



Operational risk modelling and organizational learning in structured finance operations: a Bayesian network approach

Andrew Sanford^{1*} and Imad Moosa²

¹Monash University, Victoria, Australia; and ²RMIT University, Victoria, Australia

This paper describes the development of a tool, based on a Bayesian network model, that provides *posteriori* predictions of operational risk events, aggregate operational loss distributions, and Operational Value-at-Risk, for a structured finance operations unit located within one of Australia's major banks. The Bayesian network, based on a previously developed causal framework, has been designed to model the smaller and more frequent, attritional operational loss events. Given the limited availability of risk factor event information and operational loss data, we rely on the elicitation of subjective probabilities, sourced from domain experts. Parameter sensitivity analysis is performed to validate and check the model's robustness against the beliefs of risk management and operational staff. To ensure that the domain's evolving risk profile is captured through time, a formal approach to organizational learning is investigated that employs the automatic parameter adaption features of the Bayesian network model. A hypothetical case study is then described to demonstrate model adaption and the application of the tool to operational loss forecasting by a business unit risk manager.

Journal of the Operational Research Society (2015) 66(1), 86–115. doi:10.1057/jors.2013.49

Published online 11 December 2013

Keywords: banking; finance; operational risk; probabilistic methods; artificial intelligence; Bayesian networks

Introduction

This paper presents a case study describing the development of an operational risk tool that supports a risk manager in the measurement, monitoring, reporting and control of operational risks at a local business unit level. The focus is on the smaller, more frequent, attritional loss events, rather than the larger, less frequent and potentially catastrophic events. The tool has been developed for a functioning business unit, Structured Finance Operations (SFO), located within the wholesale banking division of a major Australian commercial bank. The unit's role is to provide operational and transactional support for the portfolio of structured finance products provided to the bank's corporate clients. A core component of the tool developed in this paper is a Bayesian network model that encapsulates the probabilistic and causal features of the domain. The design of the model's network structure was developed and described previously in Sanford and Moosa (2012).

Although it may be argued that the modelling and control of larger, less frequent loss events is a more pressing priority, controlling smaller failures may also have its

benefits. When large operational failures occur, it is not uncommon to find that the cause or causes are subsequently traced to a series of smaller, innocuous contingent failures. Therefore, tools designed to support operational risk management at the local level may, in addition to minimizing the smaller more frequent losses, assist in reducing the probability of larger events.

The operational risk tool described in this paper was developed to assist a risk manager to maintain high situational awareness of a local domain and its risk drivers. A key feature of the tool is the inclusion of automated model adaption that allows a domain's changing risk profile to be captured in an efficient, low cost and timely manner, and which supports both individual and organizational learning. An important output from the tool comprises the forecasts over varying time horizons of aggregate loss distributions and Operational Value-at-Risk (OpVaR). Use of the tool also improves risk communications within the organization by making operational risk, subsequent losses and their key drivers more explicit.

Basel II and operational risk

Prior to the 1990s, the operational risks inherent in financial institutions were considered minor in comparison

*Correspondence: Andrew Sanford, Department of Accounting and Finance, Faculty of Business and Economics, Monash University, Clayton, Victoria 3800, Australia.

with their credit and market risks. Management resources for the measurement and control of operational risks were therefore allocated accordingly. But with the advent of the 1990s, and the arrival of new information and communications technologies, financial innovation, and increasing globalization, the prevailing view of the significance of operational risk required a reassessment. The significance of operational risk within this new environment was dramatically illustrated in 1995, when a single derivatives trader operating out of Singapore caused the 230-year-old London merchant bank, Barings, to collapse. Three years after this event, a report by the Basel Committee on Bank Supervision (BCBS) noted that although the awareness of operational risk among bank boards and senior management had increased, the measurement and monitoring systems were still immature, with many conceptual and data hurdles to overcome. In June 2004, BCBS introduced a new supervisory framework, replacing the previous model that had been in place since 1988. Known as Basel II, this new framework introduced for the first time operational risk as a distinct bank risk category, alongside that of market and credit risk.

The Basel Committee defines operational risk as ‘... the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events’. Set out within Basel II were three measurement methodologies that banks could adopt to calculate their operational risk capital requirements. These included the Basic Indicator approach, the Standardized approach, and the Advanced Measurement Approach (AMA). The methodology most readily adopted by major, internationally active banks was the AMA, which allowed financial institutions the discretion to develop and implement their own internal operational risk capital allocation models, provided that the models were approved by their local prudential authority.

Underlying the changes introduced in Basel II was the need for central bankers and prudential supervisors to be confident that the banks under their authority would maintain sufficient capital to absorb severe loss events regardless of origin, whether due to credit, market or operational risk. Although the implementation of Basel II was completed by 2008, the need for continued improvements in operational risk models and management systems remains. Serious operational loss events are regularly reported for example (Société Générale, 2008), while the funds tied up in risk capital allocations represent a substantial cost to institutions. It remains important therefore that organizations seek to understand and reduce their risks in order to reduce not only loss events, but the capital costs associated with absorbing them.

Following the global financial crisis, the BCBS introduced new rules under what is known as Basel III. New provisions have been put in place to deal with liquidity and

leverage, which were the reasons for the failure of financial institutions such as Bear Stearns and Northern Rock. On the way to implementing Basel III an intermediate accord, Basel 2.5, was implemented in January 2012 (see, eg, Moosa, 2010, 2011, 2012; Moosa and Burns, 2012). In this intermediate accord, emphasis was shifted back to market risk, perhaps due to the belief that the losses incurred during the global financial crisis were mostly market losses. However, the new emphasis on liquidity risk is one type of settlement risk which is a type of operational risk. Furthermore, massive operational losses are announced as frequently as ever. In June 2012, Barclay Bank was fined some \$400 million for manipulating LIBOR.

