Hoffman (2002) sounded a more optimistic note by stating “The best firms are making heroic strides toward risk definition, data collection, aggregation and first-level analysis” (p7). Some promising trends were identified, including the creation of internal loss event databases, integration of internal and external loss data for experimental risk capital calculations, and new risk information measures are emerging.

In the light of these developments, it would be interesting to see what solutions are currently available for FIs to turn to in handling this important area in risk management. Hence, we proceed to discuss some of the existing methods available for modelling OR.

Chapter Two

A Brief Survey of Existing Methods for Modelling Operational Risk

2.1 Chapter introduction

This chapter is an overview of the various existing methods for measuring and modelling OR, both in theory and what is used in practice. We will discuss briefly how useful and appropriate these methods are in meeting two vital business requirements of the whole OR modelling endeavour: (i) determine risk capital and (ii) identify causation of OR in a way that is quantifiable and actionable. The considerations will be made vis-a-vis the underlying developments in the regulatory environment, which was briefly discussed in the introductory chapter.

2.2 Top-down and bottom-up approaches

The whole spectrum of methods available for modelling OR can be broadly divided into two main categories:

- (i) *Top-down approaches* start from a firm-wide determination of OR. This is then allocated to business units according to how much each unit is thought to contribute to the overall OR. Some examples include earnings volatility, scorecard approaches and CAPM-based modelling.
- (ii) *Bottom-up approaches* model the risk of loss events at the business units and build up to an overall picture of firm-wide OR. This is where the latest modelling methods are being

developed with increasing involvement of the actuarial profession. The main examples in this area are loss distribution analysis and causal modelling⁴.

At the moment, most organizations are still in the process of classification of OR losses and gathering data. As such, top-down approaches are still the main methods of managing, measuring and modelling OR. In general, as companies accumulate more extensive internal loss databases and as modelling techniques get more sophisticated there will be a move towards bottom-up approaches.

2.3 Operational risk capital

The broad approach to setting of economic capital for banks at the moment revolves around obtaining a single figure known as the Value-at-Risk (VaR). This is often defined as a high quantile of a loss distribution that is projected over a required period based on past data. Thus, for example, BCBS may require that “the bank must be able to demonstrate that the risk measure used for regulatory capital purposes reflects a *holding period* of one-year and a *confidence level of 99.9%*” (BCBS (2001b)). A whole body of literature exists detailing the workings of VaR in setting economic capital for market risks and credit risks (see, for example, Jorion (2001) and Dowd (1998)).

Presently, if VaR is calculated for OR, the approach is commonly rather simplistic – ranging from the residual approach (firmwide VaR minus market VaR and credit VaR) to setting the VaR as a multiple of the standard error. Some of the more advanced approaches use models that essentially try to fit a loss distribution from which a VaR figure can be extracted⁵. With market VaR and credit VaR this is more straightforward as an underlying normality is assumed to exist in the data. However, the nature of OR losses are such that they either occur fairly often in a majority of minor

⁴ Ceske et al. (2000) and Marshall (2001) have fairly comprehensive tables on the advantages and disadvantages of various examples of both approaches. Both authors agree that bottom-up models are the way forward.

⁵ Several recent texts on OR give descriptions of the methodology e.g. Cruz (2002), Hoffman (2002) and King (1999). See also BCBS (2001), chapter 19 of Jorion (2001) and p198 of Dowd (1998).

cases (high frequency and low severity) or very rarely with catastrophic consequences (low frequency and high severity) thus the assumption of normality is not appropriate. The former type of losses is less of a concern as sufficient information is available on them to calculate to a fair degree of accuracy the standard risk charges (say, via pricing of products) or reserves to absorb the costs. It is the latter that is more of concern in allocating economic capital. It is also the main motivation behind most OR modelling endeavours.

The frequency of OR losses are usually modelled using counting processes. The severity of OR losses are generally modelled using ‘fat-tailed’ distributions, or distributions with high kurtosis. The more popular empirical distributions include Poisson and negative binomial distributions (for frequency), and lognormal, gamma and Weibull distributions (for severity). However, due to the lack of data at the extremes, it was found that certain asymptotic distributions which had their origins in the natural sciences (e.g. in the modelling of hurricanes and sea levels) work fairly well when applied to the tails of OR data. This branch of statistics is called Extreme Value Theory (EVT) and consists of two main families of distributions: the Generalised Extreme Value (which models maxima and minima) and the Generalised Pareto Distribution (which models data above a selected level, also known as peaks-over-threshold).

2.4 Causal modelling of OR

The approach taken by BCBS in the classification of OR losses is to deal in terms of the **causes** of the loss. This can be seen from the definition of OR by BCBS given above and their proposed loss event type classification (See Annex 2 of BCBS (2002)). The Operational Risk Working Party of the Institute of Actuaries proposed a framework for analysis of OR based on cause and consequence in

their recent report to GIRO 2002. Throughout the report, they stressed the importance of working in terms of causes rather than consequence to avoid double-counting or omissions⁶.

Modelling of causation of OR losses is not only a neat framework for comprehensive analysis. It is a vital basis for the understanding of the how the risk of OR losses arise within the structure and operations of the organization. It also provides a basis on which management may intervene to achieve the desired alteration in risk profile. Clearly, causal modelling is not only crucial in understanding and managing risks internally; it will be of necessity in the new regulatory regime.

Some existing techniques currently used for causal modelling in risk management include *time-series analysis* and *econometric* models. A common model is the *Autoregressive time-series with Conditional Heteroscedasticity* (ARCH). *Factor analysis* is also used to decompose uncertainty in profit and loss figures into various causal factors of manageable sizes. Chapter 8 of Cruz (2002) provides a brief overview of these methods. Non-linear and non-parametric models are increasingly popular for causal modelling as they offer more flexibility and as more research is done in these areas. Some examples include neural networks, fuzzy logic and Bayesian networks, which is the focus of this paper.

2.5 Challenges of modelling OR

The besetting problem with any attempt to model OR losses is the lack of data. Internal data of low frequency events are rarely sufficient to model the loss distributions to the required accuracy.

The usage of parametric loss distributions requires parameter estimation based on the existing data. As such, these methods tend to be too dependent on data quality and quantity. Even for the more promising EVT methods, the small sample sizes result in the shape of the tails being very sensitive

⁶ See Institute of Actuaries (2002) especially ¶2.2.1 and ¶3.5.

to inclusions or exclusions of single events – implying a greater degree of subjectivity than may at first appear.

There is also the challenge of model selection. Most goodness-of-fit tests use actual data to verify the theoretical model, which itself is obtained using the actual data. Lack of data exacerbates this inherent cyclical problem in model selection. As regards EVT the existing tests either require visual inspection or tend to over-fit (Cruz (2002)). The peaks-over-threshold method, for example, requires a visual inspection of the *mean excess plot* (the graph of threshold versus the mean of all points over the threshold) to identify the point where the graph begins to behave linearly to determine the threshold on which the rest of the modelling is based. King (1999) gets around this problem by fixing the threshold as a quantile of the empirical distribution. Points below the threshold are modelled by the empirical distribution whereas points above the threshold are modelled using EVT.

In an attempt to overcome the lack of internal data, the usage of external data has been suggested as a solution. The idea is that data from diverse locations may share certain features, thus making sense to combine them to obtain a larger database. To achieve this, external data need to be carefully filtered and scaled. Cruz (2002) describes a method for pooling data from different locations called ‘frequency analysis’ and Frachot & Roncalli (2002) propose using linear credibility (a method commonly used for premium rating in general insurance) to combine internal and external data. There have also been some efforts made at industry-wide level collection of data to provide a shared source of external data (e.g. British Bankers Association). Obviously, merging internal and external data is not a straightforward exercise – especially in the case of operational loss data. There are qualitative problems such as quality of data from external parties, the dissimilarity of OR management practices across firms, lack of detailed breakdown of data and lack of up-to-date data that may pose difficulties to using external data to set internal capital requirements (which are essentially prospective).

In any case, the modelling of OR losses would require an in-depth understanding of the complex nature of company operations and a model that adequately represents such complexities. Linear models may need to be mathematically intractable to sufficiently achieve this objective. As a logical step forward, various non-linear models have been proposed and experimented with in as yet a limited number of financial and non-financial institutions.

2.6 Expert input and Non-linear Methodologies

The lack of data and complexity of operations in FIs intuitively suggests the inclusion of *expert* input. An expert in this case would be anyone whose knowledge and expertise enables him/her to make sufficiently credible conjectures about how company operations affect the company's risk profile. Such input can be used as a proxy for data and yield valuable information about the complexity of company operations that is difficult to capture from data alone. The challenge of the modeler, then, is to incorporate such input into the overall OR modelling framework.

It has been found that qualitative information, such as management decisions, competencies and preferences, can be better incorporated into a measurable (and hence quantitative) framework using non-linear methods. Some of these include *fuzzy logic*, *neural networks*, *system dynamics*, and *Bayesian networks*.

Fuzzy logic uses a multivariate logical set that recognizes that human decisions are often not binary (e.g. Yes/No, Hot/Cold) by allowing gradations in its formulations (e.g. rather hot, very hot etc.). Hoffman (2002) has a brief case study on how this has been used in a bank in its OR management⁷. Although this method has advantages in its ease-of-understanding, "...the theory of fuzzy logic

⁷ See pp 326-330.

cannot replace robust statistical methods in measuring operational risk in a capital-at-risk sense...⁸
Thus, its application for our present concern is limited.

Although not strictly a method that is used for incorporation of qualitative opinion, *neural networks* are useful for modelling complex relationships between variables that would be difficult to do using linear methods. The network consists of nodes with values of input, output and intermediate variables. Data mining techniques are used to ‘train’ the model by using complex algorithms that learn the relationships between the variables. The model is then calibrated such that its output is as close to the actual data output as possible. A drawback of this approach is its heavy reliance on the availability of data. A brief example of its application can be found in Hoffman (2002), although it involves a non-financial institution⁹.

System dynamics was developed by Jay Forrester of the Massachusetts Institute of Technology and has been promoted by Tillinghast-Towers Perrin¹⁰ as an OR modelling solution. This approach involves using expert input to map a network of cause-and-effect relationships between variables affecting the OR of a business unit. The relationship between each cause-and-effect set of variables is then quantified by combining data and expert input to obtain a plot on two axes (one for each of the cause and effect variables).

This dissertation discusses the usage of *Bayesian networks* (BNs) as a framework for modelling OR. We shall see that it is an efficient and intuitively coherent methodology for incorporating expert input. In addition, BNs are useful for capturing causal dependencies. This satisfies a vital requirement of any OR modelling framework: ability to model causation. Developments in the field of graphical models (of which BNs are an example of) have made BNs very user friendly and as

⁸ Hoffman (2002) p327.

⁹ More specifically, the National Aeronautic and Space Administration (NASA) in the United States of America. See pp299-300.

¹⁰ Miccolis & Shah (2000) and Shah (2001).

such, a widely accepted tool for making inference from complex networks of causal relationships. For example, BNs have been used for almost two decades now in medical diagnoses and more recently in criminal forensic sciences using DNA analysis. Underlying BNs is the powerful Bayesian statistical property that allows the combination of subjective input and empirical observations. This lends it very well to situations with a high degree of uncertainty and where data is costly or sparse.

Alexander (2000) provides a brief introduction to modelling OR using BNs. Marshall (2001), Cruz (2002) and Hoffman (2002) give brief overviews of BNs and where they fit into the whole framework of OR modelling. There is also an illustrative albeit high-level discussion on causal modelling using BNs via a banking example in King (1999).

2.7 Chapter conclusion

We have seen in this chapter that although many new methods have been proposed in the past few years, many of these methods place heavy reliance on the data used and not enough on implicit knowledge of experts. Furthermore, most of these methods are linear and thus do not fully utilise nor capture causal relationships inherent in the data – a step that is vital considering the regulatory emphasis on causal identification and action. BNs have been proposed as a potential approach as it offers intuitive yet mathematically and computationally tractable means of dealing with these two aspects.

In the next chapter, we will proceed with a cursory recall of basic concepts of Bayesian statistics before examining how BNs work.

Chapter Three

Bayesian Networks¹¹

3.1 Chapter Introduction

We will start off by looking at some fundamental inferential properties of Bayesian statistical theory, specifically *Bayes' theorem* and *probability calculus*. Then, some of these ideas will be combined with a field of information sciences known as *graph theory* and applied to a well-known example to illustrate the basic building block behind BNs: *the Directed Acyclic Graph*. The rest of the chapter will develop the theory behind BNs and how it can be used for probabilistic inference by incorporating expert opinion and data.

3.2 Bayesian Statistics

3.2.1 Introduction

In classical probability theory, sample statistics are assumed to belong to a certain population with a specified distribution, which is defined by a set of parameters that have a fixed value. The main task of the statistician is to estimate these parameters as best as possible based on whatever data is available (i.e. the sample statistics). Where possible 'experiments' are performed repeatedly to obtain a suitably large sample to assign values to these parameters.

¹¹ The theory and techniques discussed in this chapter follow closely the ideas presented in Cowell et al (1999). An alternative text with some useful introductory examples is Jensen (1996).

Bayesian statistics allow these parameters themselves to be random variables. Furthermore, assertions made regarding the characteristics of a population are necessarily dependent not just on empirical observations or data (objective information) but also on any knowledge available to the statistician *prior* to making the observations (subjective information). This knowledge may come in the form of data from a different location that is considered to have a certain degree of relevance to the observed population. More importantly, it may also come in the form of information obtained from knowledgeable parties i.e. experts whose familiarity with the subject matter makes them a credible source.

As a consequence, where classical probability deals predominantly with assessing *unconditional* probability statements such as “the probability of event A is x”, denoted $P(A) = x$, Bayesian statistics vocabulary expand into *conditional* probabilities with statements such as “the probability of event A **given** event B has occurred is y” (or “the probability of event A **conditioned on** event B has occurred is y”) denoted $P(A|B) = y$.

3.2.2 Bayes’ Theorem and Probability Calculus

Manipulation of such probabilities involves treating them as functions of variables using certain well-known rules, collectively known as *probability calculus*. We introduce the *fundamental rule* in probability calculus:

$$P(A, B) = P(A | B)P(B) \tag{3.1}$$

This states that the probability of joint occurrence of events A and B is equal to the probability of A conditional on B *multiplied by* probability of event B. Since the function $P(A, B)$ is symmetrical, we can express it in the following form:

$$P(A, B) = P(B | A)P(A) \tag{3.2}$$

Equating the right hand sides of (3.1) and (3.2) we obtain the definitive theorem in Bayesian statistics, *Bayes’ Theorem*:

$$P(B | A) = \frac{P(A | B)P(B)}{P(A)} \quad (3.3)$$

This can be interpreted as follows: the *posterior* probability $P(B | A)$ is equal to the *prior* probability $P(B)$ multiplied by the ratio $P(A | B) / P(A)$.

3.2.3 Inference Using Bayes' Theorem

Applying Bayes' theorem to random variables (e.g. one where observations \underline{x} sampled from a population with a probability distribution defined by the parameter θ) we obtain the following form:

$$P(\theta | \underline{x}) = \frac{P(\underline{x} | \theta)P(\theta)}{P(\underline{x})} \quad (3.4)$$

This allows us to understand Bayes' theorem in terms of the discussion above on objective and subjective information. The *prior* is the subjective information alluded to earlier. In this case it is an unconditional probability representing the uncertainty about the parameter θ . The function $P(\underline{x} | \theta)$ is commonly called the *likelihood* of the dataset \underline{x} (also denoted $L(\underline{x} | \theta)$) and can be interpreted as the probability of observing a certain dataset \underline{x} given that certain characteristics of the population (in this case the parameter θ) are true.

Combining the subjective information and the empirical observations (i.e. the prior and the likelihood), we get the *posterior* probability i.e. the probability that the parameter θ takes a certain value **given** that observations \underline{x} have been made. This is denoted $P(\theta | \underline{x})$. If we sum $P(\theta | \underline{x}) * P(\underline{x})$ over the values of \underline{x} we get the *marginal probability* of the parameter θ .

Treating $P(\underline{x})$ as a constant of proportionality, (3.4) can be expressed as follows:

$$P(\theta | \underline{x}) \propto L(\underline{x} | \theta) * P(\theta) \quad (3.5)$$

Or in words:

$$\text{Posterior} \propto \text{Likelihood} * \text{Prior}$$

Intuitively, we can think of this relationship as representing the Bayesian idea that conjectures about a population is the combination of *presuppositions* (or prior knowledge) about the population and any observations made regarding the population.

3.2.4 Epistemology

The reason why this way of looking at reality is so appealing to our present concern is because information about the OR profile of a company is very often incomplete, derived from quantitative *and* qualitative sources and is continuously updated as new data is gathered. Bayesian statistics offers a practical, yet intuitively appealing, methodology to deal with such a situation.

We will now examine how BNs apply Bayesian statistics to solve a variety of problems. It is with apologies to the reader that the rest of this chapter is rather heavy on terminology.

3.3 Introduction to Graphical Models

3.3.1 Introductory Terminology and the Direct Acyclic Graph

BNs, or variously called belief networks, causal probabilistic networks, directed graphical models or generative models, are a type of graphical model. Graphical models are a combination of probability theory and graph theory. It is the result of converging developments in statistical modelling, engineering and artificial intelligence that began in the 1980's. Initially the extensive calculations in probability theory rendered these efforts unfeasible. However, the utilisation of conditional independencies in graph theory and recent developments in efficient algorithms for propagation of evidence across graphical structures has made this field much more feasible computationally.

To illustrate the usage of graphical models, consider the following diagrams:

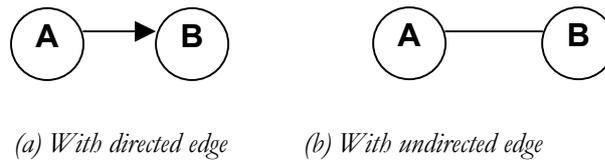


fig 3.1 Nodes in graphical models

A and B are called *nodes* and represent variables A and B. In *fig. 3.1(a)*, the directed edge from A to B implies a causal relationship between A and B. To be more precise, it states that a change in what is known about A (usually affecting the probabilities of events in A) *causes* a change in what is known about B (and thus the probabilities of events in B). This change is usually the result of new information arriving about A (henceforth we will use such statements as “information about A” and “information about the events in A” interchangeably). This new information is sometimes called *evidence*.

When variables are connected in this way, we call variable from which the edge originates the *parent* and the variable to which the edge leads the *child*. When the edge between the nodes are not directed, as illustrated in *fig. 3.1(b)*, then no causation is implied, but rather that some ‘weaker’ form of association (e.g. correlation) exists between A and B. Using the same sort of descriptive language, A and B are called *neighbours*. It follows, finally, that if no edges exist between A and B, then A and B are *independent* i.e. occurrence of events in A has no bearing on occurrence of events in B and vice versa. However, as we shall see, this may depend on whether any *intervening* variables exist between A and B.

Fig 3.1(a) and *3.1(b)* also represents the joint probability of A and B. However, *(a)* and *(b)* expresses this joint probability differently. In *fig. 3.1(a)*, the causal relationship that exists between A and B means that the joint distribution can be expressed as a product of the probability of A and the

probability of B conditional on A or simply written: $P(A)P(B | A)$. This, as we have seen in (3.1), is simply the fundamental rule in pictorial form. Since no such relationship is defined in *fig. 3.1(b)*, this graph only expresses the joint distribution itself: $P(A, B)$.

In general, graphical models comprise of a network of such nodes with edges to connect variables that have some form of relationship, whether of correlation or of causation. BNs on the other hand comprise almost entirely of causal relationships and would thus involve nodes connected by directed edges.

BNs belong to a subset of graphical models that are known as a *directed acyclic graph (DAG)*. DAGs are constructed with relationships such as those in *fig. 3.1(a)* as its basic building block. These building blocks are arranged in such a way that the variables are not cyclical i.e. moving along the edges in the directions implied, it is impossible to return to a previous node. Hence, the term “acyclic”.

Associated with each node of a DAG that has at least one parent is a set of conditional probabilities. These describe the behaviour of the node *conditioned* on all its parents. This is often written as $P(X | pa(X))$ where $pa(X)$ represents the parents of X. For example, the parent set for node D in *fig. 3.2* is the set of nodes {A, B, C}.

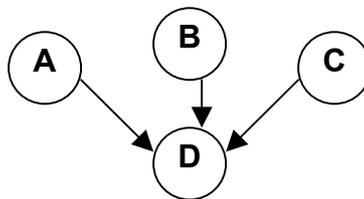


Fig. 3.2 Parent and child nodes

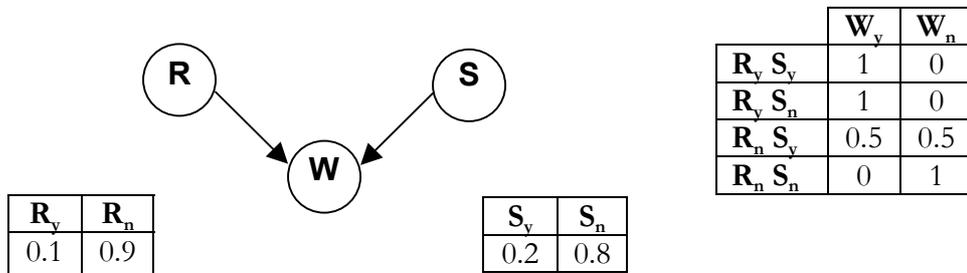
A useful property of a DAG is that it illustrates the assumption that the joint density over a set of variables U (comprising of variables denoted by the vector X) can be expressed as a product of the conditional densities at each node. This is expressed as follows:

$$P(U) = \prod_X P(X | pa(X)) \tag{3.6}$$

This factorization comes from the conditional independence property inherent in the structure of DAGs. However, even this expression can make manipulations of DAGs exceedingly complex, especially if the nodes represent variables with many states.

3.3.2 Wet grass example

We pause briefly at this point to consider a popular example to illustrate some of the concepts presented so far. In this example, grass in a garden is observed in the morning to be either wet or dry. If the grass is wet it could be due to either the sprinkler being on or some rain falling earlier on. The probability of rain is 0.1 whereas the probability that the sprinkler is on is 0.2. The graph and probabilities are shown in *fig. 3*.



The subscripts represent the state of the variable e.g. W_y = wet grass, S_n =sprinkler not on etc.

Fig. 3.3 Sprinkler example.

The probability of the grass being wet *conditional* on the state of the sprinkler and rain is also shown in *fig 3*. Thus, for example, given that neither the sprinkler was on nor was there rain, the probability of the grass being dry is 1. As a first step, we find the joint density of R, W and S. Using the

fundamental rule (3.1) and assuming that R and S are independent (as is represented in the lack of an edge between the respective nodes) we know that:

$$P(W, R, S) = P(W | R, S)P(R)P(S)$$

	W_y	W_n
$R_y S_y$	0.02	0
$R_y S_n$	0.08	0
$R_n S_y$	0.09	0.09
$R_n S_n$	0	0.72
	0.19	0.81

(a) Joint distribution

	W_y	W_n
$R_y S_y$	0.02	0
$R_y S_n$	0.08	0
$R_n S_y$	0.09	0
$R_n S_n$	0	0
	0.19	0

(b) Incorporation of information on wet grass

	W_y
$R_y S_y$	0.11
$R_y S_n$	0.42
$R_n S_y$	0.47
$R_n S_n$	0
	1

(c) Normalisation

	W_y
$R_y S_y$	0
$R_y S_n$	0
$R_n S_y$	0.9
$R_n S_n$	0
	0.9

(b) Incorporation of information on no rain

Tables 3.1 Manipulation of the joint distribution

Thus, multiplying the marginals of R and S into the table of conditional probabilities, we get the resulting joint distribution in Table 3.1(a). By summing up the columns, we get the marginal or unconditional probabilities $P(W)$. We call this *marginalisation*. We can similarly marginalise for the other variables by summing up cells containing the required variable.