Background and motivation

Driving the development of operational risk models to date has been the need for financial institutions to improve their operational risk capital allocation and insurance coverage. Modelling has been focused on aggregate losses at the firm or portfolio level, using actuarial-based models similar to those in accident and property loss insurance. Lacking the necessary organizational detail however, these models are less helpful when the focus shifts from capital allocation to operational risk control. For control purposes, models must encapsulate greater systems detail, including the key risk drivers of operational loss events. Such models would be useful for local risk managers, providing loss predictions, as well as the loss attribution of specific key risk drivers and their causal relationships. When addressing issues of task and process control, risk mitigation and operational improvements, such information would assist in the prioritization and directing of resources.

The systems level approach described in this paper draws from systems analysis and probabilistic risk analysis (PRA) techniques (see Bedford and Cooke, 2001). These techniques are used to build a ‘bottom-up’ model of the internal systems and human factors within an operational domain located within a major Australian bank. The experience of modelling at the systems level has highlighted many issues, particularly those related to model complexity, data availability, causality, information gain, and model adaption.

A systems level model requires the incorporation of greater domain information. With this, larger numbers of variables and their associations means increased model complexity. This demand for additional data also means that it is often necessary to augment the limited data with subjective information provided by local domain experts. A need for greater domain detail also makes accessing to the specifics of an organization’s internal operations necessary, increasing the importance of local cooperation and coordination between a model’s development team and the domain personnel, which potentially can introduce

additional political and cultural complexities to the model's development.

The modelling task complexity can also be increased as a result of the technical and task specifics of the environment. Modern financial institutions for example rely increasingly on sophisticated socio-technical systems to manage their business and commercial processes, distributed both locally and globally. Such domains rely on the use of information and communications technologies, and the sophisticated knowledge and skills of their staff, to ensure efficient operations. To measure the performances of people, machines and their interfaces, the system level model needs to accommodate a mixture of both quantitative and qualitative performance measures. In such environments, the coupling between, and the latency of certain variables can make their state measurement difficult. Software bugs and psychological states for example are difficult to measure as they are not readily observable. Components that are highly coupled make measurement of a system's state difficult as it may require multiple contemporaneous observations.

A key motivation for the development of a systems level model is to extend operational risk models beyond risk capital allocation and towards operational risk intervention and control. This requires a deeper understanding of causation in operational risk events that are necessary for the implementation of control. A systems level model needs to incorporate the antecedent events and causal pathways of operational losses. The antecedents, or key risk drivers, include events based on human, system or processes failures, either originating internally within the business unit or externally. A system level model therefore needs to be sufficiently flexible to accommodate a diverse set of potential risk factors. Isolating and measuring the individual risk factors, when complex causal interactions or latency exist, can make modelling them difficult. Risk drivers such as human factors can be particularly difficult to model.

The emphasis on causal relationships is a departure from the existing statistical loss models where correlation, not causation, has been the major concern. Statisticians have traditionally treated with suspicion the causal interpretation of their models when derived from non-experimental or observational data. To operationalize the concept of causation within the systems level model, we use a 'manipulationist' perspective, similar to that used in constructing fault and event trees.¹

A key component of the tool developed in this research is its automated model adaption facility. The incorporation

¹We use a manipulability account of causation whereby, if one intervenes within a system to change the state of an object A, and this results in a change in the state of another object B, then A's state is said to cause object B's state. However, if one intervenes within a system to change the state of an object B, but object A's state remains unchanged, then object B's state does not cause object A's state.

of an automated adaption facility makes the model's updating more efficient and timelier, reducing its reliance on the subjective judgements of experts. This then reduces the demands on costly operations staff, reducing the tools maintenance costs. Keeping the model updated in a consistent and timely manner in such a complex domain as SFO would be difficult without automation. A more efficient integration of new domain data with the existing model will ensure timelier benefits and model validity even in a domain whose risk profile is changing.

Of importance in determining the successful adoption of any information tool is the level of information gain that the targeted end users perceive from using the tool. Information gain is defined here as the information benefits achieved for users, less their costs, relative to the users' existing information sources. In the case of SFO, the existing information sources for operational risk include the formal categories of internal audit, the current operational risk reporting, and informal sources such as the tacit operational risk knowledge and practical experience embedded within the risk manager, operational staff and organizational networks. While the information benefits to a local risk manager have been detailed above, the costs of achieving these benefits have not. The main costs relate to the resources necessary to maintain the model. Maintenance of the model includes regular domain data capture and entry, domain monitoring, and report generation carried out by local staff. The burden of such costs can be more readily dispersed if the benefits of the tools outputs are relevant to as wider audience as possible. The model's summary and aggregate outputs, its ability to communicate and quantify risks, and its ability to adapt to a changing risk profile, suggest a wider interest beyond the local domain. The bank's broader operational risk function and its internal audit function are immediate beneficiaries, while its ability to provide information on local operational risk-adjusted returns would be of great interest to the bank's senior management.