Suppose that we received the information that the grass is wet, W_y . We might then be interested to know if this was caused by rain or the sprinkler. We would want $P(R_y|W_y)$ and $P(S_y|W_y)$. Using Bayes' theorem (3.3) we can find these as follows:

$$P(R_y|W_y) = \frac{P(R_y W_y)}{P(W_y)} = \frac{0.1}{0.19} = 0.53$$

Similarly, $P(S_y|W_y) = 0.58$. Another way to arrive at this is to set the probabilities of all events involving W_n to 0 (Table 3.1(b)). Then normalise the remaining probabilities (producing Table 3.1(c)) before extracting the required values by marginalising (e.g. $P(S_y|W_y) = 0.11 + 0.47 = 0.58$). The normalisation process has the effect of setting $P(W_y)$ to 1, which is effectively making the statement that "the grass is wet".

Thus, we see that with the evidence of wet grass the probabilities of both rain and the sprinkler being on have increased substantially, which is what we would expect.

Next, suppose we received additional information that it did not rain earlier. We would set all events involving R_y to 0, producing *Table 3.1(d)*. It is trivial to see that normalization then results in $P(S_y) = 1$. This is intuitively obvious: if the grass is wet and it was not caused by rain, then the sprinkler caused it. It is easy to show that the converse is also true.

What we have seen here is called *conditional dependence*. Before anything is known about W, R and S are independent. However, they cease to be independent once evidence is entered into W. Thus, R and S are conditionally dependent given W. This property exists wherever there are more than one parent nodes to a common child node. It can similarly be shown that in the cases where more than one child nodes share a common parent node, the child nodes are *conditionally independent* given the parent node.

3.4 Inference in Bayesian Networks

3.4.1 The Advent of Efficient Algorithms

Inference in DAGs generally involves the incorporation of evidence entered at the nodes along the graph. The evidence is then propagated along the network to the other nodes by reversing conditional distributions through forming joint distributions (using Bayes' theorem) and establishing marginal posterior distributions from the resulting distributions. The wet grass example above was a simple illustration of this process. However, for a while this approach was still computationally unfeasible since these joint densities can get very large. For example, the wet grass model involved 3 binary nodes and resulted in a joint density with 8 cells. If a model had 5 ternary nodes, the joint

density would increase to 243 cells. In reality, a BN for most applications would be much more extensive than this.

Developments in this area have shown that DAGs can be transformed in such a way that computations can be done ‘locally’ among clusters of nodes, thus involving much smaller joint distributions. Efficient algorithms were devised to facilitate the propagation of the ‘message’ to each cluster across the network. This drastically reduced the amount of computation required and has made inference using BNs fairly user-friendly. In fact, most complex operations using BNs can now be performed on laptop computers using software downloadable from the Internet.

These convenient representations of the DAG are called *junction trees*. These are non-directional graphs consisting of a collection of *maximal* sub-graphs known as *cliques*. Cliques (sometimes called *belief universes*) are groups of nodes where each node in the group is connected to every other node in the group. When a group of nodes are connected this way, we say that they are maximal. Different cliques within a junction tree are connected via *separators*. These are just common nodes shared by two adjacent cliques.

3.4.2 Triangulation

The process of transforming a DAG into a junction tree is called *triangulation*. Although this strictly refers to one of the steps in the process, we will also use it to mean the whole process. Broadly, we can look at this in three stages:

- (i) Moralization
- (ii) Triangulation
- (iii) Specifying the junction tree

We will illustrate these steps using the DAG in *Fig. 3.4(a)*

Moralization

In this step, all parent nodes of a common child node that are not connected are firstly joined with an undirected edge. For example, C, E and F are common parents of G. Thus, edges are added between all the pairs of these three nodes. Then all remaining directed edges are rendered non-directional. The result as shown in *Fig. 3.4(b)* is what is known as a ‘moral graph’ as all parents are ‘married’.

Triangulation

A *cycle* is a sequence of nodes connected by edges that start and end at the same node. A cycle of *length n* consists of a sequence of n consecutive edges. For example, the cycle C, D, H, G, C in *Fig. 3.4(b)* has length 4. *Triangulation* is the process of adding undirected edges such that any cycle that has length over 3 possesses a *chord*, where a chord is defined as an undirected edge joining two non-consecutive nodes in the cycle in question. Thus, for example, an edge has to be added between C and H or D and G. Suppose we choose to add an edge between D and G, we get the *triangulated* graph in *Fig. 3.4(c)*.

Specifying the Junction Tree

Once a triangulated graph is obtained, a junction tree can be specified. This involves identifying the cliques within the graph and the separators that connect them. Thus, for example, the cliques (A,B,C) and (C,D,G) have separator C, whereas the cliques (C,D,G) and (D,G,H) have as a separator the clique (D,G). A junction tree is usually shown as in *Fig. 3.5*.

When the graph is expressed in this form, the evidence can be incorporated locally at the cliques where the calculations involve fewer dimensions. Information from the updated cliques is then propagated in step to the other cliques of the graph via the separators. Thus, there is no need to deal with large joint distributions of the whole graph. This is possible because there is a series of

common nodes shared by adjacent cliques running throughout the whole tree – one can picture this to act like a series of interlocking chains. This property is called the *running intersection property*.

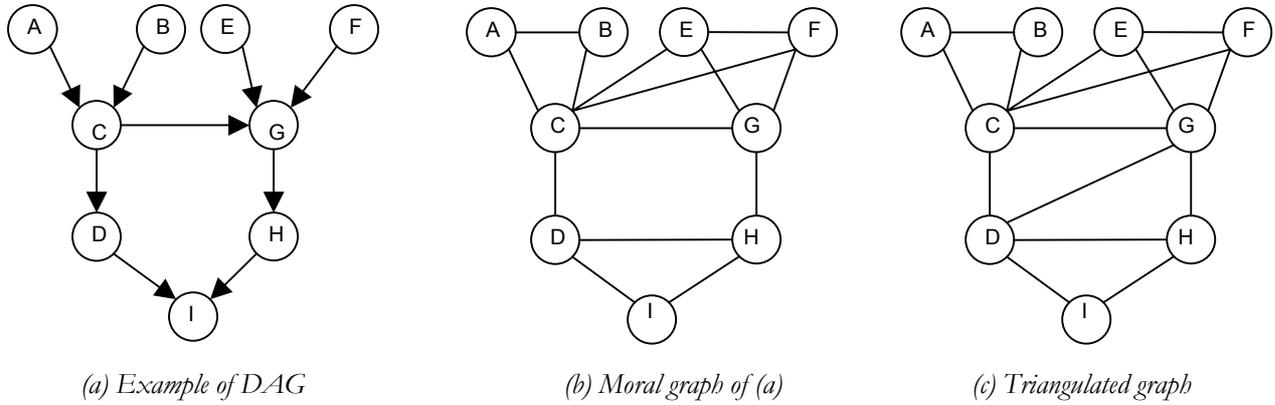


Fig 3.4 Triangulation

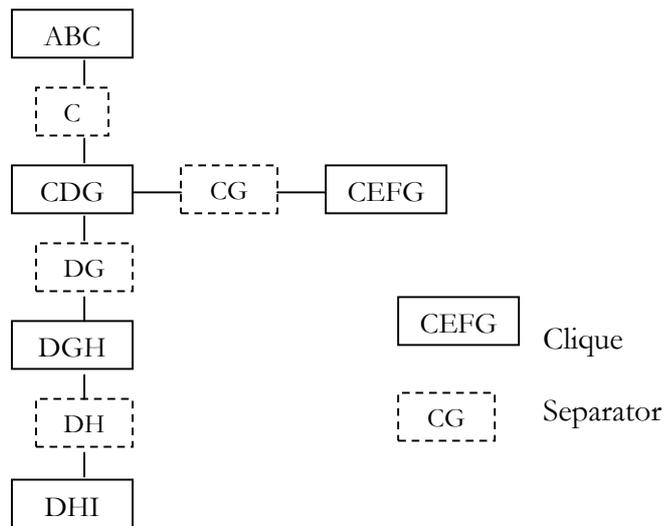


Fig 3.5 Junction tree

3.4.3 Inference along the junction tree

To facilitate the illustration of the steps involved in inference along the junction tree, we will use the example in the previous section where the junction tree is as shown in *fig. 3.5*. For simplicity, we assume that all the nodes are binary variables, with the states (Yes, No) denoted with subscripts (e.g. D_y for the state $D=Yes$). The main steps to be taken for inference using the junction tree are:

- (i) Defining prior distributions;
- (ii) Initializing the tree; and
- (iii) Two-phased propagation.

(i) Defining prior distributions

The first thing that needs to be obtained would be the prior distributions. This would be the unconditional prior distribution for nodes without parents and conditional prior distributions for child nodes – similar to the example in *fig 3.3*. For our example, we would need unconditional priors $P(A)$, $P(B)$, $P(E)$ and $P(F)$ and conditional priors $P(C|A,B)$, $P(G|C,E,F)$, $P(D|C)$, $P(H|G)$ and $P(I|D,H)$. For each prior, we would need the probabilities for each configuration of the combination of states of variables involved. Thus, for $P(C|A,B)$ we would need figures for $P(C_y | A_y, B_y)$, $P(C_n | A_y, B_y)$, $P(C_y | A_n, B_y)$ etc.

These can be determined by:

- (a) *Subjective opinion of the expert.*

Experts are interviewed in a series of questionnaires to arrive at quantified conclusions of the probabilities. Of course, sufficient confidence in the accuracy of the expert's advice is a prerequisite to use this method.

- (b) *Maximum likelihood estimation.* For conditional priors this method would entail taking a ratio of frequency of the event to the frequency of the parent configuration.

For example:
$$P(C_y | A_y, B_y) = \frac{n(C_y, A_y, B_y)}{n(A_y, B_y)}$$

For unconditional priors, this would simply be the proportion of occurrence between the various states of the variable.

The method of maximum likelihood assumes that past data is relevant and complete. In practice, there will not be a clear dichotomy between these two methods as the experts will also rely on past data but tempered with experience and knowledge regarding its applicability for future events. §3.5 below will discuss updating the priors in the light of new data.

(ii) Initializing the tree

Next, we need to *initialize* the tree. For each clique and separator, a table exists with cells that correspond to each combination of states of the variables in the clique or separator. For example, the clique (A,B,C) will have 8 cells for combinations $A_y B_y C_y$, $A_n B_y C_y$, $A_n B_n C_y$ etc. The numerical values in these cells are called *potentials*. These will change as information is passed along the junction tree and as each clique or separator is updated. We will see how this works below.

To initialize, all potentials are set to unity. Priors that factorize as per (3.6) are then multiplied into cliques that contain the variables in the factorized set of priors. For example, the potentials in the clique (A,B,C) will be multiplied by the values in the joint distribution $P(A,B,C) = P(C|A,B) * P(A) * P(B)$. This process is analogous to the steps taken to arrive at *table 3.1(a)*.

Some cliques may have more variables than any available set of factorized priors. In these cases, the priors are still multiplied into the cliques, so long as the variables in the priors are a subset of the clique variables. Cliques for which this operation is not possible are left with all potentials as unity. The potentials in the separators also remain as unity. Once this has been done, then the tree can be used for *propagation*.

(iii) Two-phased propagation

Propagation involves a messages originating from the *leaves* of the tree (i.e. the cliques at the extremities of the tree) and passed along the tree to the *root clique* and then back out again to the leaves (the choice of the root clique is arbitrary). Along the way the potentials in each clique will be updated.

Message passing between two cliques takes place as follows (we call the sending and receiving cliques *source* and *sink* respectively):

1. *Marginalisation at the source clique.* Sum up cells in the clique for common configurations of the variables of interest – these are the ones in the separator. For example, (A,B,C) is connected to (C,D,G) via separator C. Thus, we marginalize (A,B,C) for variable C by summing up all the cells with configurations containing C_y (e.g. $A_y B_y C_y$, $A_n B_y C_y$, etc) to obtain the sum margin C_y and similarly for C_n . This is analogous to how $P(W_y)$ and $P(W_n)$ were extracted from *table 3.1(a)*.
2. *Update separator node.* The separator node is updated with the sum margins obtained by marginalizing the source clique. This becomes the updated potential of the separator.
3. *Update sink clique.* The clique receiving the message is updated using *update ratio*. This is the ratio of the potentials of the separator after the message is passed to the potentials before the message is passed (if this is the collect phase, which takes place right after the initialization, the separator potentials would be unity. Hence the update ratio would just be the new potentials). The updating is performed by multiplying the update ratio into the relevant cells in the sink clique. These are the ones with variables corresponding to the separator variables.

This step is performed sequentially: the leaves are the first to be sum-marginalized, then each adjacent clique will be updated in turn until the root clique is updated. Cliques due to receive updates from more than one source only send out a message once that has occurred. This flow towards the root clique is called the *collect phase*.

Then the whole process is repeated but with the messages going out from the root clique to the leaves. This flow outwards from the root clique is called the *distribute phase*. Hence, this process is sometimes called *two-phased propagation*. Once both phases are complete, the tree has achieved equilibrium. At this stage, the potentials at each clique need to be normalised – i.e. scaled so that all the potentials add up to 1. They can then be marginalised to obtain the posterior marginal distributions for every node variable. This is the desired result.

3.4.4 Entering evidence into the junction tree

The scope of this investigation covers only evidence of *certain* events (e.g. “the grass is wet”). Junction trees can incorporate evidence very efficiently. For this purpose, the two-phased propagation method is also utilised. The one difference is that the evidence is incorporated at the initialization stage. Evidence would involve that a node variable be known for certain to assume one of the possible states (e.g. $A=Yes$ or $P(A_y)=1$), implying certain knowledge that the variable has *not* assumed any of the other states. This is incorporated into the initial clique potentials by setting all potentials involving the other states of the same variable to be equal to 0. In this example, we set the potentials for $A_n B_y C_y$, $A_n B_y C_n$, $A_n B_n C_y$ and $A_n B_n C_n$ to 0. The other initial potentials remain as they were.

Once this is done, the rest of the propagation then carries on as before. The marginal probabilities obtained now will reflect the impact of the evidence entered.

3.4.5 Simulation

A useful application of the junction tree is to generate random samples of configurations for the various probability distributions in the DAG. There are two methods of sampling:

(i) *Probabilistic logic sampling*

This method samples directly from the DAG itself. Samples are first taken from the unconditional nodes based on the marginal probability distributions at these nodes. Then, the child nodes of these nodes are sampled based on the conditional probability distributions at these nodes *conditional* on the outcome of the sampled configuration of their respective parents. For example, A and B are first sampled from $P(A)$ and $P(B)$ respectively. Supposing the outcome is A_y and B_n . Then for node C, we sample from $P(C | A_y B_n)$. This continues until all nodes have been sampled. This will be counted as *one* sample for the whole DAG. This is carried out for as many sample as required. The frequencies at each node can then be used to express the sampled marginal probabilities.

If evidence is entered, the process is carried out as before (using the prior probabilities). However, samples that include values that are not equal to the evidence entered (at the specific nodes) are rejected. This results in much redundancy in the sampling process.

(ii) *Sampling from the junction tree*

A better way to sample with evidence is to exploit the junction tree. The process is similar to two-phased propagation. The junction tree is first initialized with the evidence incorporated. Then the collect phase is carried out. When the root clique has been updated, a configuration is sampled on the root clique. This sample is then immediately entered into the root clique as evidence. The distribute phase is then carried out but with the modification of this step of sampling immediately to enter evidence at the clique

before passing the message on to the adjacent cliques. The algorithm continues in this way until the whole tree has a sampled configuration at each clique.

This method is much more efficient because none of the samples are rejected as compared to the probabilistic logic sampling method. This is due to the fact that evidence is already incorporated into the junction tree, thus the sampling process will not result in samples with values that are at odds with the evidence.

3.5 Updating the BN in the light of fresh data

The method for prior specification so far has been rather simplistic, not taking into account that new data from time to time might be useful in updating the probabilities in the BN. In this section, we will explore how probabilities derived in §3.4.3(i) can be updated in a Bayesian fashion as new data arrives. The methods described are applicable for complete data only. Incomplete data will require additional methods which will not be discussed here.

What we will see here is that the expert's opinion can be quantified in an intuitive way yet mathematically tractable in the form of a Dirichlet distribution. This distribution allows for convenient updating using Bayesian inference as new data arrives.

3.5.1 Dirichlet distribution as conjugate prior of a multinomial likelihood

Recall from (3.5) that Bayesian inference can be expressed in the following way:

$$P(\theta | \underline{x}) \propto L(\underline{x} | \theta) * P(\theta)$$

or:

$$\text{Posterior} \propto \text{Likelihood} * \text{Prior}$$

Suppose variable X in a BN that can take on the values of x_1, x_2, \dots, x_k with probabilities $\theta_1, \theta_2, \dots, \theta_k$. The likelihood for a sequence of n independent samples, where n_1 samples have the value x_1 , n_2 samples have the value x_2 , etc, can be represented by the multinomial distribution (of which the binomial distribution is a two-parameter case) :

$$p(n_1, n_2, \dots, n_k | \theta) = \frac{n!}{\prod_{i=1}^k n_i!} \prod_{i=1}^k \theta_i^{n_i} \quad (3.7)$$

A convenient candidate to be used as the conjugate prior for the θ_i 's is the Dirichlet distribution:

$D(\alpha_1, \dots, \alpha_k)$ with the density:

$$p(\theta) = \frac{\Gamma(\alpha_+)}{\prod_{i=1}^k \Gamma(\alpha_i)} \prod_{i=1}^k \theta_i^{\alpha_i - 1} \quad (3.8)$$

where $\alpha_+ = \sum_i \alpha_i$, also known as the *precision*. The Beta distribution is actually a two-parameter specific case of the Dirichlet distribution. Since binary nodes (e.g. Yes/No variables) are quite common in BNs the Beta distribution is often used.

The parameters of the Dirichlet distribution, $(\alpha_1, \dots, \alpha_k)$, have a rather intuitive interpretation: they can be seen as the implied relative sample sizes for the θ_i 's. In practice, the Dirichlet distribution would usually not be specified precisely by the expert. Rather, the expert's conjecture about the θ_i 's (e.g. expected value and range) can be used to work backwards to arrive at a Dirichlet distribution since we know helpful statistics such as:

$$E(\theta_j) = \frac{\alpha_j}{\alpha_+} \quad \text{and} \quad \text{Var}(\theta_j) = \frac{\alpha_j(\alpha_+ - \alpha_j)}{\alpha_+^2(\alpha_+ + 1)} \quad (3.9)$$

from which the α_i 's can be derived.

The posterior analysis is fairly straightforward when using the Dirichlet prior. The posterior itself is a Dirichlet with its functional form arrived at by multiplying (3.7) and (3.8):

$$p(\theta|n_1, n_2, \dots, n_k) \propto \prod_{i=1}^k \theta_i^{\alpha_i + n_i - 1} \quad (3.10)$$

This is a $D(\alpha_1 + n_1, \dots, \alpha_k + n_k)$. Thus, the posterior is a Dirichlet distribution with the new parameters, and hence the new implied relative sample sizes, equal to the prior implied relative sample sizes *adjusted* by the obtained sample counts.

3.5.2 Specifying the Dirichlet parameters

Suppose we are considering the variable $P(C|A, B)$. C is a binary variable, thus the prior distribution will be $B(\alpha_1, \alpha_2)$. Its parameters can be specified in the following way. We first gather the expert's opinion on the mean and standard deviation of the probabilities of the values in the variable. Some experts who are not familiar with statistical concepts can be asked to quote a best estimate and the range of most likely values. The best estimate can then be taken as the mean and the range can be taken to encompass two standard deviations about the mean, from which the variance can be easily obtained. The results can be stored as shown in Table 3.2.

Parent	C_y		C_n	
	$E(\theta_y)$	$Var(\theta_y)$	$E(\theta_n)$	$Var(\theta_n)$
A_y, B_y				
A_n, B_y				
A_y, B_n				
A_n, B_n				

Table 3.2 Table for recording mean and variance of probability values from prior elicitation exercise

Then, for each parent configuration we have, for a two-parameter case of (3.9), the following pairs of simultaneous equations:

$$C_y: \quad E(\theta_y) = \frac{\alpha_y}{\alpha_y + \alpha_n} \quad \text{and} \quad Var(\theta_y) = \frac{\alpha_y \alpha_n}{(\alpha_1 + \alpha_2)^2 (\alpha_y + \alpha_n + 1)}$$

$$C_n: \quad E(\theta_n) = \frac{\alpha_n}{\alpha_y + \alpha_n} \quad \text{and} \quad Var(\theta_n) = \frac{\alpha_y \alpha_n}{(\alpha_1 + \alpha_2)^2 (\alpha_y + \alpha_n + 1)}$$

These can then be solved for α_y and α_n . There will be two sets of parameters (α_1, α_2) to choose from. In general, it is more conservative to choose the one that gives a *lower* precision (and thus a higher variance).

3.5.3 Bayesian updating with new data

This can be done in a fairly straightforward manner. As shown in the posterior analysis that resulted in (3.10), all that is required to arrive at the posterior Dirichlet distribution is to add the data counts to the relevant prior parameter to obtain the posterior parameter. Supposing the prior distribution for $P(C|A_y, B_y)$ is $B(\alpha_y, \alpha_n)$ and the following set of data were obtained (only the relevant variables are shown):

$$\{(A_y, B_y, C_y \dots), (A_y, B_y, C_n \dots), (A_y, B_y, C_n \dots), (A_y, B_y, C_n \dots), (A_y, B_y, C_y \dots)\}$$

It then follows that the posterior is $B(\alpha_y + 2, \alpha_n + 3)$. It is worth noting that the same result is arrived at whether the data is updated sequentially or in batches.

The updated probability estimates can be obtained by taking the mean from the $B(\alpha_y + 2, \alpha_n + 3)$ distribution. Thus, the posterior for $P(C_y|A_y, B_y)$ is $\alpha_y/(\alpha_y + \alpha_n + 5)$ whereas the posterior for $P(C_n|A_y, B_y)$ is $\alpha_n/(\alpha_y + \alpha_n + 5)$.

This whole process of Bayesian updating is sometimes also described as *learning*. Thus, when a BN is carried out *without* learning, then the probabilities remain unchanged with data.

3.6 Usage of monitors in model assessment

Perhaps of paramount importance, the resulting model needs to be compared against actual data to verify that it adequately corresponds to reality and to assess its usefulness as a predictive tool. Very often there will be two or more alternatives with regards to model structures or the sets of probabilities. In such cases, there will also be a need to distinguish the better model with regards to these two criteria.

3.6.1 Logarithmic score

The main indicator used is the *logarithmic score (LS)*:

$$S_m = -\log p_m(y_m)$$

Where $p_m(\cdot)$ is the predictive distribution for the event, Y_m , after $m-1$ occurrences of events. If learning is allowed then $p_m(\cdot)$ incorporates all updates resulting from the $m-1$ events. The *LS* is the negative log of the probability of the event in the actual outcome y_m .