Literature survey

Bayesian network models have and continue to be applied to a diverse set of problem domains. Examples include transportation (Trucco *et al*, 2008), systems dependability (Sigurdsson *et al*, 2001; Neil *et al*, 2008), infrastructure (Willems *et al*, 2005), medical and health-care provision (van der Gaag *et al*, 2002; Lucas *et al*, 2004; Cornalba, 2009), environmental modelling (Bromley *et al*, 2005; Uusitalo, 2007), legal/evidential reasoning (Kadane and Schum, 1996), forensic science (Taroni *et al*, 2006), venture capital decision making (Kemmerer *et al*, 2002), project management (Khodakarami *et al*, 2007), customer service delivery (Anderson *et al*, 2004), new product development (Cooper, 2000), traffic accident modelling (Davis, 2003)

and national security and terrorist threats (Paté-Cornell and Guikema, 2002). A common use for Bayesian networks in many of these developments is the facilitation of reasoning and decision making under uncertainty.

Although the extant literature on Bayesian network applications to operational risk is still small, a number of papers are discussed in Sanford and Moosa (2012). These include Alexander (2000, 2003), Cornalba and Giudici (2004), Neil *et al* (2005), Adusei-Poku *et al* (2007), Cowell *et al* (2007), Mittnik and Starobinskaya (2010), Bonafede and Giudici (2007), Moosa (2008), Neil *et al* (2009), and Aquaro *et al* (2010).

Although the number of specific papers related to operational risk and Bayesian networks is small, applications to allied research fields such as quality, safety and reliability engineering have also produced models and tools that are relevant and adaptable to the problems of operational risk within financial institutions. Examples of these include Fenton *et al* (1998), Fenton and Neil (2000), Neil *et al* (2000, 2001), Bobbio *et al* (2001), Sigurdsson *et al* (2001), Bobbio *et al* (2003); and more recently, Boudali and Dugan (2005), Weber and Jouffe (2006), Wilson and Huzurbazar (2007), Neil *et al* (2008), Marquez *et al* (2010), Groth *et al* (2010), and Nordgård and Sand (2010).

As part of the development of the Bayesian network, it was necessary for the purpose of parameterizing the model to elicit the subjective beliefs of local domain experts. A great deal of material exists covering expert elicitation; however, Granger-Morgan and Henrion (1990) provide an excellent general overview. A more specific discussion covering Bayesian network development using expert elicitation is found in Renooij (2001). Kjærulff and Madsen (2008) provide a practical discussion on the construction and refinement of Bayesian network models, using a combination of domain expert elicitation, sensitivity analysis and Bayesian network adaption. Other Bayesian network development methodologies that incorporate expert elicitation are also discussed in Korb and Nicholson (2004) and Woodberry *et al* (2004). An extensive coverage of many of the cognitive biases associated with elicitation can be found in Baron (2008). Structured frameworks to expert elicitation are described in Cooke and Goossens (1999, 2004).

A model validation technique commonly used in Bayesian network development is sensitivity analysis. Two distinctive forms of sensitivity analysis are available. Parameter sensitivity analysis, which measures the sensitivity of output node probabilities $\Pr(X_i)$ and $\Pr(X_i | e)$ to changes in model CPT parameters, θ , and evidence sensitivity analysis, which measures the sensitivity of output node probabilities to changes in the domain observations or evidence e . Coupé *et al* (1999), Coupé and van der Gaag (2002), van der Gaag *et al* (2002, 2007), Bednarski *et al* (2004) provide discussions of Bayesian network model validation using sensitivity analysis.

Manual construction of all but the most simple Bayesian network model can be very costly in terms of scarce human resources and expertise. Algorithms have been developed to reduce these costs by allowing a model to be automatically constructed and parameterized using the available domain data. Such cost and benefits warrant that in the future, more automated adaption opportunities for both structural and parameter estimation are employed. Neapolitan (2004) provides a broad coverage of structural adaption and parameter updating algorithms, while Koller and Friedman (2009) cover many of the most recent technical developments in Bayesian networks, including details of model and parameter adaption.

Bayesian networks

To assist in making further technical discussions clear and consistent, a formal description of a discrete Bayesian network model is presented, and several notational conventions introduced. We begin by describing a discrete Bayesian network model, denoted \mathbf{B} , as comprising two distinct components: \mathbf{G} a qualitative component and $\mathbf{\Omega}$ a quantitative component. $\mathbf{\Omega}$ is represented as an $(n \times 1)$ vector, $\mathbf{\Omega} = \{\mathbf{\Omega}_1, \mathbf{\Omega}_2, \dots, \mathbf{\Omega}_n\}'$, where each $\mathbf{\Omega}_i$ contains the individual parameter sets defining n nodes, X_i , that comprise the variables within the Bayesian network model. For a discrete Bayesian network model, all nodes X_i have a finite set of j distinct states.

The \mathbf{G} component contains details that describe the Bayesian network model's network structure, which captures the manner in which each node X_i within the model relate. Graphically, this structure takes the form of a directed acyclic graph (DAG) within which the relevance, conditional dependence/independence, or causal influences between nodes are encoded. Any given node X_i is connected to other nodes by sets of directed arcs. When an arc proceeds from one node directly to a second node, the first node is defined as the parent of the second node and the second node is described as a child of the parent. A node may have no parents or one or more parents. Throughout the paper, we denote these sets of parent nodes for any given node X_i as $\text{Pa}(X_i)$, and we denote the k th state configuration of these same parents as $\text{Pa}_k(X_i)$. If a child node has two parent nodes, with each parent having two possible states, then the number of possible parent configurations is four.