Actually, what the *LS* does is indicate the level of ‘surprise’ resulting from the actual outcome. For example, if the event which occurred was predicted to happen with a probability of 0.1 the *LS* would be $-\log(0.1) = 2.3$. Conversely, an event expected to occur with probability of 0.9 would carry a score of 0.1. Thus, the less likely an event is predicted to happen, the more ‘surprising’ it is if it did happen.

For a series of M events, the total penalty is $S = \sum_{m=1}^M S_m$. If the probabilities $p_m(\cdot)$ were updated with each subsequent event, then the S is invariant to the order in which the y_m 's occur. A lower penalty is always more desirable.

3.6.2 Model assessment

Now we need some criteria by which to accept or reject a model. The total penalty incurred as a single figure is quite arbitrary on its own. It needs to be compared to a standard. This can be done in two ways:

(i) *Relative standardisation*

In this method the model defined by the expert is tested against a reference model, which is a pre-defined benchmark by which the model is assessed. The total penalty incurred by the model defined by the expert, S , is compared against the penalty incurred (using the same dataset) by a reference model, S_{ref} . The model is rejected if S exceeds S_{ref} . The degree to which the expert model is preferred over the reference model is indicated by:

$$\exp(S_{ref} - S) = \frac{P(\text{data} | \text{expert's prior})}{P(\text{data} | \text{reference prior})}$$

This is also known as the *Bayes' Factor* in favour of the expert's model.

(ii) *Absolute standardisation*

In this method, a test statistic is compiled using the penalties and tested against a null hypothesis (H_0) that the data fits the probabilities assumed in the BN. We define the expectation and variance of the penalty at each update:

$$E_m = -\sum_{k=1}^K p_m(d_k) \log p_m(d_k) \quad (3.11)$$

$$Var_m = \sum_{k=1}^K p_m(d_k) \log^2 p_m(d_k) - (E_m)^2 \quad (3.12)$$

where the d_k 's are all the possible states of the node(s) considered. E_m could be seen as a probability-weighted average of the logarithmic scores.

The sums of the *actual* penalties incurred by the data, S , together with the expectation and variance at each update are used to compute the following test statistic:

$$\frac{S - \sum_{m=1}^M E_m}{\sqrt{\sum_{m=1}^M Var_m}} \quad (3.13)$$

Under H_0 , the actual penalties are from same underlying distribution as the expected penalties. Thus, this test statistic will have the standard Normal distribution and for values outside the range $[-1.96, 1.96]$, H_0 will be rejected at the 5% significance level.

3.6.3 Parent-child monitors

There exists a set of monitors that can be used to diagnose the validity of BN's in the light of data. In general, these monitors are used to obtain the LS for each piece of data from which the total penalties can be compiled and tested according to the two criteria laid out above. The one thing that differentiates one type of monitor from another is the form of the LS . This allows various aspects of the BN to be diagnosed.

The first of these is known as the Parent-Child Monitor. For each parent-child set of nodes, we have the following LS :

$$-\log p_m(x_k | X_{pa(k)} = \rho)$$

where $p_m(\cdot | X_{pa(k)} = \rho)$ is the probability distribution of the child node k for the parent configuration $X_{pa(k)} = \rho$ after $m-1$ cases of complete data have arrived and x_k is the state of the k node on the m th state. This monitor measures how well the conditional probabilities in the BN that link parent and child nodes predict the outcome of the child nodes.

For example, if we are monitoring the parent-child set expressed in the conditional probability $P(C | A_y, B_y)$ then as the first piece of data arrives (say, A_y, B_y, C_y) the LS can be determined as

$S_1 = -\log P_1(C_y | A_y, B_y)$ and this is obtained from the expert prior distribution, $B(\alpha_y, \alpha_n)$. The data is then incorporated into the conditional probability as illustrated in §3.4.3 to obtain the posterior $P_2(C | A_y, B_y)$ which has the distribution $B(\alpha_y + 1, \alpha_n)$. The process then continues iteratively as each subsequent case arrives.

The total penalty, S , can then be compared with the total penalty from various alternative reference models, S_{ref} . Some examples of reference models could be using the same expert prior **without** learning or using a reference prior $B(0.5, 0.5)$ **with** learning.

For absolute standardisation, the penalties incurred would be the same. Using (3.11) and (3.12), the expectation and variance can be calculated for the first data set as follows:

$$E_1 = -P_1(C_y | A_y, B_y) \log P_1(C_y | A_y, B_y) - P_1(C_n | A_y, B_y) \log P_1(C_n | A_y, B_y)$$

$$Var_1 = P_1(C_y | A_y, B_y) [\log P_1(C_y | A_y, B_y)]^2 + P_1(C_n | A_y, B_y) [\log P_1(C_n | A_y, B_y)]^2 - E_1^2$$

This is similarly performed for all subsequent updates. The test statistic can then be found using (3.13). If the absolute value is under 1.96 then the child node is being correctly predicted by the conditional probability.

It is often useful to plot the cumulative values of the penalty (for relative standardisation) or cumulative values of the test statistic (for absolute standardisation) against the data. The graph can be used to compare the different alternative models and their paths give an indication of how well/soon the models adapt to the data. The preferred model is the one where the cumulative penalty increases very little with new data and the cumulative test statistic centers about zero.

3.6.4 Node monitors

Node monitors measure what is happening at the level of the individual nodes. There are two types of node monitors: *unconditional node monitors* and *conditional node monitors*. Unconditional node monitors detect poorly estimated marginal distributions. The *LS* is:

$$-\log p_m(x_v)$$

where $p_m(\cdot)$ is the marginal distribution of the states of the node X_v after $m-1$ cases. The score expresses the ‘surprise’ at obtaining $X_v = x_v$ on the m th case. Conditional node monitors can be used to detect poor structure, usually upon obtaining reasonable results for the unconditional node monitors first. The *LS* is:

$$-\log p_m(x_v | \varepsilon_m \setminus X_v)$$

where $p_m(\cdot | \varepsilon_m \setminus X_v)$ is the probability distribution of the node X_v after $m-1$ cases of evidence have been incorporated and the latest set of evidence ε_m has just been propagated throughout the BN *except* the evidence at node X_v . This score then measures the ‘surprise’ at obtaining $X_v = x_v$ *after* all other nodes have been propagated with evidence.

The tests can then be performed in a similar manner.

3.6.5 Global monitors

The global monitor of a BN measures the *LS* of the total evidence entered for the m th case after $m-1$ cases have been entered. This is:

$$-\log p_m(\varepsilon_m)$$

The overall global monitor is just a sum of the *LS* for all the cases:

$$G = \sum_{m=1}^M -\log p_m(\varepsilon_m)$$

The global monitors of two competing models, G^1 and G^2 are then compared, with the preferred model having the lower value. The Bayes' factor in favour of model 2 is $\exp(G^1 - G^2)$.

3.7 Chapter Conclusion

In this chapter, we have examined some theoretical foundations describing the construction of a BN, its usage for inference and ways to assess the validity of the resultant model. We shall see in the following chapter how all this theory can be put to practice in an OR example. This example will also serve to clarify the methods explained in this chapter.

Chapter Four

Applying Bayesian Networks to the Modelling of Operational Risk

4.1 Chapter Introduction

In this chapter we examine how BNs might be applied in practice for OR modeling. This will be illustrated using a hypothetical example of the risks involved in setting up an Internet on-line business. A fair amount of detail has been provided within the text to facilitate the clarification of concepts introduced in the previous chapter. More complete numerical information underlying the illustrations can be found in the Appendices. Reference to relevant sections in chapter 3 will be made at various points to allow quick recall of the theory being illustrated.

4.2 Network Risk

The example we are considering is that of a fictitious medium sized insurance company (BayeSure Insurance Co. or BSI) that decided to set up an on-line business (called BSNet) so that customers may purchase via the Internet more basic products such as travel insurance, personal accident and even some cleverly designed health insurance policies. The company has invested a significant amount of capital to set up the on-line network infrastructure, involving fancy widgets such as high-end servers and firewalls. However, as with any on-line business in a high-volume, high-value and highly competitive industry, the management is concerned that they have the adequate financial resources to deal with contingencies that may damage the business i.e. Operational Risks. We shall call the aspect of OR that results from operating this on-line business *Network Risk (NR)*.

After a year of operating BSNet, the management decided to put the I.T. managers, marketing personnel and various BSNet user departments together with some statisticians from Risk Management Unit to set up a BN to serve a two-pronged strategy:

- (i) Identify causation of NR events such as transaction downtime, server downtime and application failure; and
- (ii) Help the management to decide how much risk capital to allocate to cover NR in all but the most extreme of scenarios.

Throughout this chapter, we will view samples of output from a program for BN inference called XBaies¹² which was developed by Robert Cowell of City University.

4.3 Defining the Structure

After some deliberation the working committee was able to put together a structure that looks like *fig 4.1*¹³. *Table 4.1* is a list of the variables employed in this model. The set of values for each variable are also given along with the abbreviations to be used throughout this chapter. We now briefly introduce the causal dependencies in the BN and how it relates to the business environment.

The three main causes of NR loss are transaction downtime (TD), data loss (DL) and server downtime (SD). These can have varying degrees of severity but can be grouped into half a day or a full day to get a server or network up and running, or in the case of data loss, either 50% or complete data loss. Based on data gathered in the first year of operation with some additional expert

¹² The latest version Xbaies is still in development. However, a very similar version 2.0 is available at his website <http://www.staff.city.ac.uk/~rgc>. Various computer soft wares are available to perform the operations and calculations described in this dissertation (e.g HUGIN, Netica, XBaies). The steps involved in triangulation and updating of the cliques are usually automatically done in the program and not seen in the display. The user need not worry about the underlying calculations but instead sees only the marginal densities of each node – and how they change as evidence is entered, as we shall see below.

¹³ *Fig. 4.1* and all figures of the BN in this chapter are screen dumps from Xbaies 3.0.

input by the marketing department the total cost as a result of various combinations of such events were estimated. This would include, for example, the cost of repairs, lost business opportunities during the downtime, wages paid to idle staff, costs to recover loss data and damage to reputation (perhaps quantified as loss of future business). All this is combined into one variable, Cost, representing the bottom-line effect.

Meanwhile, TD is caused by network failure (NF) but there is a mitigating factor: whether or not a high-availability network (HAN) was employed in the running of BSNet. SD is caused by server failure (SF) which in turn is due to power surges (PS) and the server quality (SQ). Again there are mitigating factors here: the availability of uninterrupted power supply (UPS) – perhaps an internal generator, and the usage of a high availability server (HAS).

No.	Description	Values	Abbreviation
1.	Application Failure	Application corruption, Lockup, OK(No failure)	AF
2.	Cost of losses from network risk (Cost)	0.0m, 0.5m, 1.0m, 1.5m, 2.0m, 2.5m	Cost
3.	Data Loss	0%, 50%, 100%	DL
4.	End User Modification	Yes, No	EUM
5.	Firewall	Application Proxy, Packet Filter	F
6.	File Access Control	High, Low	FAC
7.	High Availability Network	Yes, No	HAN
8.	High Availability Server	Yes, No	HAS
9.	Hacker Attack	Yes, No	Hack
10.	Network Failure	Yes, No	NF
11.	Power Surge	Yes, No	PS
12.	Server Downtime	0 day, 0.5 day, 1 day	SD
13.	Server Failure	Yes, No	SF
14.	Server Hardware Quality	High, Low	SQ
15.	Transaction Downtime	0 day, 0.5 day, 1 day	TD
16.	Uninterrupted Power Supply	Yes, No	UPS
17.	Virus Attack	Yes, No	V

Table 4.1 Values and abbreviations for Network Risk model

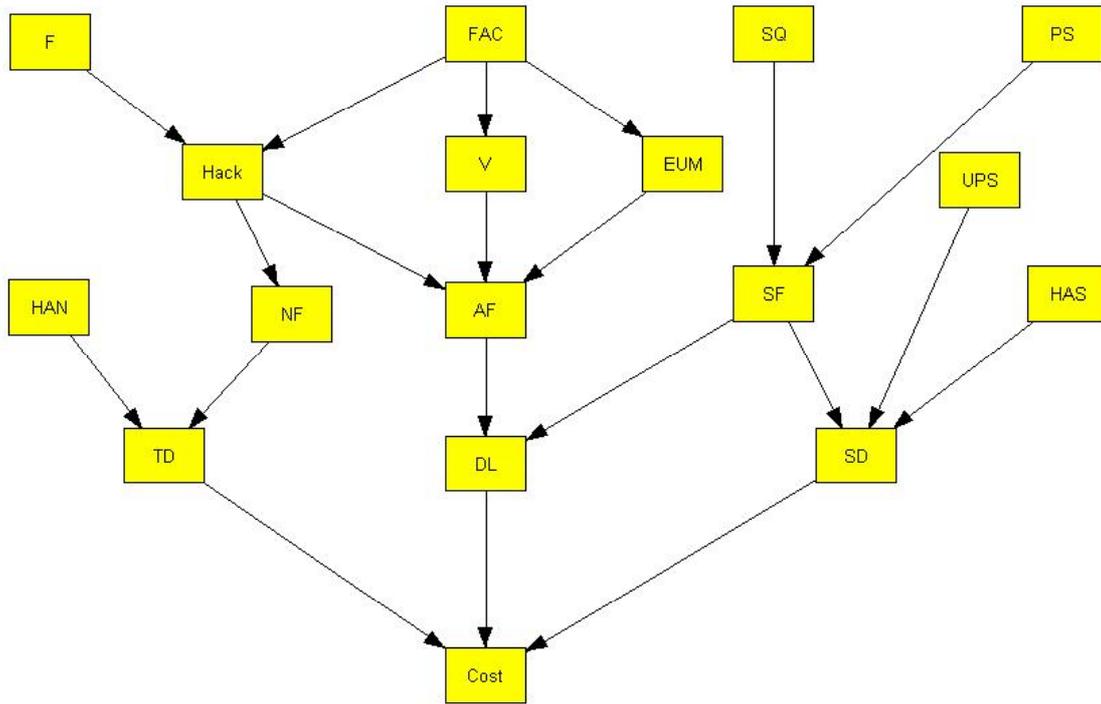


fig. 4.1 DAG of Network Risk example

DL can be caused by either SF or application failure (AF). AF has three main causes: (i) Modifications made by end users in BSI to the system, either intentionally or accidentally (EUM); (ii) Virus attacks (V); and (iii) Malicious hacking by external parties (Hack). These three events are largely controlled by the level of file access given to various parties (FAC). The type of firewall (F) used will also affect the ease with which malicious hackers can access the system. Malicious hacking not only causes AF but NF as well.

This model has been defined to reflect a ‘holding period’ of one week. Thus, for example, the marginal distribution for the variable cost would be the expected probability distribution of NR costs incurred by the company over the period of one week. The choice of holding period is arbitrary as far as the mechanics of the model is concerned. The only consideration in this case

would be the availability of data. As the system has commenced for only a year, it makes more sense to use 52 sets of weekly data than one set of annual data.

4.4 Setting up the BN

4.4.1 Prior specification

At this point the DAG will need to be populated with the prior probabilities at each node. The main methods of prior specification have been described in §3.4.3 and we assume that BSI uses these methods to elicit the priors from the experts mentioned above and from past data.

For example, since attack by computer viruses on online networks are a relatively common and well documented occurrence, sufficient data might exist for the conditional probabilities for the V|FAC node to be obtained using maximum likelihood estimation. Thus, having combined industry experience with one year's worth of BSNet experience, it might have been observed that only in 10% of cases where a high level of File Access Control is implemented was there a virus attack.

For events like end user modification, the circumstances of different companies (e.g. training and recruitment policies) might be so different as to render any external data irrelevant. In this case, the I.T. Department together with staff appraisal information from Human Resources Department would need to decide on a figure for the probability of adverse end user modification for systems with high/low levels of file access control. This corresponds with the EUM|FAC node.

Having performed these exercises, the company arrives at the probabilities detailed in Appendix I. For the purpose of this illustration, the figures have been simplified and as such, might not be realistic – with apologies to readers well versed in network risks.

For flexibility in the model, various unconditional priors have been assigned neutral probabilities. These will be treated as input variables to reflect the actual state of BSNNet once ‘evidence’ as been entered. Power surge, obviously, is outside of the control of the company. Thus past data has been used to arrive at the probabilities and it will not be treated as an input variable.

Theoretically, any node can be treated as an input variable. To prevent confusion, when evidence is entered in the other nodes it will be treated as a stress or scenario test.

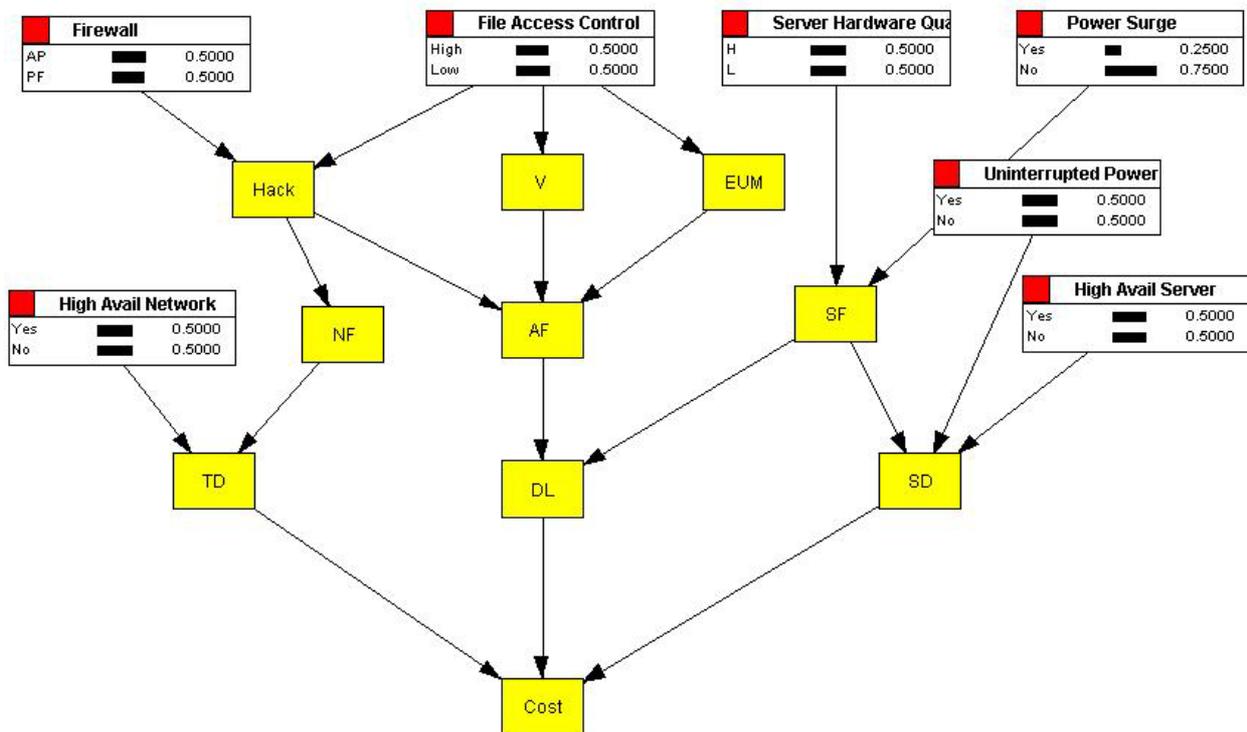


fig. 4.2 Unconditional priors

4.4.2 Junction Tree Specification

Before we go on to use the DAG for inference, we need to obtain a junction tree via the process of *triangulation* as described in §3.4.2. We begin with the moralization of the DAG by establishing moral links between parent nodes. The result is shown in fig. 4.3.

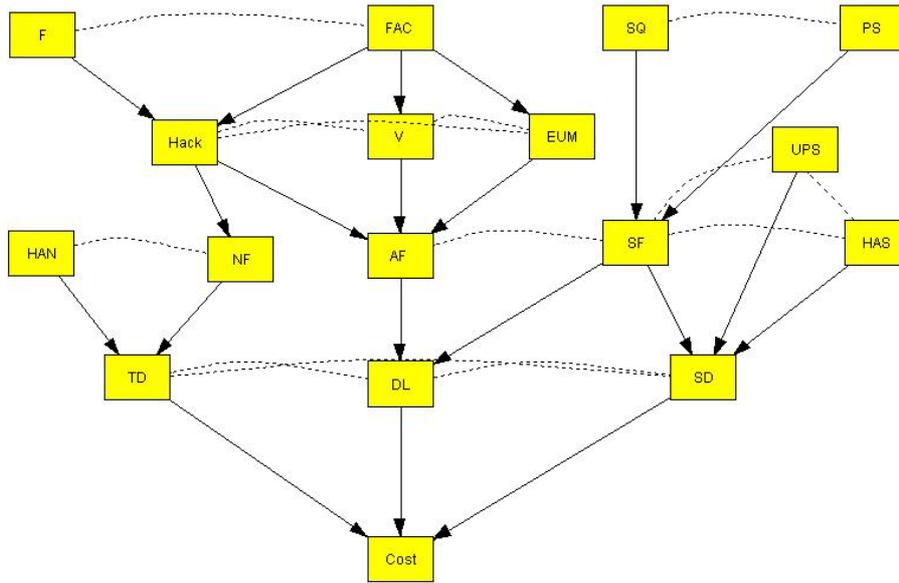


fig 4.3 Moral links

Then, replacing the directed edges with undirected edges, we get the moral graph for this model.

Additional edges are added between the following pairs of nodes: (Hack, TD), (Hack, EUM), (AF,TD), and (SF,TD). We now have the triangulated graph as shown in fig 4.4.

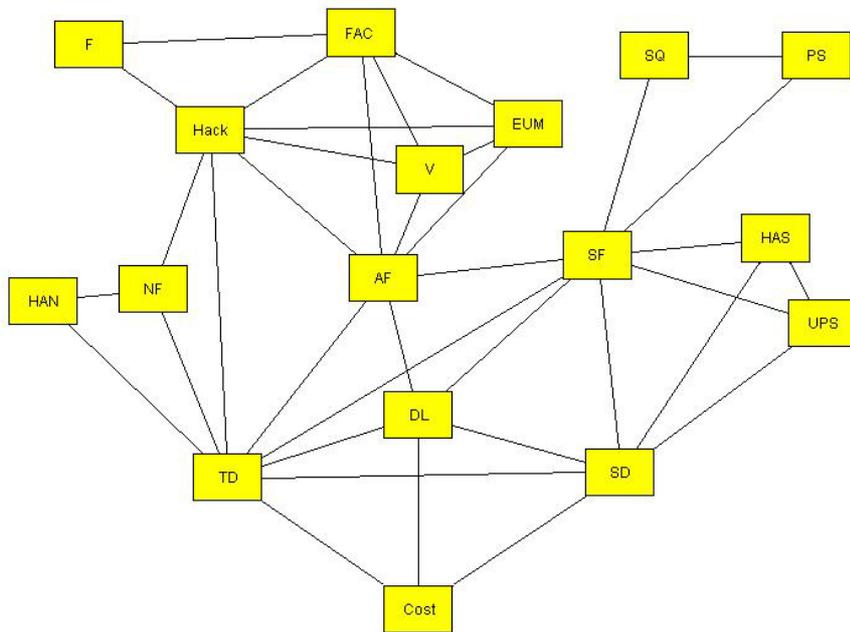


fig 4.4 Triangulated graph

We can now specify a junction tree from the cliques identified in the triangulated graph. Junction trees are not unique, though it is useful to avoid cliques that are too large. The one chosen for this example is shown in *fig 4.5*. The boxes indicate the cliques and give the members of the clique along with the clique size. The root clique is denoted by the box with a bold outline.