The parameters of node X_i , which are conditional on the k th state configuration of the parent nodes, are denoted as $p(X_i | \text{Pa}_k(X_i))$, which can be represented as a $(j \times 1)$ vector $(\theta_1, \theta_2, \dots, \theta_j)'_{ik}$. The vector subscript ik indicates that these parameters relate specifically to states of node X_i , when its parent nodes are in their k th state configuration. The complete set of probability parameters for node X_i is denoted as $\Theta_i \in \mathbf{\Omega}_i$, which is a $(j \times k)$ matrix referred to as

X_i 's conditional probability table (CPT).

$$\Theta_i = \begin{bmatrix} \theta_{11} & \theta_{12} & \cdots & \theta_{1k} \\ \theta_{21} & \theta_{22} & \vdots & \theta_{2k} \\ \vdots & \vdots & \vdots & \vdots \\ \theta_{j1} & \theta_{j2} & \cdots & \theta_{jk} \end{bmatrix}_i \quad (1)$$

When node X_i has no parents, then $k=1$ and Θ_i is a $(j \times 1)$ vector. Each parameter in node X_i 's CPT, as shown in Equation (1), is real valued on a closed interval $[0, 1]$. Each parameter value in Θ_i represents the probability that X_i is in its j th state, conditional on its parents being in their k th state configuration. Therefore, each parameter in Equation (1) represents $\theta_{ijk} = p(X_i = j | \text{Pa}_k(X_i))$. Since each θ_{ijk} is a conditional probability, each column in Equation (1) must add up to one, so the number of free parameters in each column of the CPT is $(j-1)$.

Two remaining parameters associated with node X_i that play a particularly important role with regards to automated model adaption are the experience parameter ζ_i , and the fading parameter, η_i . Both parameters take the form of $k \times 1$ vectors, $\zeta_i = \{\zeta_{i1}, \zeta_{i2}, \dots, \zeta_{ik}\}'$ and $\eta_i = \{\eta_{i1}, \eta_{i2}, \dots, \eta_{ik}\}'$. This means that there is one experience parameter and one fading parameter for each parent node configuration of node X_i . Any node within the model that requires automated updating of its CPT parameters must have at least the experience parameter initiated for each of its parent node configurations. Further information related to these important parameters is left to a later detailed discussion of the model's parameter adaption automation. The complete parameter set for a node X_i is therefore $\Omega_i = [\Theta_i, \zeta_i, \eta_i]$.

Once \mathbf{G} has been defined, and Ω has had its parameter values assigned, the Bayesian network model can be used for inference. When no state information regarding the nodes has been entered into the network, the Bayesian network will infer the *a priori* marginal distributions, $\text{Pr}(X_i)$, of each node X_i in the model. In situations where some domain states are observed, these observed states can be instantiated into the model \mathbf{B} as evidence e . The Bayesian network can then be used to infer the *a posteriori* marginal distributions, $\text{Pr}(X_i|e)$, of the remaining unobserved nodes X_i , conditional on the evidence e .²

In order to infer the *a priori* and *a posteriori* marginal distributions, it is necessary for the Bayesian network to sum over each unobserved node's CPT state probabilities. Although this is conceptually straightforward, in a reasonably large network it is computationally demanding

²To distinguish between the marginal probabilities inferred from the network, and the conditional probabilities stored in the CPTs, marginal and conditional probabilities are denoted as $\text{Pr}(\dots)$, and $\text{p}(\dots)$, respectively.

on both memory storage and time. However, for discrete Bayesian networks, fast, efficient and exact inference algorithms are now available. The Bayesian network model described in this paper uses one of the fastest and most efficient algorithms, originally developed by Lauritzen and Spiegelhalter (1988).

Although the formal representation of a discrete Bayesian network given above is sufficient for the model developed in this paper, it should not be presumed to be definitive, as many variations in implementation exist. For an extensive discussion of these alternative implementations, refer to Koller and Friedman (2009).

Methodology

The methodology used to develop the tool's central Bayesian network model draws on Renooij (2001), Korb and Nicholson (2004) and Woodberry *et al* (2004). These methodologies emphasize the use of domain experts for probability elicitation and model construction, while incorporating the iterative feedback loops and re-modelling cycles appropriate for systems level operational risk modelling.^{3,4}

The Bayesian network model's development proceeded through four distinct stages. The first three stages being based on Korb and Nicholson (2004). The first stage, which was covered in Sanford and Moosa (2012), involved the model's structural development and evaluation. The second stage of the development involved the use of local domain experts to elicit values for the CPT parameters, while the third stage proceeded with model validation using parameter sensitivity analysis. The fourth and final stage was added in recognition of the need for ongoing refinement throughout the working life of the model, which involves the regular automated updating of the CPT parameters as new operational events occurred. It is the second, third and fourth developmental stages that are detailed in this paper.

The model

A detailed description of the model and the rationale behind its structural form is set out in Sanford and Moosa (2012). The model design was developed in collaboration with staff from the Operational Risk and SFO units, located within the wholesale banking division of one of Australia's largest banks. The director of quality

³Although the Bayesian network model development tool does come with automated network construction and parameter estimation functionality, absence of historical domain data makes these automated facilities unavailable.

⁴The Bayesian network development tool used for the research was HUGIN ResearcherTM v7.0, published by HUGIN EXPERT A/S website: www.hugin.com