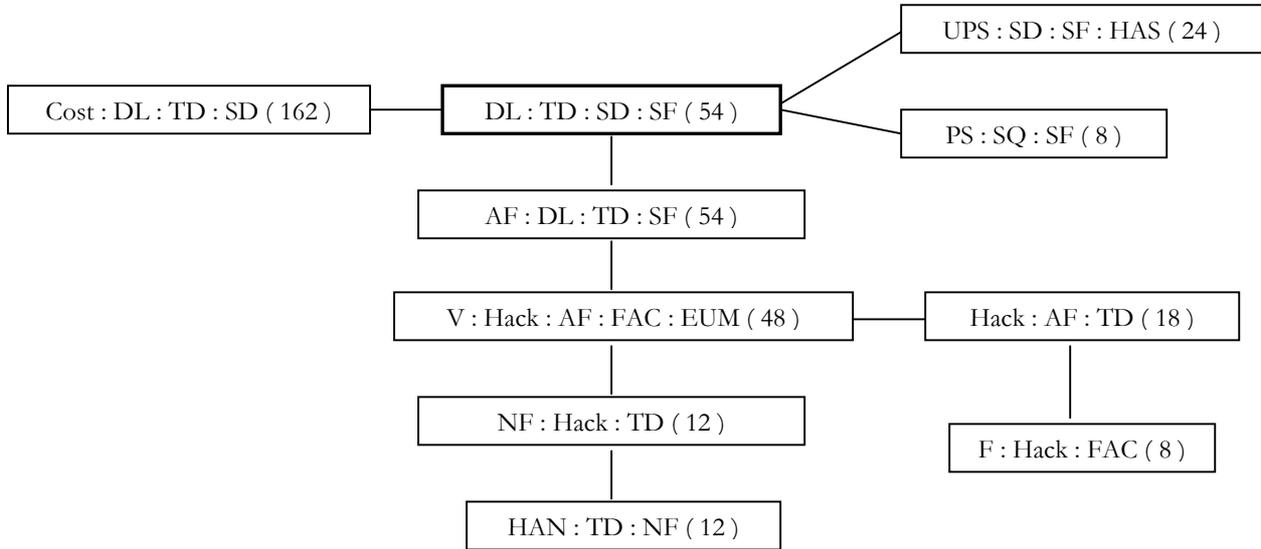


Fig 4.5 Junction tree

Inference

Having specified the junction tree and the prior distributions, we are now ready to make some inference about the posterior distributions. Firstly, we need to initialize the junction tree. As stated in §3.4.3, this entails setting all the clique potentials to 1. Then, combinations of priors that factorise as in (3.6) are multiplied into cliques that contain the corresponding set of variables. Thus, for example, the potentials in the clique (V, Hack, AF, FAC, EUM) will be multiplied by the probabilities in $P(AF|V, EUM, Hack)$ and $P(V|FAC)$.

To illustrate, consider the potential for the cell corresponding to (V="No"; EUM="No"; AF="App Corr"; FAC="Low"; Hack="Yes"). The initialized potential will be:

$$P(AF_{AppCorr} | V_{No}, EUM_{No}, Hack_{Yes}) \times P(V_{No} | FAC_{Low}) = 0.6 \times 0.3 = 0.18$$

Table 4.2 is a summary of cliques with their corresponding initial potentials:

Clique	Initial Potentials
Cost : DL : TD : SD	$P(\text{Cost} \text{DL}, \text{TD}, \text{SD})$
UPS : SD : SF : HAS	$P(\text{SD} \text{UPS}, \text{SF}, \text{Has}) * P(\text{UPS}) * P(\text{HAS})$
PS : SQ : SF	$P(\text{SF} \text{PS}, \text{SQ}) * P(\text{PS}) * P(\text{SQ})$
AF : DL : TD : SF	$P(\text{DL} \text{AF}, \text{SF})$
V : Hack : AF : FAC : EUM	$P(\text{AF} \text{V}, \text{EUM}, \text{Hack}) * P(\text{V} \text{FAC})$
NF : Hack : TD	$P(\text{NF} \text{Hack})$
HAN : TD : NF	$P(\text{TD} \text{NF}, \text{HAN}) * P(\text{HAN})$
F : Hack : FAC	$P(\text{Hack} \text{FAC}, \text{F}) * P(\text{FAC}) * P(\text{F})$

Table 4.2 Initial potentials

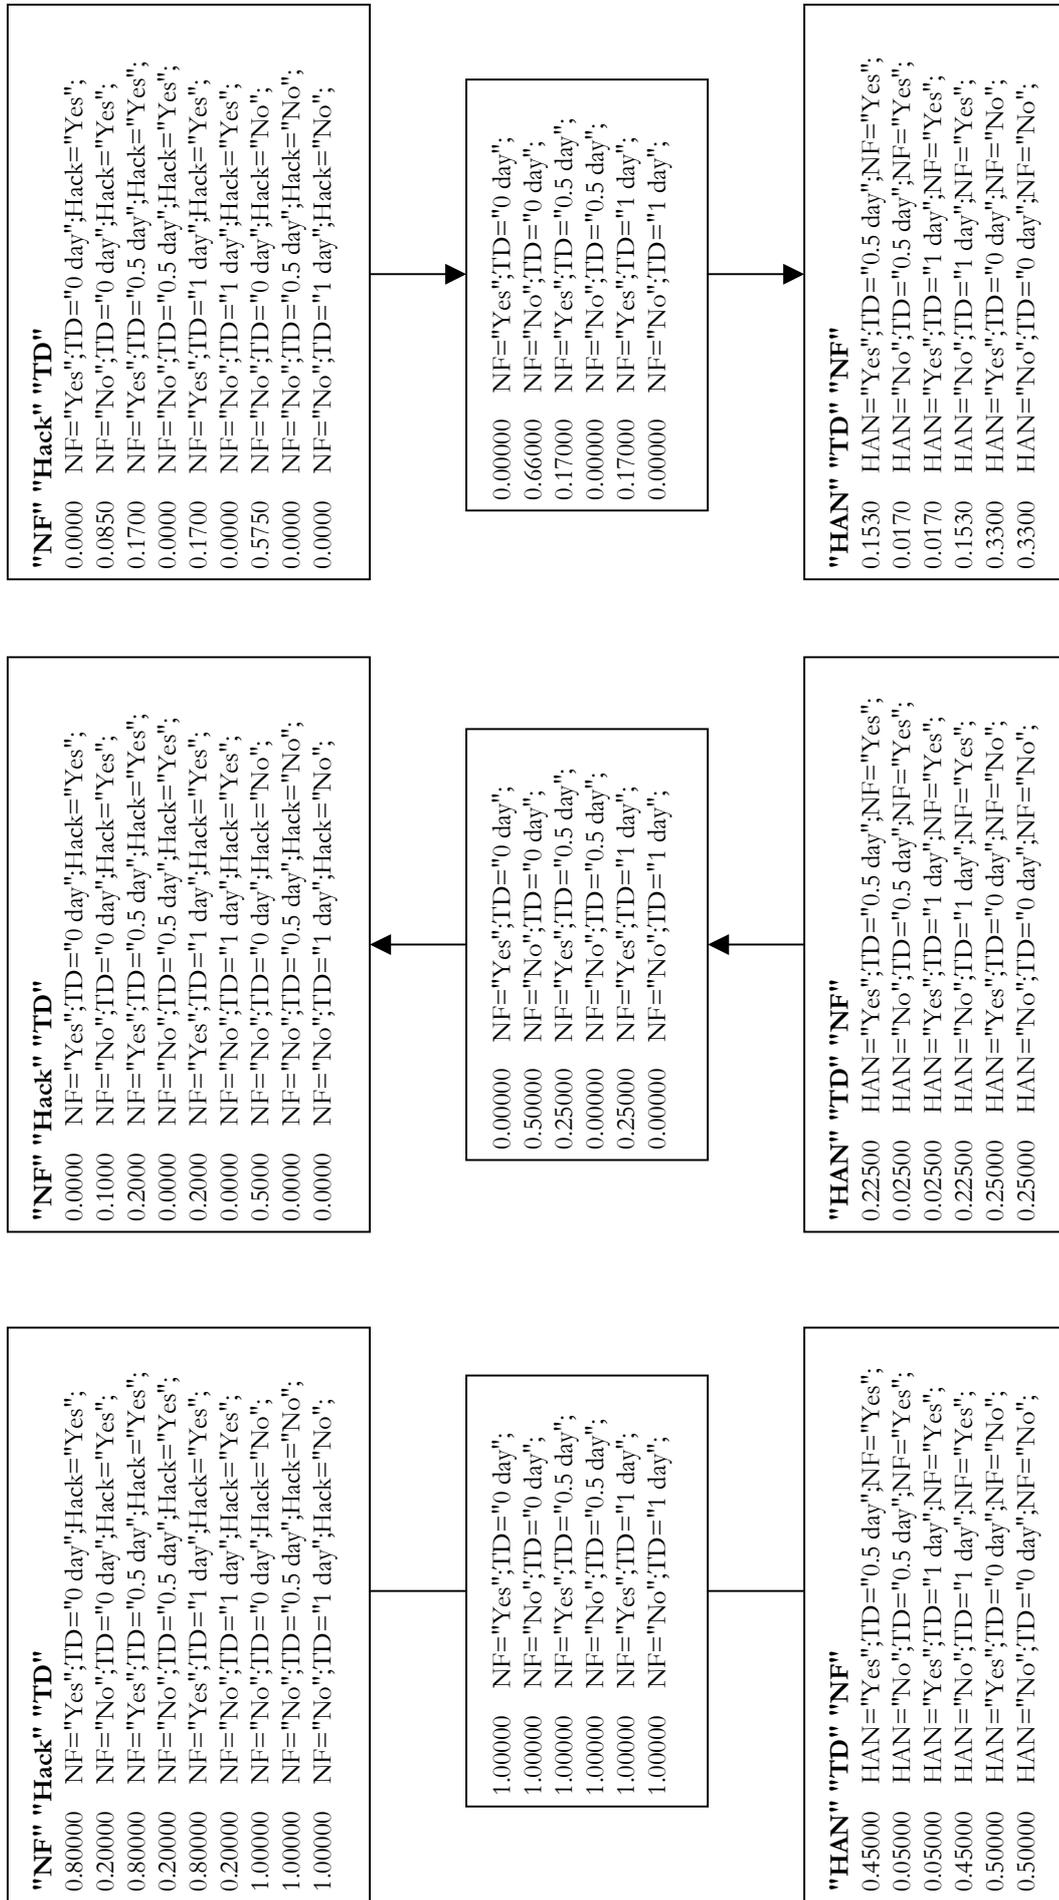
The initialized junction tree can be found in Appendix II.

Two-Phased Propagation

To illustrate the two-phased propagation algorithm for the whole junction tree, we will focus on what goes on at the local level. We shall do this by zooming into the cliques (HAN, TD, NF), denoted C1 and (NF, Hack, TD), denoted C2. The separator for these cliques consists of the variables NF and TD. First, initialization takes place according to the initial potentials given above. The separator potentials are left as unity. The result is seen in *fig 4.6(a)*.

The *collect* phase begins from the clique C1 because it is a leaf clique. Firstly, the clique potentials are normalized. The normalized potentials for C1 are shown in *fig 4.6(b)*. Then, the sum-marginals corresponding to the various combinations of states for variables NF and TD are obtained from C1 and used to update the separator.

For example, all the cells with configuration NF= “Yes” and TD = “0.5 day” are summed up: $0.2250 + 0.0250 = 0.2500$. This is used to update the corresponding potential in the separator. Once the separator has been updated, the sink clique C2 then receives the message from the separator.



(a) Initialization – factorized priors multiplied into the cliques to obtain initial potentials. Please note that zero potentials have been suppressed.

(b) Collect Phase – message is propagated to the rest of the tree and collected at the root.

(b) Distribute Phase – message returning from the root after being propagated across the tree.

Fig. 4.6 – Illustration of Two-phased Propagation

The potentials in C2 are multiplied by the update ratio, defined as:

$$\frac{\text{Separator potential after recent updating}}{\text{Separator potential before recent updating}}$$

Thus, for example, the cell (NF="No";TD="0 day";Hack="Yes")=0.2 before updating would be multiplied by 0.5/1 to get 0.1. When C2 is updated completely (as shown in *fig 4.6(b)*), it will in turn be the sender clique to the next clique in the sequence of the propagation. This algorithm carries on until the root clique is updated. Then the *collect* phase is completed. The result for the whole tree can be seen in Appendix II.

The *distribute* phase then commences and the updating proceeds outwards from the root until it reaches C2 again. In *fig 4.6(c)* we have the updated C2 after the propagation has passed through the rest of the junction tree. This time, C2 is the sender and C1 the sink. The separator potentials then take on the new set of sum-marginals from C2. C1 is then updated according to the update ratio as defined above. For example, the new potential for (HAN="Yes";TD="0.5 day";NF="Yes") would be:

$$\begin{aligned} & \text{Previous potential for (HAN="Yes";TD="0.5 day";NF="Yes")} * \text{Update ratio} \\ & = 0.225 * 0.17 / 0.25 \\ & = 0.153. \end{aligned}$$

From the updated C2, we can sum-marginalize to obtain the posterior marginals for each of the variables in the cliques. Thus, for TD the marginal distribution would be P(0 day, 0.5 day, 1 day) = (0.66, 0.17, 0.17).

Again, the updated cliques for the whole junction tree after the *distribute* phase can be seen in Appendix II. The marginal distributions are shown in *fig 4.7*.

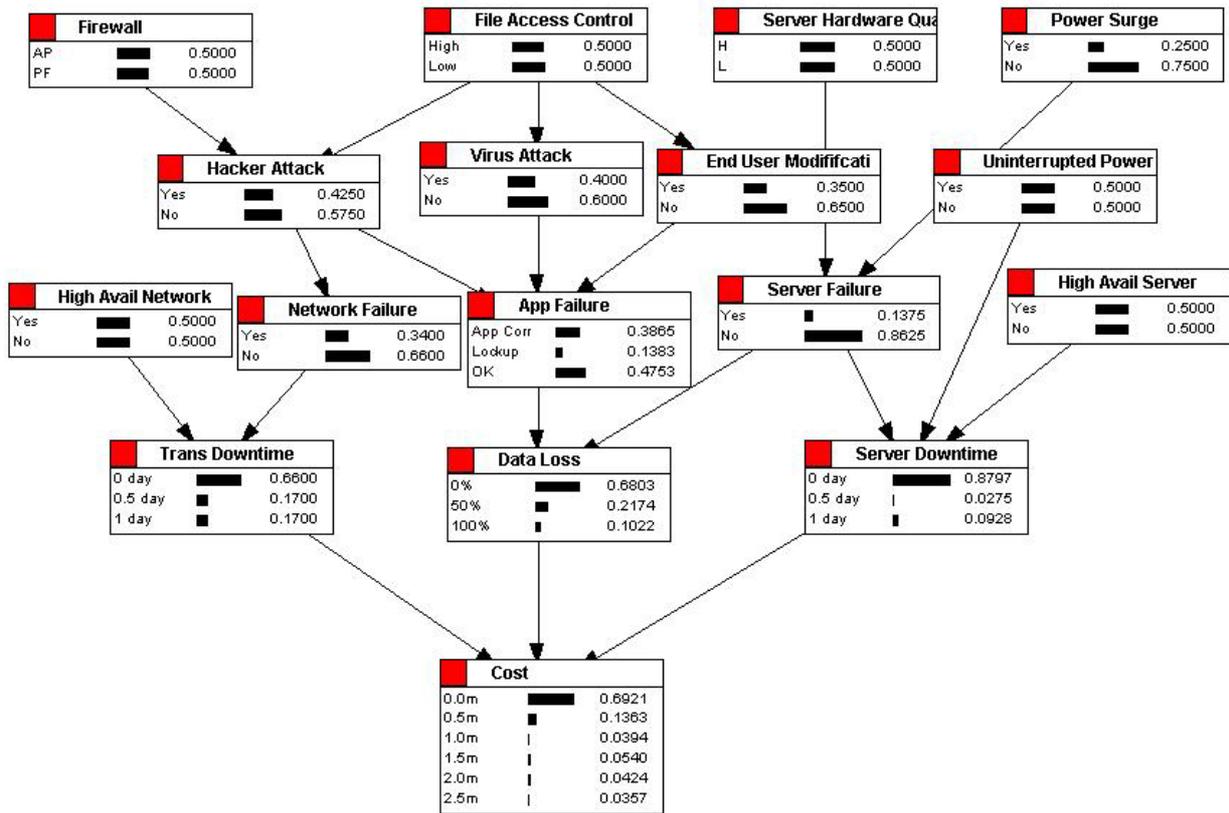


fig 4.7 Prior Marginal Distributions

At this juncture it is worth mentioning that, although this procedure appears rather tedious, the entire two-phased propagation takes place almost instantaneously at a mouse-click when using computer software such as XBaies.

4.5 Applying the BN to OR modeling

4.5.1 Risk Capital Allocation

We now come to one of the main applications of BNs in modeling of OR. Having set up the model as above, the company would like to decide on the amount of financial capital to allocate towards protecting the company from all but the most extreme cases of NR.

At this point the model can be configured to reflect the actual known state of certain variables to reflect the actual position of the company. For example, the I.T. Department may provide the information that Application Proxies are used 24 hours as the firewall for the network and that High Availability Networks have been purchased and implemented for round-the-clock availability. Such information is incorporated by inserting evidence at the relevant cliques as described in §3.4.4. For example, to indicated the usage of Application Proxies, all cells of the clique (F, Hack, FAC) corresponding to F= “PF” is set to 0. This is due to the fact that the prior distribution for $P(F_{AP}, F_{PF})$ is now (1,0) since we know for certain that the type of firewall used is an Application Proxy.

Supposing the set of evidence in *table 4.3* represents the status of BSI’s BSNNet system. This is then incorporated into the junction tree at the initialization stage. The tree is then propagated as before. The new set of posterior marginals incorporating the evidence is shown in *fig 4.8¹⁴*. The evidence was not updated for power surge, as this remains uncertain for a future period.

Variable	Evidence
F	Application Proxy
FAC	High
HAN	Yes
HAS	Yes
SQ	High
UPS	Yes

Table 4.3 Evidence entered

It is now fairly straightforward to determine the risk capital to be allocated to cover NR in BSNNet for **95% of cases over a holding period of one week**. This is simply calculated by linearly interpolating to obtain the 95th percentile of the probability distribution for cost. In this case, the result is **0.32 million**.

¹⁴ XBaies has a user-friendly graphic interface that allows the user to incorporate evidence at the nodes by clicking at the desired state of the variable of interest. E.g., to enter FAC= “High”, all that is needed is a click on the bar representing the marginal probability of the state “High” at the FAC node chart. This automatically sets the bar to 1.

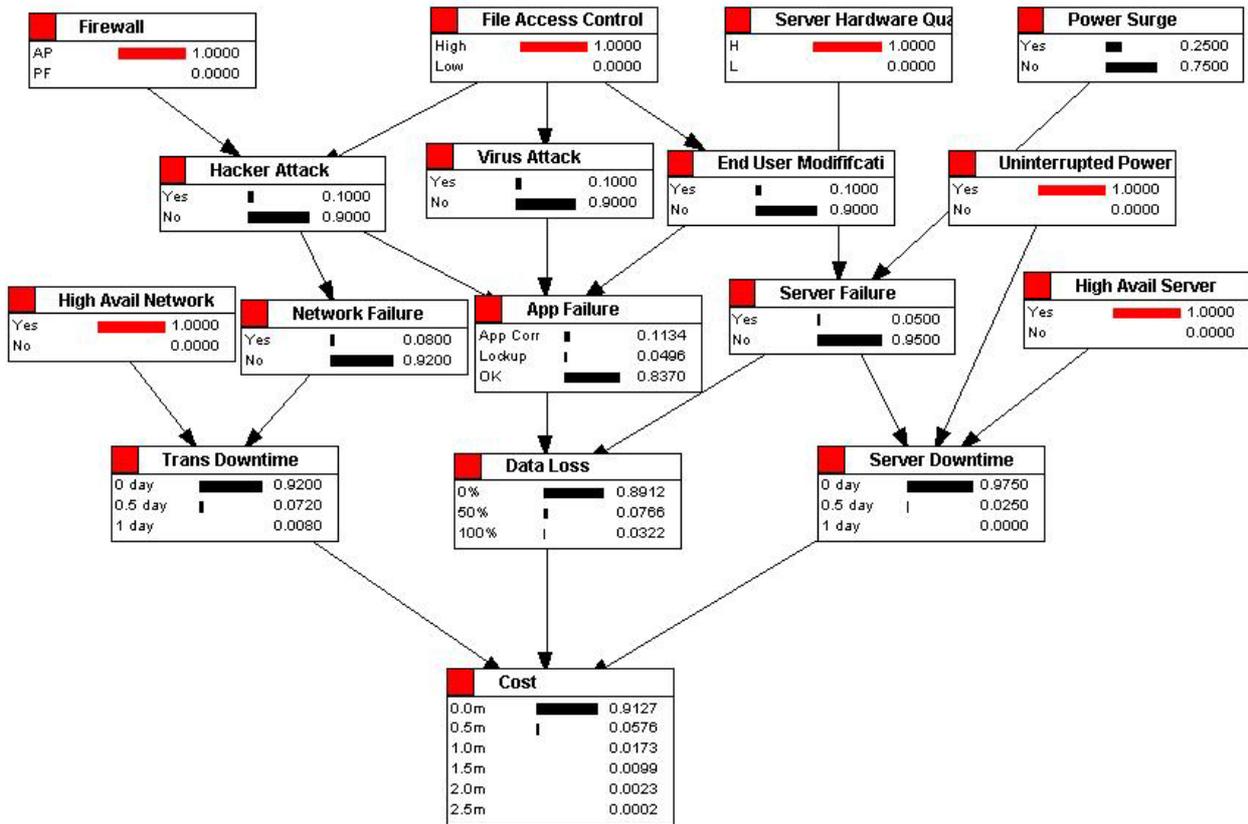


fig 4.8 Posterior marginals after incorporating evidence

4.5.2 Scenario testing and causal analysis

The model can be used to test various scenarios to help management in optimizing its risk profile. For example, if BSI wishes to reduce costs by reducing the level of File Access Control to “Low”, the effect of this action can be easily investigated by entering the evidence FAC= “Low” into the junction tree. The tree is then propagated and the resulting marginals observed as before (this is shown in fig 4.9). We can deduce that the trade-off of the lower costs is an increased capital requirement of **0.97 million**. We assume that the 95% confidence level is maintained.

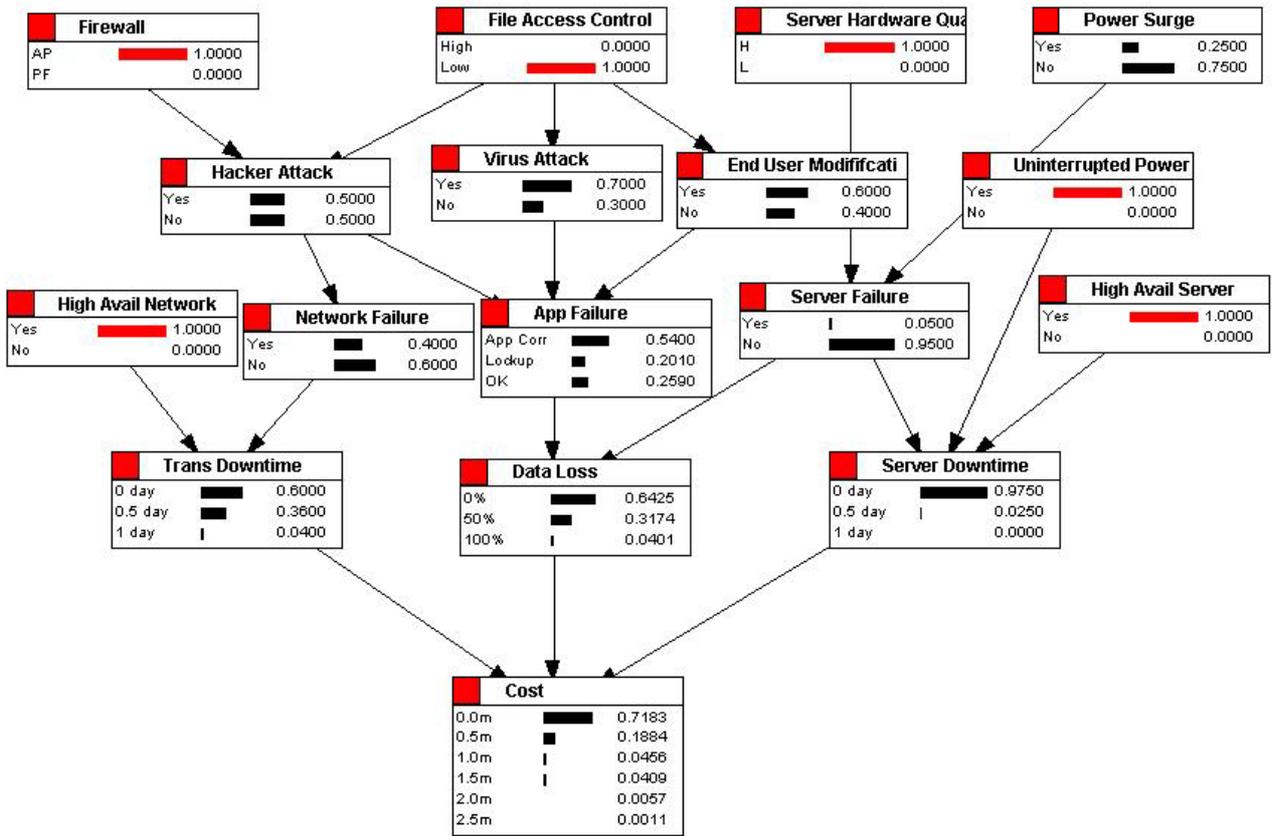


fig 4.9 Posterior marginals of scenario testing

The company may wish to perform some stress testing by investigating the impact of adverse events. Supposing the management is interested to examine the effects of situation of **complete loss of data**. This can be done by setting the node DL to “100%”. The effect is then propagated again and each node is updated accordingly. The final result can be seen in fig 4.10.

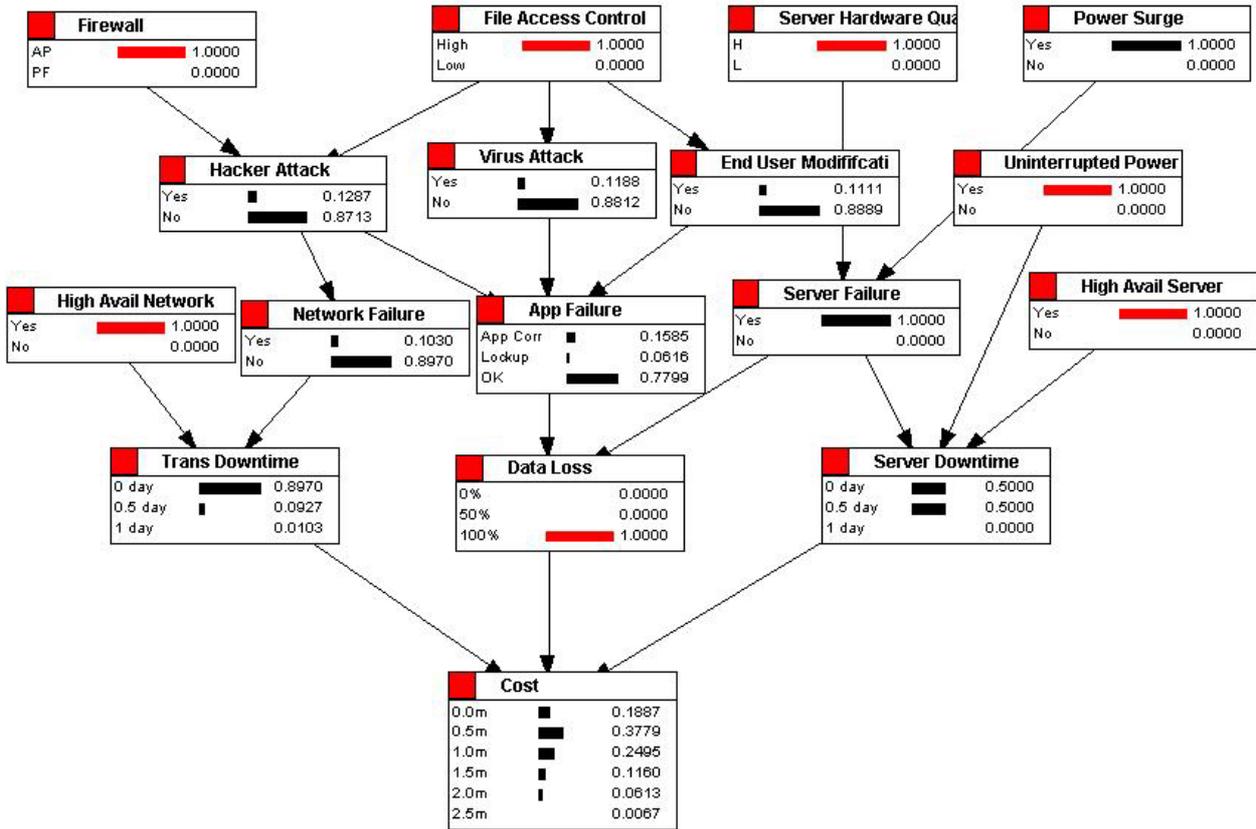


fig 4.10 Cause and effect of a 100% data loss

Note the shift in the marginal probability distribution of the cost. We can now read off a few statistics from this new distribution. The expected cost is now **0.75 million** with standard deviation **0.57 million**. The 95th percentile has now increased to **1.64 million**.

Note also the changes in the other nodes. For example, we now have $P(SF_y)$ and $P(PS_y)$ equal to 1. Conversely, the probability of “No” has increased significantly for Hack, V, EUM and NF. Similarly, AF is overwhelmingly “OK”. We can read from this that server failure is the definite cause of data loss. Consequently, we also deduced that a power surge was the definite cause of server failure. As a result, other events that might have been the causation of data loss have been *explained away*, giving the increase in “No” probabilities for these nodes.

This is an example of how BNs can be useful in causal analysis.

4.5.3 Simulation

We might want to use the BN to generate random sample configurations to simulate actual 52-week periods in which events may occur with probabilities specified in the BN. This can be useful to determine aggregate capital requirements for one-year holding periods.

Using the algorithm for sampling from a junction tree as described in §3.4.5, Xbaies was used to generate 52 sample configurations and the results are displayed in Appendix III. This can be compared to a sample set of size 9999 (using the same evidence) – also shown in Appendix III. We can see that the larger sample size is much closer to the probabilities shown in fig 4.8 from which the samples have been obtained.

4.6 Updating the BN with new data

As the model is implemented, there will be fresh data arriving every week. These can be used to update the probabilities in the model using Dirichlet priors as discussed in §3.6. As an example we consider the conditional probability $P(NF | Hack_y)$ underlying the node NF.

Initially, we have to specify the prior Beta distribution (since this is a binary node). Assuming that the process of prior elicitation is done indirectly via the I.T. Department’s opinion about the respective ranges of the probabilities specified in Appendix I, we obtain the results in Table 4.4.

Parent	NF							
	NF= “Yes”				NF= “No”			
	$\mu_{P(NF_y Hack_y)}$	Range	$\sigma^2_{P(NF_y Hack_y)}$	α_y	$\mu_{P(NF_n Hack_y)}$	Range	$\sigma^2_{P(NF_n Hack_y)}$	α_n
Hack =Yes	0.8	0.70-0.90	0.1 ²	4.3	0.2	0.15-0.25	0.05 ²	1.1

Table 4.4 Prior elicitation to specify Beta parameters

The statistics for the case of NF= “Yes” result in B(12.0,3.0) whereas the statistics for NF= “No” result in B(4.3,1.1) . The second set of results was chosen as it results in a lower precision.

Suppose that over the course of a year, the system was penetrated by malicious hackers in 6 out of the 52 weeks and half of those resulted in network failure. The Beta parameters then can be updated accordingly and the latest probability estimates can be obtained from the mean of the posterior Beta. The results are shown in Table 4.5.

Parent	NF									
	NF= “Yes”					NF= “No”				
	Prior mean	Data	α_y	$\alpha_y + n_y$	Posterior mean	Prior mean	Data	α_n	$\alpha_n + n_n$	Posterior mean
Hack =Yes	0.8	3	4.3	7.3	0.64	0.2	3	1.1	4.1	0.36

Table 4.5 Bayesian updating with data

The results obtained are fairly intuitive: the initial prior estimate of network failure in the event of a hacker attack was reduced from 80% to 64% when the data, which indicates a 50% occurrence, was incorporated. Similar updates can be performed simultaneously for all the nodes with the complete dataset.

We see that Bayesian updating for BNs given new data is fairly easy to perform and would be cost effective to conduct regularly as new cases arrive.

4.7 Model Assessment

To illustrate how the model may be assessed, we shall consider the application of the Parent-Child Monitor at the NF node. The conditional probability underlying this node is P(NF|Hack). We shall investigate the particular case where malicious hacking has occurred. Supposing sufficient time has elapsed such that data is available for 16 weeks where this was the case. We shall see how well the model fits the data.

What I have done is to simulate 50 cases of data from a slightly different population. The distributions are given in *table 4.6*. This was done to investigate how well the model adapts to the data, given that the original prior was not specified very well.

Distribution	$P(NF = Yes Hack = Yes)$	$P(NF = No Hack = Yes)$
Prior	0.8	0.2
Simulated	0.5	0.5

Table 4.6 Probability distribution of $NF | Hack$

Two alternatives were compared: one where the probabilities were updated in the same way shown in §4.6 as each new cases arrive and the other where the probabilities were left as originally specified.

At each update, the following statistics were compiled:

- (i) The logarithmic score (or penalty);
- (ii) The expectation of penalty; and
- (iii) The variance of penalty.

The cumulative penalties and the absolute standardisation test statistics were calculated and then plotted. The results are found in *fig. 4.11* and *fig. 4.12*. Details of the calculations are in Appendix IV.

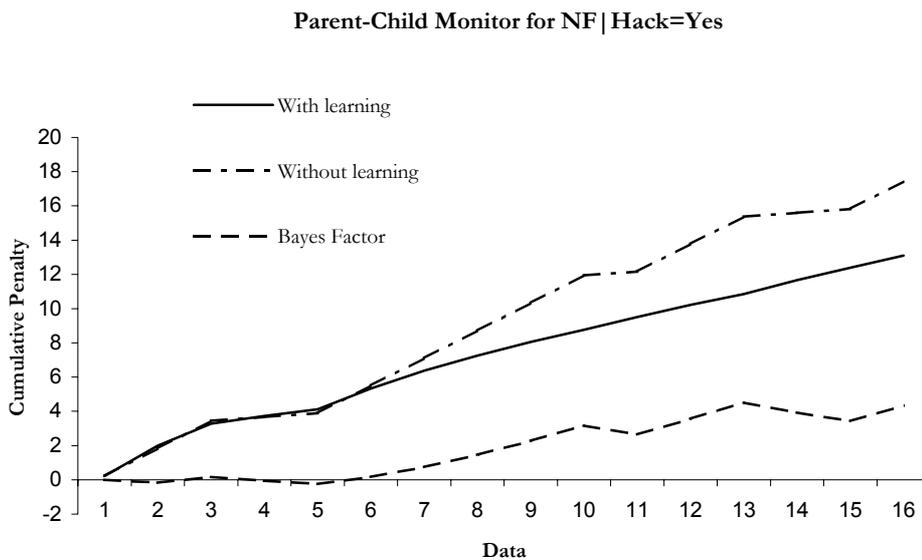


fig. 4.11 Parent-child monitor for NF node: relative standardisation

We can see that just after 16 cases, the model with learning has an advantage of a Bayes factor of 4.34, which implies an odds of 76:1 in its favour.

For the absolute standardisation, the model with learning shows much better results when compared to the model without learning. However, in both cases, the hypothesis test rejects H_0 since the threshold of 1.96 had been exceeded. Some modifications might be necessary in the structure of the model as this parent-child structure might not be a good predictor of the NF node.

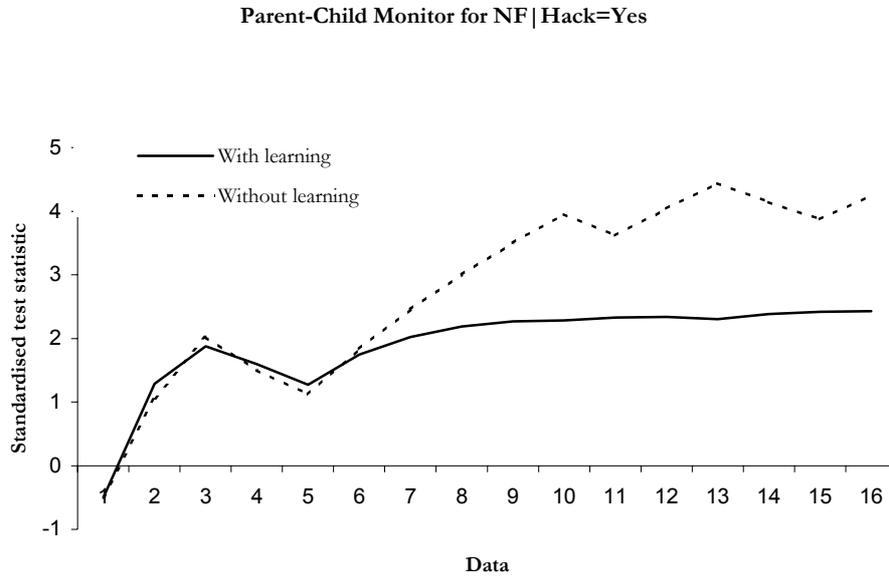


fig. 4.11 Parent-child monitor for NF node: absolute standardization

Similar tests can be performed using node and global monitors in forming an opinion on the viability of the model. In each case, an alternative structure can be tested in parallel with the results plot on the same graph for comparison. This visualisation also facilitates communication of the results.

4.8 Chapter conclusion

We have used a fictitious company's Internet business to illustrate how a BN can be set up using a combination of past data and expert input. In this illustration we have also demonstrated the application of BNs to the business needs of financial institutions in the following areas:

- (i) Setting of regulatory capital for OR;
- (ii) Scenario testing for causal analysis; and
- (iii) Simulation of future scenarios.

The model has been shown to be easily adaptable to incorporate new input. Techniques for assessing the suitability of the model have also been briefly demonstrated.

Chapter Five

Discussion

5.1 Why Bayesian networks?

We shall now summarise the reasons for adopting BNs as an approach to modelling OR in financial institutions.

The main advantage of using BN for modelling OR is the incorporating of expert opinion through:

- (i) Choosing the variables of interest;
- (ii) Defining the structure of the model via the causal dependencies; and
- (iii) Specification of the prior distributions and the conditional probabilities at each node.

Bayesian probability updating ensures that the model is not static, but quickly adapts to new input and incorporates it with prior expert opinion in a mathematically tractable manner. Monitors are also available to enable the efficacy of this process to be observed in real-time, thus facilitating informed model criticism and choice.

We have seen the usage of BNs to model an OR loss distribution, on which business decisions could be based – particularly in the allocation of economic capital. Thus, there is tremendous potential here for an internal model for the setting of regulatory risk capital for OR, such as is required by the emerging regulatory regime for financial institutions. Stress and scenario testing is often a feature of early warning systems in regulatory regimes. We have seen how these can be done fairly quickly on BNs.

Finally, the graphical presentation of BNs aids the understanding of the causal structure and presents the risk profile of the company in an intuitive way – improving management understanding of, and hence participation in, the management of OR.

5.2 Some caveats

Having sung these praises, there are various challenges to overcome in the usage of BN to model OR. Firstly, the model can get very complex with many nodes to specify – this is especially so if the nodes have many parents. In such cases, there can be many conditional probabilities to specify – if the maximum likelihood method of prior elicitation is used, significant volume of data might actually be required, thus reducing one of the main advantages of using Bayesian methods.

The alternative to that would be a rigorous exercise in prior elicitation from experts through costly methods such as the Delphi method, which involves many rounds of questionnaires. A main challenge in this area is in dealing with experts who may not be comfortable thinking in terms of frequencies, although one would hope that this is not the case in financial institutions.

There is also the issue of the non-uniqueness of the causal structure of the model. It is easily seen that choosing a suitable model structure can be as much an art as a science. Although a BN is cheap and easy to run, the whole process of setting up one can be costly, resource consuming and potentially politically messy if many business units/cost centres are involved. In many cases, a fairly complex BN is required to capture all the necessary variables. In addition, accuracy of results might be pursued at the expense of model parsimony especially when business (or even regulatory) decisions are at stake.

The main advantage of BNs – ability to incorporate subjective knowledge – can be a disadvantage when it comes to setting regulatory capital. Regulators who require an objective standard to approve of internal models may find it difficult to find a standard for acceptance of BNs due to its high subjective content. Regulators might need to specify rules on the process of model specification and prior elicitation to reduce the subjectivity.

5.3 Extensions

In the Network Risk example, the distributions have been expressed as discrete probabilities for simplicity. More advanced modelling can be used to deal with continuous distributions in BNs – these are also known as conditional Gaussian distributions. Chapter 7 of Cowell et al (1999) provides a fairly comprehensive treatment of CG.

In the example above, we have assumed the existence of complete data. In reality, this is seldom the case. Certain data cells will be missing or just impossible to obtain. Cowell et al (1999) suggests various methods such as the EM algorithm and the Gibbs Sampler to deal with updating of the BN with incomplete data. The Gibbs sampler uses a Markov Chain Monte Carlo process to recursively update the parameters of a Bayesian network to obtain a predictive distribution.

As discussed above, specification of the structure is often subject to debate. Thus, a logical progression is to move towards structural learning – i.e. letting the data speak for itself not just with regards the probability distributions of the variables but even the very structure itself. This is currently an active area of research and chapter 11 of Cowell et al (1999) provides useful suggestions of work being done.

Post-script

Probabilistic reasoning using expert systems, of which BNs are an example, is a specific case of a larger universe of scientific modelling known as Artificial Intelligence. This exciting confluence of various streams of scientific knowledge – statistics, engineering, econometrics, and information technology – will continue to be a fertile space for collaboration to discover newer and better ways to understand the world we live in. In researching for this dissertation, I am humbled by the amount of thought that has been put into this field yet I am also inspired to be part of its evolution through my pilgrimage within the actuarial profession and beyond.

Acknowledgements

I would like to thank the following people for their invaluable input and assistance in this dissertation:

Professor Richard Verrall – for being as excited as I have been by this whole project and for being such an encouraging supervisor.

Dr. Robert Cowell – for introducing me to the world of probabilistic reasoning and for the priceless signed copy of *Probabilistic Networks and Expert Systems* (see Cowell et. al (1999)).

Professor Peter England – for the industry perspective and for suggesting I.T. related risks as an example.

Professor Steven Haberman – for having spotted this as a doable dissertation in the first place.

Kevin Lai of Network Associates Incorporated, Kuala Lumpur – for helping me set up the Network Risk example and for your friendship.

My colleagues and seniors at the Insurance Supervision Department and Cik Roslina Hitam of Human Resources, Bank Negara Malaysia – for all the support from home.

Marilyn Parris – for running so many of my errands and for inspiration to be salt and light in the marketplace.

City University Student Union – for maintaining the photocopier.

All the staff at the Faculty of Actuarial Science, CASS Business School.

Special thanks to my Lord and Saviour, Jesus Christ, without whom none of this would be possible. To Him be the Glory.

References

- ALEXANDER, CAROL (2000). Bayesian methods for measuring operational risk. *Discussion Papers in Finance 2000-02*. ISMA Centre, University of Reading.
- BASEL COMMITTEE ON BANKING SUPERVISION (1999). A new capital adequacy framework. Bank for International Settlements.
- BASEL COMMITTEE ON BANKING SUPERVISION (2001a). Working paper on the regulatory treatment of operational risk. Bank for International Settlements.
- BASEL COMMITTEE ON BANKING SUPERVISION (2001b). Consultative Paper 2.5. Bank for International Settlements.
- BASEL COMMITTEE ON BANKING SUPERVISION (2002). Operational risk data collection exercise. Bank for International Settlements.
- CESKE, R., HERNANDEZ, J.V., SANCHEZ, L.M. (2000). Quantifying event risk: the next convergence. *Journal of Risk Finance*. Spring 2000, 1 – 13.
- COWELL, R. G., DAWID, A. P., LAURITZEN, S. L., SPIEGELHALTER, D. J., (1999). Probabilistic networks and expert systems. Springer-Verlag.
- CRUZ, MARCELO G. (2002) Modeling, measuring and hedging operational risk. Wiley.
- DOWD, KEVIN (1998). Beyond value at risk: the new science of risk management. Wiley.
- FRACHOT, ANTOINE & RONCALLI, THIERRY (2002). Mixing internal and external data for managing operational risk. Groupe de Recherche Operationelle, Credit Lyonnais.
- HOFFMAN, DOUGLAS (2002). Managing operational risk: 20 firmwide best practice strategies. Wiley.
- INSTITUTE OF ACTUARIES (2002). Report of the operational risk working party to GIRO 2002. Available to members at <http://www.actuaries.org.uk>
- JENSEN, FINN V. (1996). An introduction to Bayesian networks. Springer.

JORION, PHILIPPE (2001). Value at risk: the new benchmark for managing financial risk. McGraw-Hill.

KING, JACK L. (1999). Operational Risk. Wiley.

MARSHALL, CHRISTOPHER (2001). Measuring and managing operational risks in financial institutions. Wiley

MICCOLIS, JEREMY A. & SHAH, SAMIR (2000). Getting a handle on operational risks. *Emphasis 2000/1*. Tillinghast-Towers Perrin.

SHAH, SAMIR (2001). Operational risk management. *Casualty Actuarial Society 2001 Seminar on Understanding the Enterprise Risk Management Process*. Powerpoint slides available from Casualty Actuarial Society at <http://www.casact.org>

VAN DEN BRINK, GERRIT J. (2002). Operational risk: the new challenge for banks. Palgrave.