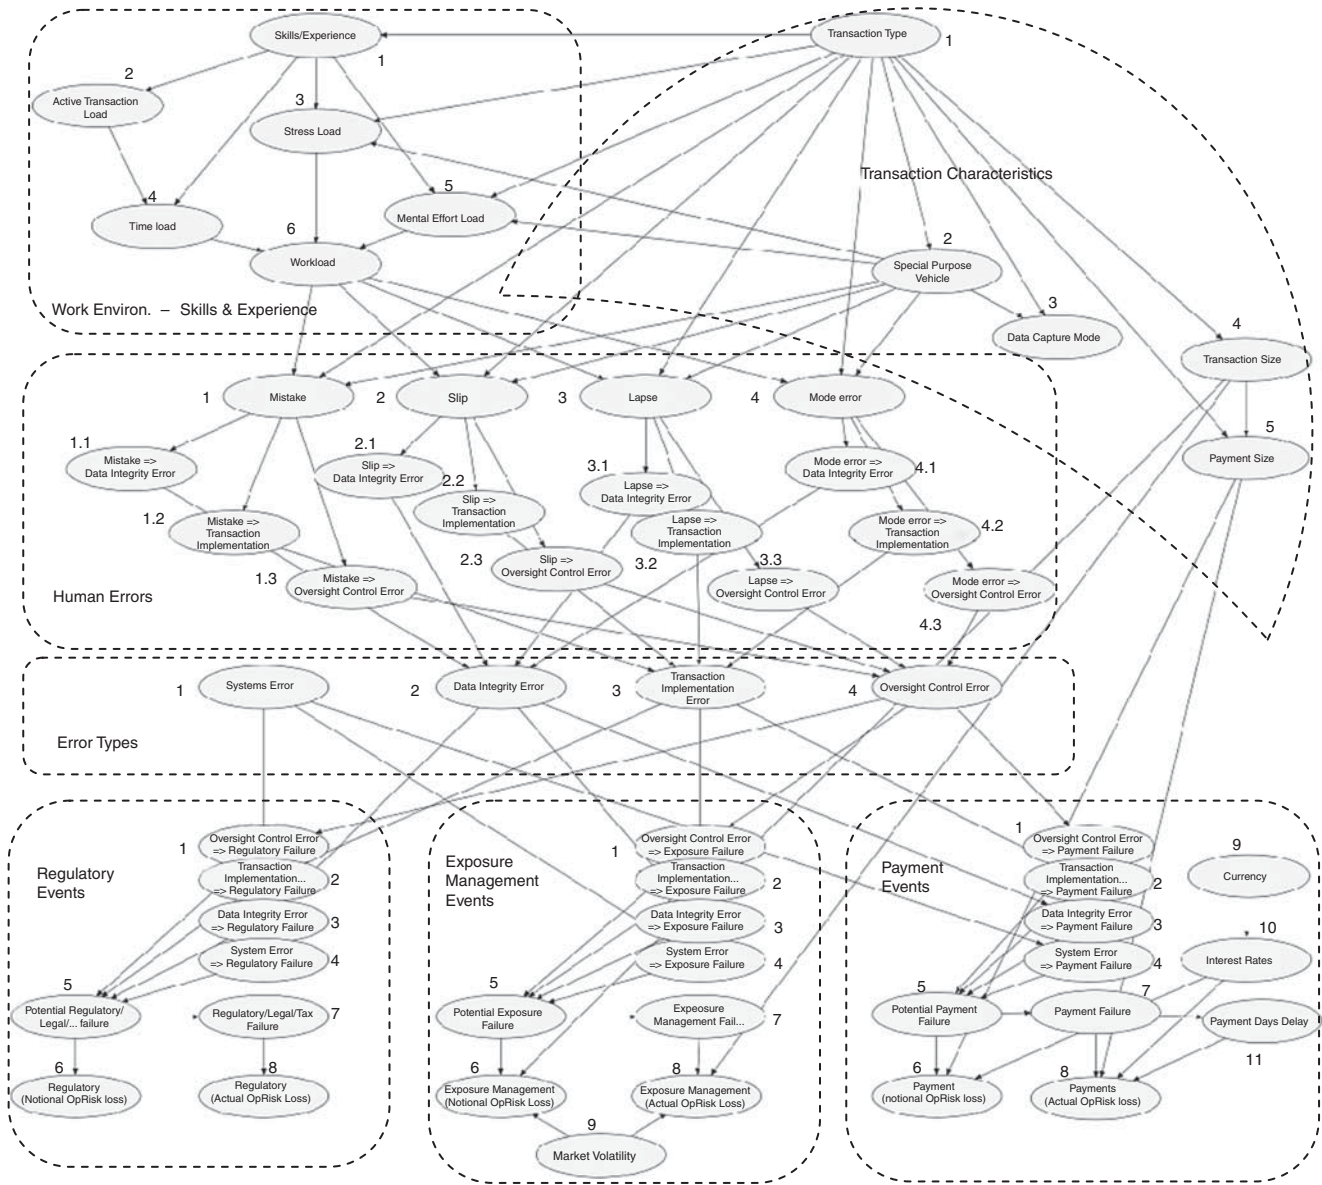


Figure 1 SFO operational risk network model.

assurance for SFO was the unit’s point of contact for the tool’s central model development. This person is referred to as the ‘risk manager’ throughout the paper. The risk manager represented the targeted end user for which the Bayesian network model was conceived and designed. They had extensive knowledge and experience of SFO’s domain, its operational processes, its risk exposures, and its staffing.

A graphical representation of the central model design and a brief description of its nodes are given in Figure 1 and Table 1. The design is unique in that it incorporates two existing analytical frameworks, one related to human error modes, and the other to socio-technical task analysis. The human error framework incorporates the generic human error categories developed previously by Reason (1990), while the second framework uses human factor loadings associated

with the subjective workload assessment technique (SWAT) developed previously by Reid and Nygren (1988).

The model is partly taxonomic (see Pearl, 1988, pp 333–337) and partly causal. Its taxonomic features classify operational risk events and their various causes via the two taxonomic frameworks of Reid and Nygren (1988) and Reason (1990). The network is not strictly taxonomic in that some classifications are not mutually exclusive for any given transaction. For example, it is possible for the risk manager to identify more than one human error type as a causal antecedent in the recording of any one operational loss event.

The model is partitioned into seven broad categories: (i) skills, experience and working environment; (ii) transaction characteristics; (iii) human errors; (iv) error types;

Table 1 Individual nodes [number of states] {number of parameters}

Category	Node	States	Parameters
<i>Work environment—skills & experience</i>	1 Skills/Experience	[3]	{26}
	2 Active Transaction Load	[6]	{15}
	3 Stress Load	[3]	{156}
	4 Time Load	[3]	{36}
	5 Mental Effort Load	[3]	{156}
	6 Work Load	[3]	{0}
<i>Transaction characteristics</i>	1 Transaction Type	[13]	{12}
	2 Special Purpose Vehicle	[2]	{13}
	3 Data Capture Mode	[2]	{26}
	4 Transaction Size	[7]	{78}
	5 Payment Size	[7]	{546}
<i>Exposure management events</i>	1 Oversight Control Error => Exposure Failure	[2]	{1}
	2 Trans.Impl.Error => Exposure Failure	[2]	{1}
	3 Data Integrity Error => Exposure Failure	[2]	{1}
	4 System Error => Exposure Failure	[2]	{1}
	5 Exposure Failure Near Miss	[2]	{0}
	6 Exposure Management (Notional Op Risk Loss)	[11]	{0}
	7 Exposure Management Failure	[2]	{1}
	8 Exposure Management (Actual Op Risk Loss)	[11]	{0}
	9 Market Volatility	[8]	{7}
<i>Regulatory events</i>	1 Oversight Control Error => Regulatory Event	[2]	{1}
	2 Trans.Impl.Error => Regulatory Event	[2]	{1}
	3 Data Integrity Error => Regulatory Event	[2]	{1}
	4 System Error => Regulatory Event	[2]	{1}
	5 Regulatory Near Miss	[2]	{0}
	6 Regulatory (Notional Op Risk Loss)	[4]	{2}
	7 Regulatory/Legal/Tax Failure	[2]	{1}
	8 Regulatory (Actual Op Risk Loss)	[4]	{2}
<i>Human errors</i>	1 Mistake	[2]	{78}
	1.1 Mistake => Data Integrity Error	[2]	{2}
	1.2 Mistake => Transact.Implement.Error	[2]	{2}
	1.3 Mistake => Oversight Control Error	[2]	{2}
	2 Slip	[2]	{78}
	2.1 Slip => Data Integrity Error	[2]	{2}
	2.2 Slip => Transact.Implement.Error	[2]	{2}
	2.3 Slip => Oversight Control Error	[2]	{2}
	3 Lapse	[2]	{78}
	3.1 Lapse => Data Integrity Error	[2]	{2}
	3.2 Lapse => Transact.Implement.Error	[2]	{2}
3.3 Lapse => Oversight Control Error	[2]	{2}	
4 Mode Error	[2]	{78}	
4.1 Mode Error => Data Integrity Error	[2]	{2}	
4.2 Mode Error => Transact.Implement.Error	[2]	{2}	
4.3 Mode Error => Oversight Control Error	[2]	{2}	
<i>Error types</i>	1 System Error	[2]	{1}
	2 Data Integrity Error	[2]	{0}
	3 Transact.Implement.Error	[2]	{0}
	4 Oversight Control Error	[2]	{0}
<i>Payment events</i>	1 Oversight Control Error => Payment Failure	[2]	{1}
	2 Transact.Implement.Error => Payment Failure	[2]	{1}
	3 Data Integrity Error => Payment Failure	[2]	{1}
	4 System Error => Payment Failure	[2]	{1}
	5 Payment Near Miss	[2]	{0}
	6 Payments (Notional Op Risk Loss)	[21]	{0}
	7 Payment Failure	[2]	{1}
	8 Payments (Actual Op Risk Loss)	[21]	{0}
	9 Currency	[9]	{8}
	10 Interest Rates	[9]	{72}
11 Payment Days Delay	[6]	{4}	

(v) payment failure events; (vi) exposure management events; and (vii) regulatory/legal/tax events. The nodes within each broad category represent a domain random variable of interest to the risk manager. They are therefore considered to be risk factors that influence the occurrence of operational loss events. Arrows are interpreted as representing causal relations between node states (see Table 1). The model is designed to generate probabilities of identified operational loss events within SFO. These include payment failures, exposure management failures and regulatory/legal/tax failures. The model also provides probability distributions over the severity of actual operational losses, as well as notional loss amounts associated with near miss events.⁵

In the Bayesian network model, continuous valued nodes were approximated by static discretization. For the initial settings, the number of partitions, or the degree

of static discretization, was influenced by the desire to minimize model complexity and reduce the elicitation and maintenance demands on SFO staff. Having fewer rather than more node state partitions means of course that the discrete approximations of the continuous variables within the domain are less accurate. Being static partitions also means that they will have to be reassessed from time to time during the working life of the model to ensure that excessive information loss is not occurring. Neil *et al* (2008) note that static discretization is unacceptable for critical systems, where reasonable accuracy is expected. In defence of the model's existing discretization scheme, we invoke Simon's concept of 'bounded rationality' (Simon, 1996), acknowledging that at this stage in the model's development the current static discretization is a compromise of 'satisficing' over 'optimizing'.

Following completion of the network design, specification of all nodes and their state spaces, and validation of the design, as detailed in Sanford and Moosa (2012), the next stage of development required the elicitation of the CPT probabilities for each of the nodes.

⁵An extensive, detailed description document—containing all node definitions, states, probabilities, as well as experience and fading table settings—is available on request from the corresponding author.