Network Risk: Prior Distributions

Unconditional Priors

FAC		HAN		HAS		F	
High	Low	Yes	No	Yes	No	AP	PF
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5

PS		SQ		UPS	
Yes	No	High	Low	Yes	No
0.25	0.75	0.5	0.5	0.5	0.5

Conditional Priors

SF PS, SQ				Hack FAC, F				TD HAN, NF				
		SF				Hack				TD		
SQ	PS	Y	N	F	FAC	Y	N	HAN	NF	0 day	0.5 day	1 day
H	Y	0.2	0.8	AP	High	0.1	0.9	Y	Y	0	0.9	0.1
	N	0	1		Low	0.5	0.5		N	1	0	0
L	Y	0.9	0.1	PF	High	0.3	0.7	N	Y	0	0.1	0.9
	N	0	1		Low	0.8	0.2		N	1	0	0

AF V, EUM, Hack						SD UPS, SF, HAS					
		AF						SD			
Hack	EUM	V	App Corr	Lockup	OK	HAS	SF	UPS	0 day	0.5 day	1 day
Y	Y	Y	0.9	0.1	0	Y	Y	Y	0.5	0.5	0
		N	0.6	0.4	0			N	0	0.2	0.8
	N	Y	0.7	0.3	0		N	Y	1	0	0
		N	0.6	0.2	0.2			N	1	0	0
	Y	Y	0.4	0.3	0.3	N		Y	0	0.1	0.9
		N	0.1	0.3	0.6			N	0	0	1
		Y	0.5	0	0.5			Y	1	0	0
		N	0	0	1			N	1	0	0

V FAC			EUM FAC			NF Hack		
		V			EUM			
FAC	Y	N	FAC	Y	N	Hack	Y	N
High	0.1	0.9	High	0.1	0.9	Y	0.8	0.2
Low	0.7	0.3	Low	0.6	0.4	N	0	1

DL AF, SF				
		DL		
SF	AF	0%	50%	100%
Y	App Corr	0.1	0	0.9
	Lockup	0.1	0.1	0.8
	OK	0.2	0.2	0.6
N	App Corr	0.5	0.5	0
	Lockup	0.7	0.3	0
	OK	1	0	0

Cost | SD, TD, DL

DL	TD	SD	Cost					
			0.0m	0.5m	1.0m	1.5m	2.0m	2.5m
0%	0 day	0 day	1	0	0	0	0	0
		0.5 day	0	0.6	0.4	0	0	0
		1 day	0	0.1	0.2	0.3	0.2	0.2
	0.5 day	0 day	0.7	0.3	0	0	0	0
		0.5 day	0	0.5	0.3	0.2	0	0
		1 day	0	0	0.2	0.4	0.2	0.2
	1 day	0 day	0.6	0.4	0	0	0	0
		0.5 day	0	0.4	0.3	0.2	0.1	0
		1 day	0	0	0.1	0.2	0.4	0.3
50%	0 day	0 day	0.6	0.4	0	0	0	0
		0.5 day	0	0.5	0.4	0.1	0	0
		1 day	0	0	0.1	0.3	0.4	0.2
	0.5 day	0 day	0.3	0.3	0.2	0.2	0	0
		0.5 day	0	0.3	0.3	0.2	0.2	0
		1 day	0	0	0.1	0.2	0.5	0.2
	1 day	0 day	0.3	0.2	0.2	0.2	0.1	0
		0.5 day	0	0.3	0.2	0.2	0.2	0.1
		1 day	0	0	0	0.1	0.4	0.5
100%	0 day	0 day	0.4	0.4	0.2	0	0	0
		0.5 day	0	0.4	0.3	0.2	0.1	0
		1 day	0	0	0	0.3	0.4	0.3
	0.5 day	0 day	0.2	0.2	0.3	0.2	0.1	0
		0.5 day	0	0.2	0.2	0.3	0.2	0.1
		1 day	0	0	0	0.2	0.4	0.4
	1 day	0 day	0	0.1	0.2	0.3	0.2	0.2
		0.5 day	0	0	0.2	0.3	0.3	0.2
		1 day	0	0	0	0	0.2	0.8

**NETWORK RISK: PROPAGATION EXAMPLE
INITIALIZATION (Continued)**

DL:TD:SD:SF node
From previous page

"AF" "DL" "TD" "SF"

```

0.10000 AF="App Corr";TD="0 day";DL="0%";SF="Yes";
0.10000 AF="Lockup";TD="0 day";DL="0%";SF="Yes";
0.20000 AF="OK";TD="0 day";DL="0%";SF="Yes";
0.10000 AF="App Corr";TD="0.5 day";DL="0%";SF="Yes";
0.10000 AF="Lockup";TD="0.5 day";DL="0%";SF="Yes";
0.20000 AF="OK";TD="0.5 day";DL="0%";SF="Yes";
0.10000 AF="App Corr";TD="1 day";DL="0%";SF="Yes";
0.10000 AF="Lockup";TD="1 day";DL="0%";SF="Yes";
0.20000 AF="OK";TD="1 day";DL="0%";SF="Yes";
0.10000 AF="App Corr";TD="0 day";DL="50%";SF="Yes";
0.20000 AF="OK";TD="0 day";DL="50%";SF="Yes";
0.10000 AF="App Corr";TD="0.5 day";DL="50%";SF="Yes";
0.20000 AF="OK";TD="0.5 day";DL="50%";SF="Yes";
0.10000 AF="App Corr";TD="1 day";DL="50%";SF="Yes";
0.20000 AF="OK";TD="1 day";DL="50%";SF="Yes";
0.90000 AF="App Corr";TD="0 day";DL="100%";SF="Yes";
0.80000 AF="Lockup";TD="0 day";DL="100%";SF="Yes";
0.60000 AF="OK";TD="0 day";DL="100%";SF="Yes";
0.90000 AF="App Corr";TD="0.5 day";DL="100%";SF="Yes";
0.80000 AF="Lockup";TD="0.5 day";DL="100%";SF="Yes";
0.60000 AF="OK";TD="0.5 day";DL="100%";SF="Yes";
0.90000 AF="App Corr";TD="1 day";DL="100%";SF="Yes";
0.80000 AF="Lockup";TD="1 day";DL="100%";SF="Yes";
0.60000 AF="OK";TD="1 day";DL="100%";SF="Yes";
0.50000 AF="App Corr";TD="0 day";DL="0%";SF="No";
0.70000 AF="Lockup";TD="0 day";DL="0%";SF="No";
1.00000 AF="OK";TD="0 day";DL="0%";SF="No";
0.50000 AF="App Corr";TD="0.5 day";DL="0%";SF="No";
0.70000 AF="Lockup";TD="0.5 day";DL="0%";SF="No";
1.00000 AF="OK";TD="0.5 day";DL="0%";SF="No";
0.50000 AF="App Corr";TD="1 day";DL="0%";SF="No";
0.70000 AF="Lockup";TD="1 day";DL="0%";SF="No";
1.00000 AF="OK";TD="1 day";DL="0%";SF="No";
0.50000 AF="App Corr";TD="0 day";DL="50%";SF="No";
0.30000 AF="Lockup";TD="0 day";DL="50%";SF="No";
0.50000 AF="App Corr";TD="0.5 day";DL="50%";SF="No";
0.30000 AF="Lockup";TD="0.5 day";DL="50%";SF="No";
0.50000 AF="App Corr";TD="1 day";DL="50%";SF="No";
0.30000 AF="Lockup";TD="1 day";DL="50%";SF="No";
    
```

"NF" "Hack" "TD"

```

0.80000 NF="Yes";TD="0 day";Hack="Yes";
0.20000 NF="No";TD="0 day";Hack="Yes";
0.80000 NF="Yes";TD="0.5 day";Hack="Yes";
0.20000 NF="No";TD="0.5 day";Hack="Yes";
0.80000 NF="Yes";TD="1 day";Hack="Yes";
0.20000 NF="No";TD="1 day";Hack="Yes";
1.00000 NF="No";TD="0 day";Hack="No";
1.00000 NF="No";TD="0.5 day";Hack="No";
1.00000 NF="No";TD="1 day";Hack="No";
    
```

"HAN" "TD" "NF"

```

0.45000 HAN="Yes";TD="0.5 day";NF="Yes";
0.05000 HAN="No";TD="0.5 day";NF="Yes";
0.05000 HAN="Yes";TD="1 day";NF="Yes";
0.45000 HAN="No";TD="1 day";NF="Yes";
0.50000 HAN="Yes";TD="0 day";NF="No";
0.50000 HAN="No";TD="0 day";NF="No";
    
```

"V" "Hack" "AF" "FAC" "EUM"

```

0.09000 V="Yes";EUM="Yes";AF="App Corr";FAC="High";Hack="Yes";
0.54000 V="No";EUM="Yes";AF="App Corr";FAC="High";Hack="Yes";
0.07000 V="Yes";EUM="No";AF="App Corr";FAC="High";Hack="Yes";
0.54000 V="No";EUM="No";AF="App Corr";FAC="High";Hack="Yes";
0.01000 V="Yes";EUM="Yes";AF="Lockup";FAC="High";Hack="Yes";
0.36000 V="No";EUM="Yes";AF="Lockup";FAC="High";Hack="Yes";
0.03000 V="Yes";EUM="No";AF="Lockup";FAC="High";Hack="Yes";
0.18000 V="No";EUM="No";AF="Lockup";FAC="High";Hack="Yes";
0.18000 V="No";EUM="No";AF="OK";FAC="High";Hack="Yes";
0.63000 V="Yes";EUM="Yes";AF="App Corr";FAC="Low";Hack="Yes";
0.18000 V="No";EUM="Yes";AF="App Corr";FAC="Low";Hack="Yes";
0.49000 V="Yes";EUM="No";AF="App Corr";FAC="Low";Hack="Yes";
0.18000 V="No";EUM="No";AF="App Corr";FAC="Low";Hack="Yes";
0.07000 V="Yes";EUM="Yes";AF="Lockup";FAC="Low";Hack="Yes";
0.12000 V="No";EUM="Yes";AF="Lockup";FAC="Low";Hack="Yes";
0.21000 V="Yes";EUM="No";AF="Lockup";FAC="Low";Hack="Yes";
0.06000 V="No";EUM="No";AF="Lockup";FAC="Low";Hack="Yes";
0.06000 V="No";EUM="No";AF="OK";FAC="Low";Hack="Yes";
0.04000 V="Yes";EUM="Yes";AF="App Corr";FAC="High";Hack="No";
0.09000 V="No";EUM="Yes";AF="App Corr";FAC="High";Hack="No";
0.05000 V="Yes";EUM="No";AF="App Corr";FAC="High";Hack="No";
0.03000 V="Yes";EUM="Yes";AF="Lockup";FAC="High";Hack="No";
0.27000 V="No";EUM="Yes";AF="Lockup";FAC="High";Hack="No";
0.03000 V="Yes";EUM="Yes";AF="OK";FAC="High";Hack="No";
0.54000 V="No";EUM="Yes";AF="OK";FAC="High";Hack="No";
0.05000 V="Yes";EUM="No";AF="OK";FAC="High";Hack="No";
0.90000 V="No";EUM="No";AF="OK";FAC="High";Hack="No";
0.28000 V="Yes";EUM="Yes";AF="App Corr";FAC="Low";Hack="No";
0.03000 V="No";EUM="Yes";AF="App Corr";FAC="Low";Hack="No";
0.35000 V="Yes";EUM="No";AF="App Corr";FAC="Low";Hack="No";
0.21000 V="Yes";EUM="Yes";AF="Lockup";FAC="Low";Hack="No";
0.09000 V="No";EUM="Yes";AF="Lockup";FAC="Low";Hack="No";
0.21000 V="Yes";EUM="Yes";AF="OK";FAC="Low";Hack="No";
0.18000 V="No";EUM="Yes";AF="OK";FAC="Low";Hack="No";
0.35000 V="Yes";EUM="No";AF="OK";FAC="Low";Hack="No";
0.30000 V="No";EUM="No";AF="OK";FAC="Low";Hack="No";
    
```

"Hack" "AF" "TD"

```

1.00000 Hack="Yes";TD="0 day";AF="App Corr";
1.00000 Hack="No";TD="0 day";AF="App Corr";
1.00000 Hack="Yes";TD="0.5 day";AF="App Corr";
1.00000 Hack="No";TD="0.5 day";AF="App Corr";
1.00000 Hack="Yes";TD="1 day";AF="App Corr";
1.00000 Hack="No";TD="1 day";AF="App Corr";
1.00000 Hack="Yes";TD="0 day";AF="Lockup";
1.00000 Hack="No";TD="0 day";AF="Lockup";
1.00000 Hack="Yes";TD="0.5 day";AF="Lockup";
1.00000 Hack="No";TD="0.5 day";AF="Lockup";
1.00000 Hack="Yes";TD="1 day";AF="Lockup";
1.00000 Hack="No";TD="1 day";AF="Lockup";
1.00000 Hack="Yes";TD="0 day";AF="OK";
1.00000 Hack="No";TD="0 day";AF="OK";
1.00000 Hack="Yes";TD="0.5 day";AF="OK";
1.00000 Hack="No";TD="0.5 day";AF="OK";
1.00000 Hack="Yes";TD="1 day";AF="OK";
1.00000 Hack="No";TD="1 day";AF="OK";
    
```

"F" "Hack" "FAC"

```

0.05000 F="AP";FAC="High";Hack="Yes";
0.15000 F="PF";FAC="High";Hack="Yes";
0.25000 F="AP";FAC="Low";Hack="Yes";
0.40000 F="PF";FAC="Low";Hack="Yes";
0.45000 F="AP";FAC="High";Hack="No";
0.35000 F="PF";FAC="High";Hack="No";
0.25000 F="AP";FAC="Low";Hack="No";
0.10000 F="PF";FAC="Low";Hack="No";
    
```

NETWORK RISK: PROPAGATION EXAMPLE
COLLECT PHASE

"Cost" "DL" "TD" "SD"

0.03704 Cost="0.0m";TD="0 day";SD="0 day";DL="0%";
 0.02593 Cost="0.0m";TD="0.5 day";SD="0 day";DL="0%";
 0.01111 Cost="0.5m";TD="0.5 day";SD="0 day";DL="0%";
 0.02222 Cost="0.0m";TD="1 day";SD="0 day";DL="0%";
 0.01481 Cost="0.5m";TD="1 day";SD="0 day";DL="0%";
 0.02222 Cost="0.5m";TD="0 day";SD="0.5 day";DL="0%";
 0.01481 Cost="1.0m";TD="0 day";SD="0.5 day";DL="0%";
 0.01852 Cost="0.5m";TD="0.5 day";SD="0.5 day";DL="0%";
 0.01111 Cost="1.0m";TD="0.5 day";SD="0.5 day";DL="0%";
 0.00741 Cost="1.5m";TD="0.5 day";SD="0.5 day";DL="0%";
 0.01481 Cost="0.5m";TD="1 day";SD="0.5 day";DL="0%";
 0.01111 Cost="1.0m";TD="1 day";SD="0.5 day";DL="0%";
 0.00741 Cost="1.5m";TD="1 day";SD="0.5 day";DL="0%";
 0.01111 Cost="2.0m";TD="1 day";SD="0.5 day";DL="0%";
 0.00370 Cost="0.5m";TD="0 day";SD="1 day";DL="0%";
 0.00741 Cost="1.0m";TD="0 day";SD="1 day";DL="0%";
 0.01111 Cost="1.5m";TD="0 day";SD="1 day";DL="0%";
 0.00741 Cost="2.0m";TD="0 day";SD="1 day";DL="0%";
 0.00741 Cost="2.5m";TD="0 day";SD="1 day";DL="0%";
 0.00741 Cost="1.0m";TD="0.5 day";SD="1 day";DL="0%";
 0.01481 Cost="1.5m";TD="0.5 day";SD="1 day";DL="0%";
 0.00741 Cost="2.0m";TD="0.5 day";SD="1 day";DL="0%";
 0.00741 Cost="2.5m";TD="0.5 day";SD="1 day";DL="0%";
 0.00370 Cost="1.0m";TD="1 day";SD="1 day";DL="0%";
 0.00741 Cost="1.5m";TD="1 day";SD="1 day";DL="0%";
 0.01481 Cost="2.0m";TD="1 day";SD="1 day";DL="0%";
 0.01111 Cost="2.5m";TD="1 day";SD="1 day";DL="0%";
 0.02222 Cost="0.0m";TD="0 day";SD="0 day";DL="50%";
 0.01481 Cost="0.5m";TD="0 day";SD="0 day";DL="50%";
 0.01111 Cost="0.0m";TD="0.5 day";SD="0 day";DL="50%";
 0.01111 Cost="0.5m";TD="0.5 day";SD="0 day";DL="50%";
 0.00741 Cost="1.0m";TD="0.5 day";SD="0 day";DL="50%";
 0.00741 Cost="1.5m";TD="0.5 day";SD="0 day";DL="50%";
 0.01111 Cost="2.0m";TD="0.5 day";SD="0 day";DL="50%";
 0.00741 Cost="2.5m";TD="0.5 day";SD="0 day";DL="50%";
 0.00741 Cost="0.0m";TD="1 day";SD="0 day";DL="50%";
 0.00741 Cost="0.5m";TD="1 day";SD="0 day";DL="50%";
 0.00741 Cost="1.0m";TD="1 day";SD="0 day";DL="50%";
 0.00741 Cost="1.5m";TD="1 day";SD="0 day";DL="50%";
 0.00370 Cost="2.0m";TD="1 day";SD="0 day";DL="50%";
 0.01852 Cost="0.5m";TD="0 day";SD="0.5 day";DL="50%";
 0.01481 Cost="1.0m";TD="0 day";SD="0.5 day";DL="50%";
 0.00370 Cost="1.5m";TD="0 day";SD="0.5 day";DL="50%";
 0.01111 Cost="2.0m";TD="0 day";SD="0.5 day";DL="50%";
 0.01481 Cost="2.5m";TD="0 day";SD="0.5 day";DL="50%";
 0.00741 Cost="1.0m";TD="0 day";SD="1 day";DL="50%";
 0.00741 Cost="0.0m";TD="0.5 day";SD="0 day";DL="100%";
 0.00741 Cost="0.5m";TD="0.5 day";SD="0 day";DL="100%";
 0.01111 Cost="1.0m";TD="0.5 day";SD="0 day";DL="100%";
 0.00741 Cost="1.5m";TD="0.5 day";SD="0 day";DL="100%";
 0.00370 Cost="2.0m";TD="0.5 day";SD="0 day";DL="100%";
 0.00370 Cost="0.5m";TD="1 day";SD="0 day";DL="100%";
 0.00741 Cost="1.0m";TD="1 day";SD="0 day";DL="100%";
 0.01111 Cost="1.5m";TD="1 day";SD="0 day";DL="100%";
 0.00741 Cost="2.0m";TD="1 day";SD="0 day";DL="100%";
 0.00741 Cost="2.5m";TD="1 day";SD="0 day";DL="100%";
 0.01111 Cost="1.0m";TD="0 day";SD="0.5 day";DL="100%";
 0.01481 Cost="1.5m";TD="0 day";SD="0.5 day";DL="100%";
 0.01111 Cost="2.0m";TD="0 day";SD="0.5 day";DL="100%";
 0.00741 Cost="2.5m";TD="0 day";SD="0.5 day";DL="100%";
 0.00741 Cost="1.5m";TD="0.5 day";SD="1 day";DL="100%";
 0.01481 Cost="2.0m";TD="0.5 day";SD="1 day";DL="100%";
 0.01481 Cost="2.5m";TD="0.5 day";SD="1 day";DL="100%";
 0.00741 Cost="2.0m";TD="1 day";SD="1 day";DL="100%";
 0.02963 Cost="2.5m";TD="1 day";SD="1 day";DL="100%";

"DL" "TD" "SD" "SF"

0.0186 DL="0%";TD="0 day";SD="0 day";SF="Yes";
 0.0162 DL="50%";TD="0 day";SD="0 day";SF="Yes";
 0.0752 DL="100%";TD="0 day";SD="0 day";SF="Yes";
 0.0030 DL="0%";TD="0.5 day";SD="0 day";SF="Yes";
 0.0010 DL="50%";TD="0.5 day";SD="0 day";SF="Yes";
 0.0244 DL="100%";TD="0.5 day";SD="0 day";SF="Yes";
 0.0030 DL="0%";TD="1 day";SD="0 day";SF="Yes";
 0.0010 DL="50%";TD="1 day";SD="0 day";SF="Yes";
 0.0244 DL="100%";TD="1 day";SD="0 day";SF="Yes";
 0.0186 DL="0%";TD="0 day";SD="0.5 day";SF="Yes";
 0.0162 DL="50%";TD="0 day";SD="0.5 day";SF="Yes";
 0.0752 DL="100%";TD="0 day";SD="0.5 day";SF="Yes";
 0.0030 DL="0%";TD="0.5 day";SD="0.5 day";SF="Yes";
 0.0010 DL="50%";TD="0.5 day";SD="0.5 day";SF="Yes";
 0.0244 DL="100%";TD="0.5 day";SD="0.5 day";SF="Yes";
 0.0030 DL="0%";TD="1 day";SD="0.5 day";SF="Yes";
 0.0010 DL="50%";TD="1 day";SD="0.5 day";SF="Yes";
 0.0244 DL="100%";TD="1 day";SD="0.5 day";SF="Yes";
 0.0186 DL="0%";TD="0 day";SD="1 day";SF="Yes";
 0.0162 DL="50%";TD="0 day";SD="1 day";SF="Yes";
 0.0752 DL="100%";TD="0 day";SD="1 day";SF="Yes";
 0.0030 DL="0%";TD="0.5 day";SD="1 day";SF="Yes";
 0.0010 DL="50%";TD="0.5 day";SD="1 day";SF="Yes";
 0.0244 DL="100%";TD="0.5 day";SD="1 day";SF="Yes";
 0.0030 DL="0%";TD="1 day";SD="1 day";SF="Yes";
 0.0010 DL="50%";TD="1 day";SD="1 day";SF="Yes";
 0.0244 DL="100%";TD="1 day";SD="1 day";SF="Yes";
 0.0951 DL="0%";TD="0 day";SD="0 day";SF="No";
 0.0149 DL="50%";TD="0 day";SD="0 day";SF="No";
 0.0000 DL="100%";TD="0 day";SD="0 day";SF="No";
 0.0162 DL="0%";TD="0.5 day";SD="0 day";SF="No";
 0.0121 DL="50%";TD="0.5 day";SD="0 day";SF="No";
 0.0000 DL="100%";TD="0.5 day";SD="0 day";SF="No";
 0.0162 DL="0%";TD="1 day";SD="0 day";SF="No";
 0.0121 DL="50%";TD="1 day";SD="0 day";SF="No";
 0.0000 DL="100%";TD="1 day";SD="0 day";SF="No";
 0.0951 DL="0%";TD="0 day";SD="0.5 day";SF="No";
 0.0149 DL="50%";TD="0 day";SD="0.5 day";SF="No";
 0.0000 DL="100%";TD="0 day";SD="0.5 day";SF="No";
 0.0162 DL="0%";TD="0.5 day";SD="0.5 day";SF="No";
 0.0121 DL="50%";TD="0.5 day";SD="0.5 day";SF="No";
 0.0000 DL="100%";TD="0.5 day";SD="0.5 day";SF="No";
 0.0162 DL="0%";TD="1 day";SD="0.5 day";SF="No";
 0.0121 DL="50%";TD="1 day";SD="0.5 day";SF="No";
 0.0000 DL="100%";TD="1 day";SD="0.5 day";SF="No";
 0.0951 DL="0%";TD="0 day";SD="1 day";SF="No";
 0.0149 DL="50%";TD="0 day";SD="1 day";SF="No";
 0.0000 DL="100%";TD="0 day";SD="1 day";SF="No";
 0.0162 DL="0%";TD="0.5 day";SD="1 day";SF="No";
 0.0121 DL="50%";TD="0.5 day";SD="1 day";SF="No";
 0.0000 DL="100%";TD="0.5 day";SD="1 day";SF="No";
 0.0162 DL="0%";TD="1 day";SD="1 day";SF="No";
 0.0121 DL="50%";TD="1 day";SD="1 day";SF="No";
 0.0000 DL="100%";TD="1 day";SD="1 day";SF="No";

AL:DL:TD:SF node
 Continued on next page

"PS" "SQ" "SF"

0.05000 PS="Yes";SQ="H";SF="Yes";
 0.22500 PS="Yes";SQ="L";SF="Yes";
 0.20000 PS="Yes";SQ="H";SF="No";
 0.75000 PS="No";SQ="H";SF="No";
 0.02500 PS="Yes";SQ="L";SF="No";
 0.75000 PS="No";SQ="L";SF="No";

"UPS" "SD" "SF" "HAS"

0.06250 UPS="Yes";SD="0 day";HAS="Yes";SF="Yes";
 0.06250 UPS="Yes";SD="0.5 day";HAS="Yes";SF="Yes";
 0.02500 UPS="No";SD="0.5 day";HAS="Yes";SF="Yes";
 0.10000 UPS="No";SD="1 day";HAS="Yes";SF="Yes";
 0.01250 UPS="Yes";SD="0.5 day";HAS="No";SF="Yes";
 0.11250 UPS="Yes";SD="1 day";HAS="No";SF="Yes";
 0.12500 UPS="No";SD="1 day";HAS="No";SF="Yes";
 0.12500 UPS="Yes";SD="0 day";HAS="Yes";SF="No";
 0.12500 UPS="No";SD="0 day";HAS="Yes";SF="No";
 0.12500 UPS="Yes";SD="0 day";HAS="No";SF="No";
 0.12500 UPS="No";SD="0 day";HAS="No";SF="No";

NETWORK RISK: PROPAGATION EXAMPLE
COLLECT PHASE (Continued)

DL:TD:SD:SF node
From previous page

"AF" "DL" "TD" "SF"

0.00707 AF="App Corr";TD="0 day";DL="0%";SF="Yes";
 0.00312 AF="Lockup";TD="0 day";DL="0%";SF="Yes";
 0.04561 AF="OK";TD="0 day";DL="0%";SF="Yes";
 0.00612 AF="App Corr";TD="0.5 day";DL="0%";SF="Yes";
 0.00189 AF="Lockup";TD="0.5 day";DL="0%";SF="Yes";
 0.00096 AF="OK";TD="0.5 day";DL="0%";SF="Yes";
 0.00612 AF="App Corr";TD="1 day";DL="0%";SF="Yes";
 0.00189 AF="Lockup";TD="1 day";DL="0%";SF="Yes";
 0.00096 AF="OK";TD="1 day";DL="0%";SF="Yes";
 0.00312 AF="Lockup";TD="0 day";DL="50%";SF="Yes";
 0.04561 AF="OK";TD="0 day";DL="50%";SF="Yes";
 0.00189 AF="Lockup";TD="0.5 day";DL="50%";SF="Yes";
 0.00096 AF="OK";TD="0.5 day";DL="50%";SF="Yes";
 0.00189 AF="Lockup";TD="1 day";DL="50%";SF="Yes";
 0.00096 AF="OK";TD="1 day";DL="50%";SF="Yes";
 0.06367 AF="App Corr";TD="0 day";DL="100%";SF="Yes";
 0.02498 AF="Lockup";TD="0 day";DL="100%";SF="Yes";
 0.13681 AF="OK";TD="0 day";DL="100%";SF="Yes";
 0.05512 AF="App Corr";TD="0.5 day";DL="100%";SF="Yes";
 0.01516 AF="Lockup";TD="0.5 day";DL="100%";SF="Yes";
 0.00288 AF="OK";TD="0.5 day";DL="100%";SF="Yes";
 0.05512 AF="App Corr";TD="1 day";DL="100%";SF="Yes";
 0.01516 AF="Lockup";TD="1 day";DL="100%";SF="Yes";
 0.00288 AF="OK";TD="1 day";DL="100%";SF="Yes";
 0.03537 AF="App Corr";TD="0 day";DL="0%";SF="No";
 0.02186 AF="Lockup";TD="0.5 day";DL="0%";SF="No";
 0.22802 AF="OK";TD="0 day";DL="0%";SF="No";
 0.03062 AF="App Corr";TD="0.5 day";DL="0%";SF="No";
 0.01326 AF="Lockup";TD="0.5 day";DL="0%";SF="No";
 0.00480 AF="OK";TD="0.5 day";DL="0%";SF="No";
 0.03062 AF="App Corr";TD="1 day";DL="0%";SF="No";
 0.01326 AF="Lockup";TD="1 day";DL="0%";SF="No";
 0.00480 AF="OK";TD="1 day";DL="0%";SF="No";
 0.03537 AF="App Corr";TD="0 day";DL="50%";SF="No";
 0.00937 AF="Lockup";TD="0 day";DL="50%";SF="No";
 0.03062 AF="App Corr";TD="0.5 day";DL="50%";SF="No";
 0.00568 AF="Lockup";TD="0.5 day";DL="50%";SF="No";
 0.03062 AF="App Corr";TD="1 day";DL="50%";SF="No";
 0.00568 AF="Lockup";TD="1 day";DL="50%";SF="No";

"NF" "Hack" "TD"

0.0000 NF="Yes";TD="0 day";Hack="Yes";
 0.1000 NF="No";TD="0 day";Hack="Yes";
 0.2000 NF="Yes";TD="0.5 day";Hack="Yes";
 0.0000 NF="No";TD="0.5 day";Hack="Yes";
 0.2000 NF="Yes";TD="1 day";Hack="Yes";
 0.0000 NF="No";TD="1 day";Hack="Yes";
 0.5000 NF="No";TD="0 day";Hack="No";
 0.0000 NF="No";TD="0.5 day";Hack="No";
 0.0000 NF="No";TD="1 day";Hack="No";

"HAN" "TD" "NF"

0.22500 HAN="Yes";TD="0.5 day";NF="Yes";
 0.02500 HAN="No";TD="0.5 day";NF="Yes";
 0.02500 HAN="Yes";TD="1 day";NF="Yes";
 0.22500 HAN="No";TD="1 day";NF="Yes";
 0.25000 HAN="Yes";TD="0 day";NF="No";
 0.25000 HAN="No";TD="0 day";NF="No";

"V" "Hack" "AF" "FAC" "EUM"

0.01125 V="Yes";EUM="Yes";AF="App Corr";FAC="High";Hack="Yes";
 0.06750 V="No";EUM="Yes";AF="App Corr";FAC="High";Hack="Yes";
 0.00875 V="Yes";EUM="No";AF="App Corr";FAC="High";Hack="Yes";
 0.06750 V="No";EUM="No";AF="App Corr";FAC="High";Hack="Yes";
 0.00125 V="Yes";EUM="Yes";AF="Lockup";FAC="High";Hack="Yes";
 0.04500 V="No";EUM="Yes";AF="Lockup";FAC="High";Hack="Yes";
 0.00375 V="Yes";EUM="No";AF="Lockup";FAC="High";Hack="Yes";
 0.02250 V="No";EUM="No";AF="Lockup";FAC="High";Hack="Yes";
 0.02250 V="No";EUM="No";AF="OK";FAC="High";Hack="Yes";
 0.07875 V="Yes";EUM="Yes";AF="App Corr";FAC="Low";Hack="Yes";
 0.02250 V="No";EUM="Yes";AF="App Corr";FAC="Low";Hack="Yes";
 0.06125 V="Yes";EUM="No";AF="App Corr";FAC="Low";Hack="Yes";
 0.02250 V="No";EUM="No";AF="App Corr";FAC="Low";Hack="Yes";
 0.00875 V="Yes";EUM="Yes";AF="Lockup";FAC="Low";Hack="Yes";
 0.01500 V="No";EUM="Yes";AF="Lockup";FAC="Low";Hack="Yes";
 0.02625 V="Yes";EUM="No";AF="Lockup";FAC="Low";Hack="Yes";
 0.00750 V="No";EUM="No";AF="Lockup";FAC="Low";Hack="Yes";
 0.00750 V="No";EUM="No";AF="OK";FAC="Low";Hack="Yes";
 0.00500 V="Yes";EUM="Yes";AF="App Corr";FAC="High";Hack="No";
 0.01125 V="No";EUM="Yes";AF="App Corr";FAC="Low";Hack="No";
 0.00625 V="Yes";EUM="No";AF="App Corr";FAC="High";Hack="No";
 0.00375 V="Yes";EUM="Yes";AF="Lockup";FAC="High";Hack="No";
 0.03375 V="No";EUM="Yes";AF="Lockup";FAC="High";Hack="No";
 0.00375 V="Yes";EUM="Yes";AF="OK";FAC="High";Hack="No";
 0.06750 V="No";EUM="Yes";AF="OK";FAC="High";Hack="No";
 0.00625 V="Yes";EUM="No";AF="OK";FAC="High";Hack="No";
 0.11250 V="No";EUM="No";AF="OK";FAC="High";Hack="No";
 0.03500 V="Yes";EUM="Yes";AF="App Corr";FAC="Low";Hack="No";
 0.00375 V="No";EUM="Yes";AF="App Corr";FAC="Low";Hack="No";
 0.04375 V="Yes";EUM="No";AF="App Corr";FAC="Low";Hack="No";
 0.02625 V="Yes";EUM="Yes";AF="Lockup";FAC="Low";Hack="No";
 0.01125 V="No";EUM="Yes";AF="Lockup";FAC="Low";Hack="No";
 0.02625 V="Yes";EUM="Yes";AF="OK";FAC="Low";Hack="No";
 0.02250 V="No";EUM="Yes";AF="OK";FAC="Low";Hack="No";
 0.04375 V="Yes";EUM="No";AF="OK";FAC="Low";Hack="No";
 0.03750 V="No";EUM="No";AF="OK";FAC="Low";Hack="No";

"Hack" "AF" "TD"

0.0613 Hack="Yes";TD="0 day";AF="App Corr";
 0.0803 Hack="No";TD="0 day";AF="App Corr";
 0.1225 Hack="Yes";TD="0.5 day";AF="App Corr";
 0.0000 Hack="No";TD="0.5 day";AF="App Corr";
 0.1225 Hack="Yes";TD="1 day";AF="App Corr";
 0.0000 Hack="No";TD="1 day";AF="App Corr";
 0.0190 Hack="Yes";TD="0 day";AF="Lockup";
 0.0435 Hack="No";TD="0 day";AF="Lockup";
 0.0379 Hack="Yes";TD="0.5 day";AF="Lockup";
 0.0000 Hack="No";TD="0.5 day";AF="Lockup";
 0.0379 Hack="Yes";TD="1 day";AF="Lockup";
 0.0000 Hack="No";TD="1 day";AF="Lockup";
 0.0048 Hack="Yes";TD="0 day";AF="OK";
 0.4513 Hack="No";TD="0 day";AF="OK";
 0.0096 Hack="Yes";TD="0.5 day";AF="OK";
 0.0000 Hack="No";TD="0.5 day";AF="OK";
 0.0096 Hack="Yes";TD="1 day";AF="OK";
 0.0000 Hack="No";TD="1 day";AF="OK";

"F" "Hack" "FAC"

0.02500 F="AP";FAC="High";Hack="Yes";
 0.07500 F="PF";FAC="High";Hack="Yes";
 0.12500 F="AP";FAC="Low";Hack="Yes";
 0.20000 F="PF";FAC="Low";Hack="Yes";
 0.22500 F="AP";FAC="High";Hack="No";
 0.17500 F="PF";FAC="High";Hack="No";
 0.12500 F="AP";FAC="Low";Hack="No";
 0.05000 F="PF";FAC="Low";Hack="No";

NETWORK RISK: PROPAGATION EXAMPLE
DISTRIBUTE PHASE

"Cost"	"DL"	"TD"	"SD"
0.49399	Cost="0.0m";TD="0 day";SD="0 day";DL="0%";		
0.05901	Cost="0.0m";TD="0.5 day";SD="0 day";DL="0%";		
0.02529	Cost="0.5m";TD="0.5 day";SD="0 day";DL="0%";		
0.05058	Cost="0.0m";TD="1 day";SD="0 day";DL="0%";		
0.03372	Cost="0.5m";TD="1 day";SD="0 day";DL="0%";		
0.00184	Cost="0.5m";TD="0 day";SD="0.5 day";DL="0%";		
0.00123	Cost="1.0m";TD="0 day";SD="0.5 day";DL="0%";		
0.00025	Cost="0.5m";TD="0.5 day";SD="0.5 day";DL="0%";		
0.00015	Cost="1.0m";TD="0.5 day";SD="0.5 day";DL="0%";		
0.00010	Cost="1.5m";TD="0.5 day";SD="0.5 day";DL="0%";		
0.00020	Cost="0.5m";TD="1 day";SD="0.5 day";DL="0%";		
0.00015	Cost="1.0m";TD="1 day";SD="0.5 day";DL="0%";		
0.00010	Cost="1.5m";TD="1 day";SD="0.5 day";DL="0%";		
0.00005	Cost="2.0m";TD="1 day";SD="0.5 day";DL="0%";		
0.00104	Cost="0.5m";TD="0 day";SD="1 day";DL="0%";		
0.00207	Cost="1.0m";TD="0 day";SD="1 day";DL="0%";		
0.00311	Cost="1.5m";TD="0 day";SD="1 day";DL="0%";		
0.00207	Cost="2.0m";TD="0 day";SD="1 day";DL="0%";		
0.00207	Cost="2.5m";TD="0 day";SD="1 day";DL="0%";		
0.00033	Cost="1.0m";TD="0.5 day";SD="1 day";DL="0%";		
0.00067	Cost="1.5m";TD="0.5 day";SD="1 day";DL="0%";		
0.00033	Cost="2.0m";TD="0.5 day";SD="1 day";DL="0%";		
0.00033	Cost="2.5m";TD="0.5 day";SD="1 day";DL="0%";		
0.00017	Cost="1.0m";TD="1 day";SD="1 day";DL="0%";		
0.00033	Cost="1.5m";TD="1 day";SD="1 day";DL="0%";		
0.00067	Cost="2.0m";TD="1 day";SD="1 day";DL="0%";		
0.00050	Cost="2.5m";TD="1 day";SD="1 day";DL="0%";		
0.04731	Cost="0.0m";TD="0 day";SD="0 day";DL="50%";		
0.03154	Cost="0.5m";TD="0 day";SD="0 day";DL="50%";		
0.01882	Cost="0.0m";TD="0.5 day";SD="0 day";DL="50%";		
0.01882	Cost="0.5m";TD="0.5 day";SD="0 day";DL="50%";		
0.01255	Cost="1.0m";TD="0.5 day";SD="0 day";DL="50%";		
0.01255	Cost="1.5m";TD="0.5 day";SD="0 day";DL="50%";		
0.01882	Cost="0.0m";TD="1 day";SD="0 day";DL="50%";		
0.01255	Cost="0.5m";TD="1 day";SD="0 day";DL="50%";		
0.01255	Cost="1.0m";TD="1 day";SD="0 day";DL="50%";		
0.01255	Cost="1.5m";TD="1 day";SD="0 day";DL="50%";		
0.00627	Cost="2.0m";TD="1 day";SD="0 day";DL="50%";		
0.00134	Cost="0.5m";TD="0 day";SD="0.5 day";DL="50%";		
0.00107	Cost="1.0m";TD="0 day";SD="0.5 day";DL="50%";		
0.00027	Cost="1.5m";TD="0 day";SD="0.5 day";DL="50%";		
0.00005	Cost="0.5m";TD="0.5 day";SD="0.5 day";DL="50%";		
0.00005	Cost="1.0m";TD="0.5 day";SD="0.5 day";DL="50%";		
0.00003	Cost="1.5m";TD="0.5 day";SD="0.5 day";DL="50%";		
0.00003	Cost="2.0m";TD="0.5 day";SD="0.5 day";DL="50%";		
0.00003	Cost="2.5m";TD="0.5 day";SD="0.5 day";DL="50%";		
0.00005	Cost="0.5m";TD="1 day";SD="0.5 day";DL="50%";		
0.00003	Cost="1.0m";TD="1 day";SD="0.5 day";DL="50%";		
0.00003	Cost="1.5m";TD="1 day";SD="0.5 day";DL="50%";		
0.00003	Cost="2.0m";TD="1 day";SD="0.5 day";DL="50%";		
0.00002	Cost="2.5m";TD="1 day";SD="0.5 day";DL="50%";		
0.00090	Cost="1.0m";TD="0 day";SD="1 day";DL="50%";		
0.00271	Cost="1.5m";TD="0 day";SD="1 day";DL="50%";		
0.00362	Cost="2.0m";TD="0 day";SD="1 day";DL="50%";		
0.00181	Cost="2.5m";TD="0 day";SD="1 day";DL="50%";		
0.00005	Cost="1.0m";TD="0.5 day";SD="1 day";DL="50%";		
0.00011	Cost="1.5m";TD="0.5 day";SD="1 day";DL="50%";		
0.00026	Cost="2.0m";TD="0.5 day";SD="1 day";DL="50%";		
0.00011	Cost="2.5m";TD="0.5 day";SD="1 day";DL="50%";		
0.00005	Cost="1.5m";TD="1 day";SD="1 day";DL="50%";		
0.00021	Cost="2.0m";TD="1 day";SD="1 day";DL="50%";		
0.00026	Cost="2.5m";TD="1 day";SD="1 day";DL="50%";		
0.00310	Cost="0.0m";TD="0 day";SD="0 day";DL="100%";		
0.00310	Cost="0.5m";TD="0 day";SD="0 day";DL="100%";		
0.00155	Cost="1.0m";TD="0 day";SD="0 day";DL="100%";		
0.00050	Cost="0.0m";TD="0.5 day";SD="0 day";DL="100%";		
0.00050	Cost="0.5m";TD="0.5 day";SD="0 day";DL="100%";		
0.00075	Cost="1.0m";TD="0.5 day";SD="0 day";DL="100%";		
0.00050	Cost="1.5m";TD="0.5 day";SD="0 day";DL="100%";		
0.00025	Cost="2.0m";TD="0.5 day";SD="0 day";DL="100%";		
0.00025	Cost="0.5m";TD="1 day";SD="0 day";DL="100%";		
0.00050	Cost="1.0m";TD="1 day";SD="0 day";DL="100%";		
0.00075	Cost="1.5m";TD="1 day";SD="0 day";DL="100%";		
0.00050	Cost="2.0m";TD="1 day";SD="0 day";DL="100%";		
0.00050	Cost="2.5m";TD="1 day";SD="0 day";DL="100%";		
0.00496	Cost="0.5m";TD="0 day";SD="0.5 day";DL="100%";		
0.00372	Cost="1.0m";TD="0 day";SD="0.5 day";DL="100%";		
0.00248	Cost="1.5m";TD="0 day";SD="0.5 day";DL="100%";		
0.00124	Cost="2.0m";TD="0 day";SD="0.5 day";DL="100%";		
0.00080	Cost="0.5m";TD="0.5 day";SD="0.5 day";DL="100%";		
0.00080	Cost="1.0m";TD="0.5 day";SD="0.5 day";DL="100%";		
0.00121	Cost="1.5m";TD="0.5 day";SD="0.5 day";DL="100%";		
0.00080	Cost="2.0m";TD="0.5 day";SD="0.5 day";DL="100%";		
0.00040	Cost="2.5m";TD="0.5 day";SD="0.5 day";DL="100%";		
0.00080	Cost="1.0m";TD="1 day";SD="0.5 day";DL="100%";		
0.00121	Cost="1.5m";TD="1 day";SD="0.5 day";DL="100%";		
0.00121	Cost="2.0m";TD="1 day";SD="0.5 day";DL="100%";		
0.00080	Cost="2.5m";TD="1 day";SD="0.5 day";DL="100%";		
0.01256	Cost="1.5m";TD="0 day";SD="1 day";DL="100%";		
0.01674	Cost="2.0m";TD="0 day";SD="1 day";DL="100%";		
0.01256	Cost="2.5m";TD="0 day";SD="1 day";DL="100%";		
0.00272	Cost="1.5m";TD="0.5 day";SD="1 day";DL="100%";		
0.00543	Cost="2.0m";TD="0.5 day";SD="1 day";DL="100%";		
0.00543	Cost="2.5m";TD="0.5 day";SD="1 day";DL="100%";		
0.00272	Cost="2.0m";TD="1 day";SD="1 day";DL="100%";		
0.01087	Cost="2.5m";TD="1 day";SD="1 day";DL="100%";		

"DL"	"TD"	"SD"	"SF"
0.00192	DL="0%";TD="0 day";SD="0 day";SF="Yes";		
0.00168	DL="50%";TD="0 day";SD="0 day";SF="Yes";		
0.00775	DL="100%";TD="0 day";SD="0 day";SF="Yes";		
0.00031	DL="0%";TD="0.5 day";SD="0 day";SF="Yes";		
0.00010	DL="50%";TD="0.5 day";SD="0 day";SF="Yes";		
0.00252	DL="100%";TD="0.5 day";SD="0 day";SF="Yes";		
0.00031	DL="0%";TD="1 day";SD="0 day";SF="Yes";		
0.00010	DL="50%";TD="1 day";SD="0 day";SF="Yes";		
0.00252	DL="100%";TD="1 day";SD="0 day";SF="Yes";		
0.00307	DL="0%";TD="0 day";SD="0.5 day";SF="Yes";		
0.00268	DL="50%";TD="0 day";SD="0.5 day";SF="Yes";		
0.01240	DL="100%";TD="0 day";SD="0.5 day";SF="Yes";		
0.00049	DL="0%";TD="0.5 day";SD="0.5 day";SF="Yes";		
0.00016	DL="50%";TD="0.5 day";SD="0.5 day";SF="Yes";		
0.00402	DL="100%";TD="0.5 day";SD="0.5 day";SF="Yes";		
0.00049	DL="0%";TD="1 day";SD="0.5 day";SF="Yes";		
0.00016	DL="50%";TD="1 day";SD="0.5 day";SF="Yes";		
0.00402	DL="100%";TD="1 day";SD="0.5 day";SF="Yes";		
0.01036	DL="0%";TD="0 day";SD="1 day";SF="Yes";		
0.00905	DL="50%";TD="0 day";SD="1 day";SF="Yes";		
0.04185	DL="100%";TD="0 day";SD="1 day";SF="Yes";		
0.00167	DL="0%";TD="0.5 day";SD="1 day";SF="Yes";		
0.00053	DL="50%";TD="0.5 day";SD="1 day";SF="Yes";		
0.01358	DL="100%";TD="0.5 day";SD="1 day";SF="Yes";		
0.00167	DL="0%";TD="1 day";SD="1 day";SF="Yes";		
0.00053	DL="50%";TD="1 day";SD="1 day";SF="Yes";		
0.01358	DL="100%";TD="1 day";SD="1 day";SF="Yes";		
0.49207	DL="0%";TD="0 day";SD="0 day";SF="No";		
0.07718	DL="50%";TD="0 day";SD="0 day";SF="No";		
0.00000	DL="100%";TD="0 day";SD="0 day";SF="No";		
0.08399	DL="0%";TD="0.5 day";SD="0 day";SF="No";		
0.06263	DL="50%";TD="0.5 day";SD="0 day";SF="No";		
0.00000	DL="100%";TD="0.5 day";SD="0 day";SF="No";		
0.08399	DL="0%";TD="1 day";SD="0 day";SF="No";		
0.06263	DL="50%";TD="1 day";SD="0 day";SF="No";		
0.00000	DL="100%";TD="1 day";SD="0 day";SF="No";		
0.00000	DL="0%";TD="0 day";SD="0.5 day";SF="No";		
0.00000	DL="50%";TD="0 day";SD="0.5 day";SF="No";		
0.00000	DL="100%";TD="0 day";SD="0.5 day";SF="No";		
0.00000	DL="0%";TD="0.5 day";SD="0.5 day";SF="No";		
0.00000	DL="50%";TD="0.5 day";SD="0.5 day";SF="No";		
0.00000	DL="100%";TD="0.5 day";SD="0.5 day";SF="No";		
0.00000	DL="0%";TD="1 day";SD="0.5 day";SF="No";		
0.00000	DL="50%";TD="1 day";SD="0.5 day";SF="No";		
0.00000	DL="100%";TD="1 day";SD="0.5 day";SF="No";		
0.00000	DL="0%";TD="0 day";SD="1 day";SF="No";		
0.00000	DL="50%";TD="0 day";SD="1 day";SF="No";		
0.00000	DL="100%";TD="0 day";SD="1 day";SF="No";		
0.00000	DL="0%";TD="0.5 day";SD="1 day";SF="No";		
0.00000	DL="50%";TD="0.5 day";SD="1 day";SF="No";		
0.00000	DL="100%";TD="0.5 day";SD="1 day";SF="No";		
0.00000	DL="0%";TD="1 day";SD="1 day";SF="No";		
0.00000	DL="50%";TD="1 day";SD="1 day";SF="No";		
0.00000	DL="100%";TD="1 day";SD="1 day";SF="No";		