Table 2 Sample of elicitation capture screen for Lapse error probabilities conditional on transaction type = LOAN, and the presence or absence of an SPV, and the various workload levels

1. Given your experience of SFO, if at each **different level of work load, and with either an SPV involved or not involved**, how many of 1000 transactions would you expect to result in a **lapse error, material enough to be recorded as causing an actual OpRisk Event or potential OpRisk Event?** Give your ‘Very Surprised if count is less than this’ (5th percentile), Expected Count (50th percentile), ‘Very surprised if count is greater than this’ (95th percentile). You can revise your understanding of ‘Lapse’ errors and Workload definitions by referring to the accompanying definitions.

Transaction type	SPV?	Workload			
LOAN	TRUE	3- < 5	0	3	6
		5- < 7	1	4	7
		7-9	2	5	8
	FALSE	3- < 5	0	2	5
		5- < 7	1	3	6
		7-9	2	4	7
			Very surprised if count is less than this	Expected count	Very surprised if count is greater than this
<i>Elicited probabilities</i>					
LOAN	TRUE	3- < 5	0.00000	0.00300	0.00600
		5- < 7	0.00100	0.00400	0.00700
		7-9	0.00200	0.00500	0.00800
	FALSE	3- < 5	0.00000	0.00200	0.00500
		5- < 7	0.00100	0.00300	0.00600
		7-9	0.00200	0.00400	0.00700

Further elicitation capture cells (not shown) were prepared for all the remaining transaction types and human errors ‘Mistake’, ‘Slip’ and ‘Mode error’.

Parameter elicitation

The elicitation exercise required the determination of more than 1800 CPT parameter values. In carrying out this task, the following steps as described in Renooij (2001) were followed. The elicitation process involved (i) selection and motivation of the domain experts, (ii) training of the experts, (iii) structuring of the elicitation questions, (iv) eliciting and documenting of the domain experts judgements, and finally (v) verification of the elicited judgements.

Selection of the elicitation team was carried out by the risk manager. Criteria for team selection included the candidate’s technical knowledge of SFO operations, their familiarity with SFO operational risk, their operational experience and maturity, and their overall interest and motivation to be involved in the project. On the completion of the selection process the elicitation team comprised three personnel, one operational risk analyst from the bank’s Operational Risk department, an SFO manager, and an SFO analyst. At the selection stage, a decision was made not to include the risk manager as part of the elicitation team, despite their detailed knowledge of the domain. This decision freed the risk manager to provide later a secondary source of subjective expertise for cross-validation of the elicitation team’s judgements.

Because the elicitation team members were full-time operational staff, with existing work commitments and responsibilities, it was unavoidable that most of their elicitation tasks would have to be completed independently and remotely of the knowledge engineer and risk manager.

In response to this situation, the knowledge engineer developed a comprehensive data capture spread-sheet for use by all elicitation team members, which would provide them with additional support during the process.

Within the data capture spread-sheet, each node requiring elicitation was represented by a separate worksheet. Contained within each worksheet were the specific elicitation questions related to that node, and its response capture cells. Tables 2 and 3 provide examples of elicitation questions and data entry cells. These examples show how elicitation questions were typically framed, with language and concepts that were both clear and familiar to SFO staff, and free of any unfamiliar statistical terms. All elicitation questions required the expert to provide an estimate of either a count or a proportion as a response. As is shown in Table 2, the resultant node CPT parameter probability was determined from the expert’s expected count of lapse errors, over the 1000 transactions described in the elicitation question, including upper and lower bound probabilities derived from the expert’s beliefs as to what counts were unlikely to be observed either above or below. All worksheets displayed example responses and any relevant definitions (see the appendix for examples of reference documentation). Clear definitions of any conceptual constructs, such as the human error categories ‘mistake’, ‘slip’, ‘lapse’, or ‘mode error’, were detailed within the relevant worksheets for quick reference.

During the initial elicitation meeting, each elicitation team member was allocated a data capture spread-sheet that contained the nodes for which they were to be responsible.

Table 3 Sample of elicitation capture screen with recorded entries for Mental Effort Loads conditional on transaction type = LOAN, the presences or otherwise of a special purpose vehicle and the staff experience levels (Analyst/Senior Analyst/Associate)

1. Given your experience of working as an Analyst/Senior/Associate in SFO, when working on a typical LOAN transaction, with or without a special purpose vehicle, what would you describe as the most common, or typical, distribution of mental effort loads associated with this task.

Each row must sum to 1

<i>Transaction type</i>	<i>Using a SPV?</i>	<i>Skill/Experience</i>	<i>Very little conscious mental effort or concentration required. Activity is almost automatic, requiring little or no attention</i>	<i>Moderate conscious mental effort or concentration required. Complexity of activity is moderately high due to uncertainty, unpredictability, or unfamiliarity</i>	<i>Extensive mental effort and concentration necessary. Very complex activity requiring total attention</i>	<i>Uncertainty level (Very high, high, medium, low)</i>
Example: LOAN	Yes	Analyst	0.3	0.55	0.15	High
		Senior analyst	0.45	0.45	0.1	Medium
		Associate	0.6	0.4	0	High
LOAN	Yes	Analyst	0.0	0.5	0.5	High
		Senior analyst	0.0	0.5	0.5	High
		Associate	0.2	0.8	0.0	High
LOAN	No	Analyst	0.3	0.5	0.2	High
		Senior analyst	0.3	0.5	0.2	High
		Associate	0.2	0.8	0.0	High

Further elicitation capture cells (not shown) were prepared for all the remaining transaction types.

Such node allocations had been previously determined by the knowledge engineer and risk manager, based on their assessment of the specific knowledge and experience of each team member relative to the SFO domain. Following these allocations, the knowledge engineer, with assistance from the risk manager, instructed the team on the elicitation process, its iterative nature, and its relationship to the overall project.