AL:DL:TD:SF node
 Continued on next page

"PS"	"SQ"	"SF"
0.0250	PS="Yes";SQ="H";SF="Yes";	
0.1125	PS="Yes";SQ="L";SF="Yes";	
0.1000	PS="Yes";SQ="H";SF="No";	
0.3750	PS="No";SQ="H";SF="No";	
0.0125	PS="Yes";SQ="L";SF="No";	
0.3750	PS="No";SQ="L";SF="No";	

"UPS"	"SD"	"SF"	"HAS"
0.0172	UPS="Yes";SD="0 day";HAS="Yes";SF="Yes";		
0.0172	UPS="Yes";SD="0.5 day";HAS="Yes";SF="Yes";		
0.0069	UPS="No";SD="0.5 day";HAS="Yes";SF="Yes";		
0.0275	UPS="No";SD="1 day";HAS="Yes";SF="Yes";		
0.0034	UPS="Yes";SD="0.5 day";HAS="No";SF="Yes";		
0.0309	UPS="Yes";SD="1 day";HAS="No";SF="Yes";		
0.0344	UPS="No";SD="1 day";HAS="No";SF="Yes";		
0.2156	UPS="Yes";SD="0 day";HAS="Yes";SF="No";		
0.2156	UPS="No";SD="0 day";HAS="Yes";SF="No";		
0.2156	UPS="Yes";SD="0 day";HAS="No";SF="No";		
0.2156	UPS="No";SD="0 day";HAS="No";SF="No";		

NETWORK RISK: PROPAGATION EXAMPLE
DISTRIBUTE PHASE (Continued)

DL:TD:SD:SF node
From previous page

"AF" "DL" "TD" "SF"

```

0.00195 AF="App Corr";TD="0 day";DL="0%";SF="Yes";
0.00086 AF="Lockup";TD="0 day";DL="0%";SF="Yes";
0.01254 AF="OK";TD="0 day";DL="0%";SF="Yes";
0.00168 AF="App Corr";TD="0.5 day";DL="0%";SF="Yes";
0.00052 AF="Lockup";TD="0.5 day";DL="0%";SF="Yes";
0.00026 AF="OK";TD="0.5 day";DL="0%";SF="Yes";
0.00168 AF="App Corr";TD="1 day";DL="0%";SF="Yes";
0.00052 AF="Lockup";TD="1 day";DL="0%";SF="Yes";
0.00026 AF="OK";TD="1 day";DL="0%";SF="Yes";
0.00086 AF="Lockup";TD="0 day";DL="50%";SF="Yes";
0.01254 AF="OK";TD="0 day";DL="50%";SF="Yes";
0.00052 AF="Lockup";TD="0.5 day";DL="50%";SF="Yes";
0.00026 AF="OK";TD="0.5 day";DL="50%";SF="Yes";
0.00052 AF="Lockup";TD="1 day";DL="50%";SF="Yes";
0.00026 AF="OK";TD="1 day";DL="50%";SF="Yes";
0.01751 AF="App Corr";TD="0 day";DL="100%";SF="Yes";
0.00687 AF="Lockup";TD="0 day";DL="100%";SF="Yes";
0.03762 AF="OK";TD="0 day";DL="100%";SF="Yes";
0.01516 AF="App Corr";TD="0.5 day";DL="100%";SF="Yes";
0.00417 AF="Lockup";TD="0.5 day";DL="100%";SF="Yes";
0.00079 AF="OK";TD="0.5 day";DL="100%";SF="Yes";
0.01516 AF="App Corr";TD="1 day";DL="100%";SF="Yes";
0.00417 AF="Lockup";TD="1 day";DL="100%";SF="Yes";
0.00079 AF="OK";TD="1 day";DL="100%";SF="Yes";
0.06102 AF="App Corr";TD="0 day";DL="0%";SF="No";
0.03770 AF="Lockup";TD="0 day";DL="0%";SF="No";
0.39334 AF="OK";TD="0 day";DL="0%";SF="No";
0.05283 AF="App Corr";TD="0.5 day";DL="0%";SF="No";
0.02288 AF="Lockup";TD="0.5 day";DL="0%";SF="No";
0.00828 AF="OK";TD="0.5 day";DL="0%";SF="No";
0.05283 AF="App Corr";TD="1 day";DL="0%";SF="No";
0.02288 AF="Lockup";TD="1 day";DL="0%";SF="No";
0.00828 AF="OK";TD="1 day";DL="0%";SF="No";
0.06102 AF="App Corr";TD="0 day";DL="50%";SF="No";
0.01616 AF="Lockup";TD="0 day";DL="50%";SF="No";
0.05283 AF="App Corr";TD="0.5 day";DL="50%";SF="No";
0.00981 AF="Lockup";TD="0.5 day";DL="50%";SF="No";
0.05283 AF="App Corr";TD="1 day";DL="50%";SF="No";
0.00981 AF="Lockup";TD="1 day";DL="50%";SF="No";
    
```

"NF" "Hack" "TD"

```

0.0000 NF="Yes";TD="0 day";Hack="Yes";
0.0850 NF="No";TD="0 day";Hack="Yes";
0.1700 NF="Yes";TD="0.5 day";Hack="Yes";
0.0000 NF="No";TD="0.5 day";Hack="Yes";
0.1700 NF="Yes";TD="1 day";Hack="Yes";
0.0000 NF="No";TD="1 day";Hack="Yes";
0.5750 NF="No";TD="0 day";Hack="No";
0.0000 NF="No";TD="0.5 day";Hack="No";
0.0000 NF="No";TD="1 day";Hack="No";
    
```

"HAN" "TD" "NF"

```

0.1530 HAN="Yes";TD="0.5 day";NF="Yes";
0.0170 HAN="No";TD="0.5 day";NF="Yes";
0.0170 HAN="Yes";TD="1 day";NF="Yes";
0.1530 HAN="No";TD="1 day";NF="Yes";
0.3300 HAN="Yes";TD="0 day";NF="No";
0.3300 HAN="No";TD="0 day";NF="No";
    
```

"V" "Hack" "AF" "FAC" "EUM"

```

0.00090 V="Yes";EUM="Yes";AF="App corr";FAC="High";Hack="Yes";
0.00540 V="No";EUM="Yes";AF="App Corr";FAC="High";Hack="Yes";
0.00630 V="Yes";EUM="No";AF="App Corr";FAC="High";Hack="Yes";
0.04860 V="No";EUM="No";AF="App Corr";FAC="High";Hack="Yes";
0.00010 V="Yes";EUM="Yes";AF="Lockup";FAC="High";Hack="Yes";
0.00360 V="No";EUM="Yes";AF="Lockup";FAC="High";Hack="Yes";
0.00270 V="Yes";EUM="No";AF="Lockup";FAC="High";Hack="Yes";
0.01620 V="No";EUM="No";AF="Lockup";FAC="High";Hack="Yes";
0.01620 V="No";EUM="No";AF="OK";FAC="High";Hack="Yes";
0.12285 V="Yes";EUM="Yes";AF="App Corr";FAC="Low";Hack="Yes";
0.03510 V="No";EUM="Yes";AF="App Corr";FAC="Low";Hack="Yes";
0.06370 V="Yes";EUM="No";AF="App Corr";FAC="Low";Hack="Yes";
0.02340 V="No";EUM="No";AF="App Corr";FAC="Low";Hack="Yes";
0.01365 V="Yes";EUM="Yes";AF="Lockup";FAC="Low";Hack="Yes";
0.00340 V="No";EUM="Yes";AF="Lockup";FAC="Low";Hack="Yes";
0.02730 V="Yes";EUM="No";AF="Lockup";FAC="Low";Hack="Yes";
0.00780 V="No";EUM="No";AF="Lockup";FAC="Low";Hack="Yes";
0.00780 V="No";EUM="No";AF="OK";FAC="Low";Hack="Yes";
0.00160 V="Yes";EUM="Yes";AF="App Corr";FAC="High";Hack="No";
0.00360 V="No";EUM="Yes";AF="App Corr";FAC="High";Hack="No";
0.01800 V="Yes";EUM="No";AF="App Corr";FAC="Low";Hack="No";
0.00120 V="Yes";EUM="Yes";AF="Lockup";FAC="High";Hack="No";
0.01080 V="No";EUM="Yes";AF="Lockup";FAC="High";Hack="No";
0.00120 V="Yes";EUM="Yes";AF="OK";FAC="High";Hack="No";
0.02160 V="No";EUM="Yes";AF="OK";FAC="High";Hack="No";
0.01800 V="Yes";EUM="No";AF="OK";FAC="High";Hack="No";
0.32400 V="No";EUM="No";AF="OK";FAC="High";Hack="No";
0.02940 V="Yes";EUM="Yes";AF="App Corr";FAC="Low";Hack="No";
0.00315 V="No";EUM="Yes";AF="App Corr";FAC="Low";Hack="No";
0.02450 V="Yes";EUM="No";AF="App Corr";FAC="Low";Hack="No";
0.02205 V="Yes";EUM="Yes";AF="Lockup";FAC="Low";Hack="No";
0.00945 V="No";EUM="Yes";AF="Lockup";FAC="Low";Hack="No";
0.02205 V="Yes";EUM="Yes";AF="OK";FAC="Low";Hack="No";
0.01890 V="No";EUM="Yes";AF="OK";FAC="Low";Hack="No";
0.02450 V="Yes";EUM="No";AF="OK";FAC="Low";Hack="No";
0.02100 V="No";EUM="No";AF="OK";FAC="Low";Hack="No";
    
```

"Hack" "AF" "TD"

```

0.0613 Hack="Yes";TD="0 day";AF="App Corr";
0.0803 Hack="No";TD="0 day";AF="App Corr";
0.1225 Hack="Yes";TD="0.5 day";AF="App Corr";
0.0000 Hack="No";TD="0.5 day";AF="App Corr";
0.1225 Hack="Yes";TD="1 day";AF="App Corr";
0.0000 Hack="No";TD="1 day";AF="App Corr";
0.0190 Hack="Yes";TD="0 day";AF="Lockup";
0.0435 Hack="No";TD="0 day";AF="Lockup";
0.0379 Hack="Yes";TD="0.5 day";AF="Lockup";
0.0000 Hack="No";TD="0.5 day";AF="Lockup";
0.0379 Hack="Yes";TD="1 day";AF="Lockup";
0.0000 Hack="No";TD="1 day";AF="Lockup";
0.0048 Hack="Yes";TD="0 day";AF="OK";
0.4513 Hack="No";TD="0 day";AF="OK";
0.0096 Hack="Yes";TD="0.5 day";AF="OK";
0.0000 Hack="No";TD="0.5 day";AF="OK";
0.0096 Hack="Yes";TD="1 day";AF="OK";
0.0000 Hack="No";TD="1 day";AF="OK";
    
```

"F" "Hack" "FAC"

```

0.0250 F="AP";FAC="High";Hack="Yes";
0.0750 F="PF";FAC="High";Hack="Yes";
0.1250 F="AP";FAC="Low";Hack="Yes";
0.2000 F="PF";FAC="Low";Hack="Yes";
0.2250 F="AP";FAC="High";Hack="No";
0.1750 F="PF";FAC="High";Hack="No";
0.1250 F="AP";FAC="Low";Hack="No";
0.0500 F="PF";FAC="Low";Hack="No";
    
```

Network Risk: Simulation Results

Sample size 52

AF		
	Counts	Prob
App Corr	7	13.5%
Lock Up	2	3.8%
OK	43	82.7%
	52	100.0%

DL		
	Counts	Prob
0%	45	86.5%
50%	4	7.7%
100%	3	5.8%
	52	100.0%

TD		
	Counts	Prob
0 day	49	94.2%
0.5 day	3	5.8%
1 day	0	0.0%
	52	100.0%

SD		
	Counts	Prob
0 day	50	96.2%
0.5 day	2	3.8%
1 day	0	0.0%
	52	100.0%

EUM		
	Counts	Prob
Yes	4	7.7%
No	48	92.3%
	52	100.0%

HAN		
	Counts	Prob
Yes	52	100.0%
No	0	0.0%
	52	100.0%

HAS		
	Counts	Prob
Yes	52	100.0%
No	0	0.0%
	52	100.0%

Hack		
	Counts	Prob
Yes	6	11.5%
No	46	88.5%
	52	100.0%

NF		
	Counts	Prob
Yes	3	5.8%
No	49	94.2%
	52	100.0%

PS		
	Counts	Prob
Yes	13	25.0%
No	39	75.0%
	52	100.0%

SF		
	Counts	Prob
Yes	5	9.6%
No	47	90.4%
	52	100.0%

UPS		
	Counts	Prob
Yes	52	100.0%
No	0	0.0%
	52	100.0%

V		
	Counts	Prob
Yes	4	7.7%
No	48	92.3%
	52	100.0%

F		
	Counts	Prob
AP	52	100.0%
PF	0	0.0%
	52	100.0%

FAC		
	Counts	Prob
High	52	100.0%
Low	0	0.0%
	52	100.0%

SQ		
	Counts	Prob
High	52	100.0%
Low	0	0.0%
	52	100.0%

Cost		
	Counts	Prob
0.0m	49	94.23%
0.5m	2	3.85%
1.0m	0	0.00%
1.5m	0	0.00%
2.0m	1	1.92%
2.5m	0	0.00%
	52	100.00%

Sample size 9999

AF		
	Counts	Prob
App Corr	1027	10.27%
Lock Up	473	4.73%
OK	8499	85.00%
	9999	100.00%

DL		
	Counts	Prob
0%	8991	89.92%
50%	688	6.88%
100%	320	3.20%
	9999	100.00%

TD		
	Counts	Prob
0 day	9280	92.81%
0.5 day	643	6.43%
1 day	76	0.76%
	9999	100.00%

SD		
	Counts	Prob
0 day	9740	97.41%
0.5 day	259	2.59%
1 day	0	0.00%
	9999	100.00%

EUM		
	Counts	Prob
Yes	949	9.5%
No	9050	90.5%
	9999	100.0%

HAN		
	Counts	Prob
Yes	9999	100.0%
No	0	0.0%
	9999	100.0%

HAS		
	Counts	Prob
Yes	9999	100.0%
No	0	0.0%
	9999	100.0%

Hack		
	Counts	Prob
Yes	899	9.0%
No	9100	91.0%
	9999	100.0%

NF		
	Counts	Prob
Yes	719	7.2%
No	9280	92.8%
	9999	100.0%

PS		
	Counts	Prob
Yes	2403	24.0%
No	7596	76.0%
	9999	100.0%

SF		
	Counts	Prob
Yes	501	5.0%
No	9498	95.0%
	9999	100.0%

UPS		
	Counts	Prob
Yes	9999	100.0%
No	0	0.0%
	9999	100.0%

V		
	Counts	Prob
Yes	963	9.6%
No	9036	90.4%
	9999	100.0%

F		
	Counts	Prob
AP	9999	100.0%
PF	0	0.0%
	9999	100.0%

FAC		
	Counts	Prob
High	9999	100.0%
Low	0	0.0%
	9999	100.0%

SQ		
	Counts	Prob
High	9999	100.0%
Low	0	0.0%
	9999	100.0%

Cost		
	Counts	Prob
0.0m	9186	91.87%
0.5m	529	5.29%
1.0m	162	1.62%
1.5m	97	0.97%
2.0m	24	0.24%
2.5m	1	0.01%
	9999	100.00%

Network Risk: Calculations for Parent-Child Monitor for NF Node

With learning:

m	Data NF	α_y	α_n	$E(\theta_y)$	$E(\theta_n)$	$-\log \theta_y$	$-\log \theta_n$	S_m	S	E_m	Var_m	$\sum E_m$	$\sum Var_m$	Test Statistic
0	Prior	4.32	1.08	0.800	0.200	0.223	1.609							
1	Yes	5.32	1.08	0.831	0.169	0.185	1.779	0.223	0.223	0.500	0.307	0.500	0.307	-0.500
2	No	5.32	2.08	0.719	0.281	0.330	1.269	1.779	2.002	0.454	0.357	0.954	0.664	1.286
3	No	5.32	3.08	0.633	0.367	0.457	1.003	1.269	3.272	0.594	0.178	1.548	0.842	1.878
4	Yes	6.32	3.08	0.672	0.328	0.397	1.116	0.457	3.728	0.657	0.069	2.205	0.912	1.595
5	Yes	7.32	3.08	0.704	0.296	0.351	1.217	0.397	4.125	0.633	0.114	2.838	1.026	1.271
6	No	7.32	4.08	0.642	0.358	0.443	1.028	1.217	5.342	0.608	0.156	3.446	1.182	1.745
7	No	7.32	5.08	0.590	0.410	0.527	0.892	1.028	6.370	0.652	0.079	4.098	1.260	2.024
8	No	7.32	6.08	0.546	0.454	0.605	0.790	0.892	7.262	0.677	0.032	4.774	1.293	2.188
9	No	7.32	7.08	0.508	0.492	0.677	0.710	0.790	8.052	0.689	0.009	5.463	1.301	2.270
10	No	7.32	8.08	0.475	0.525	0.744	0.645	0.710	8.762	0.693	0.000	6.156	1.301	2.284
11	Yes	8.32	8.08	0.507	0.493	0.679	0.708	0.744	9.506	0.692	0.002	6.848	1.304	2.328
12	No	8.32	9.08	0.478	0.522	0.738	0.650	0.708	10.214	0.693	0.000	7.541	1.304	2.341
13	No	8.32	10.08	0.452	0.548	0.794	0.602	0.650	10.864	0.692	0.002	8.233	1.306	2.302
14	Yes	9.32	10.08	0.480	0.520	0.733	0.655	0.794	11.658	0.689	0.009	8.922	1.315	2.386
15	Yes	10.32	10.08	0.506	0.494	0.681	0.705	0.733	12.391	0.692	0.002	9.614	1.317	2.420
16	No	10.32	11.08	0.482	0.518	0.729	0.658	0.705	13.096	0.693	0.000	10.308	1.317	2.430

Where $\theta_y = P(NF = Yes | Hack = Yes)$ and $\theta_n = P(NF = No | Hack = Yes)$

Table A4.1 Penalties and test statistics for model with learning

With learning:

m	Data NF	$E(\theta_y)$	$E(\theta_n)$	$-\log \theta_y$	$-\log \theta_n$	S_m	S	E_m	Var_m	$\sum E_m$	$\sum Var_m$	Test Statistic
0	Prior	0.800	0.200	0.223	1.609							
1	Yes	0.800	0.200	0.223	1.609	0.223	0.223	0.500	0.307	0.500	0.307	- 0.500
2	No	0.800	0.200	0.223	1.609	1.609	1.833	0.500	0.307	1.001	0.615	1.061
3	No	0.800	0.200	0.223	1.609	1.609	3.442	0.500	0.307	1.501	0.922	2.021
4	Yes	0.800	0.200	0.223	1.609	0.223	3.665	0.500	0.307	2.002	1.230	1.500
5	Yes	0.800	0.200	0.223	1.609	0.223	3.888	0.500	0.307	2.502	1.537	1.118
6	No	0.800	0.200	0.223	1.609	1.609	5.498	0.500	0.307	3.002	1.845	1.837
7	No	0.800	0.200	0.223	1.609	1.609	7.107	0.500	0.307	3.503	2.152	2.457
8	No	0.800	0.200	0.223	1.609	1.609	8.717	0.500	0.307	4.003	2.460	3.005
9	No	0.800	0.200	0.223	1.609	1.609	10.326	0.500	0.307	4.504	2.767	3.500
10	No	0.800	0.200	0.223	1.609	1.609	11.935	0.500	0.307	5.004	3.075	3.953
11	Yes	0.800	0.200	0.223	1.609	0.223	12.159	0.500	0.307	5.504	3.382	3.618
12	No	0.800	0.200	0.223	1.609	1.609	13.768	0.500	0.307	6.005	3.690	4.041
13	No	0.800	0.200	0.223	1.609	1.609	15.378	0.500	0.307	6.505	3.997	4.438
14	Yes	0.800	0.200	0.223	1.609	0.223	15.601	0.500	0.307	7.006	4.305	4.143
15	Yes	0.800	0.200	0.223	1.609	0.223	15.824	0.500	0.307	7.506	4.612	3.873
16	No	0.800	0.200	0.223	1.609	1.609	17.433	0.500	0.307	8.006	4.920	4.250

Table A4.2 Penalties and test statistics for model without learning

m	Sm		Bayes Factor	Test Statistic	
	With learning	Without learning		With learning	Without learning
1	0.223	0.223	0.000	-0.500	-0.500
2	2.002	1.833	-0.170	1.286	1.061
3	3.272	3.442	0.170	1.878	2.021
4	3.728	3.665	-0.063	1.595	1.500
5	4.125	3.888	-0.237	1.271	1.118
6	5.342	5.498	0.156	1.745	1.837
7	6.370	7.107	0.737	2.024	2.457
8	7.262	8.717	1.455	2.188	3.005
9	8.052	10.326	2.274	2.270	3.500
10	8.762	11.935	3.173	2.284	3.953
11	9.506	12.159	2.653	2.328	3.618
12	10.214	13.768	3.554	2.341	4.041
13	10.864	15.378	4.513	2.302	4.438
14	11.658	15.601	3.943	2.386	4.143
15	12.391	15.824	3.433	2.420	3.873
16	13.096	17.433	4.337	2.430	4.250

Table A4.3 Comparison of model alternatives