All nodes were explained to the team during the training session. Those nodes that were less familiar to SFO staff, such as error type and human error classifications, were given particularly detailed explanation. Common elicitation biases, such as availability, anchoring, representativeness, and overconfidence (see Renooij, 2001), were also discussed, with descriptions of each included within the data capture spread-sheet for future reference as required.

Emphasized during the training session was the importance of capturing the uncertainty level associated with each expert's response. To capture uncertainty, two distinct methods were used. The first required the expert to give a response that represented the value that they would expect to observe, and two upper and lower bounded responses that they would be 'very surprised' to see exceeded. The lower, expected and upper values would be re-interpreted later by the knowledge engineer as the 5th percentile (lower bound), the 50th percentile (expected) and the 95th percentile (upper bound), over the distribution of the CPT parameter value concerned. This approach can be seen in the elicitation template of Table 2.

With the second, alternative uncertainty capture method each expert was asked to nominate their subjective level of uncertainty based on four pre-specified levels: 'very high', 'high', 'medium', and 'low' uncertainty, for each node state probability elicited, for a given parent node configuration. This method is shown in the elicitation template of Table 3. Each nominated uncertainty level would then be reinterpreted as a setting for the experience parameter ζ_i value. Very high, high, medium, and low uncertainties were allocated experience value settings of 50, 100, 200, and 400. The lower the experience parameter value, the higher the uncertainty considered. Following the elicitations, all uncertainty levels nominated by the elicitation team were found to be either 'very high', 'high', or 'medium'.

Although experience values are important for uncertainty capture, the relationship between them cannot be made clearer without first providing a more detail explanation of the Bayesian network inferential algorithm and the assumptions underlying the adaption method used by the model. While a discussion is left to the later section 'Organizational Learning through Model Adaption', in the interim an appreciation of the relationship can be had by consideration of the following. If it is assumed that node X_i 's CPT probability parameters for a given parent configuration k , that is, θ_{ijk} , are distributed under a j -dimensional multivariate Dirichlet, then by using Equation (10) below,

the variances of each of the j components of θ_{ik} can be calculated using knowledge of the k th experience value in ζ_i . According to Equation (10), the variance of parameter θ_{ijk} , *ceteris paribus*, will be smaller, that is, less uncertain (larger, ie more uncertain), when the value of the k th component of the experience parameter ζ_i becomes larger (smaller). Therefore, the values assigned to the experience parameters, as in those assigned for each uncertainty level, have an impact on the variance of the probability parameter, if it is assumed to be distributed under a Dirichlet distribution.

The correspondence between the second method of uncertainty capture, and the technical requirements for incorporating second-order uncertainty into the CPT parameters meant that the second capture method was the preferred approach. Those nodes that had initially used the first uncertainty capture approach were later adjusted to incorporate the uncertainty levels associated with the second capture method. To ensure continuity between the capture methods, the adjustment was performed by the relevant elicitation team members with support from the knowledge engineer.

On the completion of the elicitation training, team members dispersed to carry out their elicitation tasks, primarily as a background activity to their regular work responsibilities. The operational risk analyst and SFO staff completed their responses in three and five weeks, respectively. During the elicitation period, team members maintained communications with the knowledge engineer and risk manager, to ensure that any issues could be addressed as they arose.

On the completion of the data capture phase of the elicitation, validation of responses for each node was carried out by the risk manager and knowledge engineer using a formal walk-through of the captured responses. For any responses judged to be higher or lower than expectations, an additional elicitation meeting was arranged to discuss the deviation with the relevant team member and risk manager. Iterations between the expert responses and response validations continued over several weeks until all elicitation team members, the knowledge engineer, and risk manager were satisfied with the elicitation responses provided.

Once all the counts, proportions and uncertainty levels had been documented and cross-validated, the knowledge engineer converted the expert's final responses into probabilities and experience values for entry into the CPTs. Given the manner in which the elicitation questions were framed, and the response format of either counts or proportions, conversion to probabilities was relatively straightforward. For example, as with the Lapse Error node, the elicitation question asked the expert to specify their judgement of the count of Lapse Error events under three different workload levels, based on 1000 loan transactions. Probabilities were generated by taking the

response count and dividing by the number of 1000 loan transactions.⁶ Implicit in this approach was that the responses provided by the experts represented their beliefs with regards to long-term averages within the domain.

Once all CPT inputs had been completed, the Bayesian network was then able to infer the marginal probability distributions for each node. The knowledge engineer and risk manager reviewed the marginal distributions generated for each of the main output nodes, payment failure, exposure management failure and regulatory/legal/tax failure. These marginal distributions were compared against the risk manager's own subjective judgements, and available SFO documentation. As is common when parameterizing a Bayesian network model, the initial marginal distribution did not immediately fit with the expert's subjective marginal. This was true for the payment failure node marginal, suggesting the need for additional CPT parameter tuning. This tuning process is described in the next section.

Parameter tuning

In the initial parameterization of the Bayesian network model, the domain experts were asked to provide judgement on events that were highly conditional on other existing domain states. That is how the CPT parameters were assessed, as the CPT parameter values are conditional on their parent node values. However, domain experts will also have judgements, which are based on what they consider to be the marginal probabilities of events. That is those events where background conditionals are still considered, but rather 'averaged' over a typical period of time, or transaction volume. These marginal judgements are also of value when parameterizing the Bayesian network model as they provide a standard on which to compare the outputs of the model against. If there is some disagreement between the experts, and the model, then adjustments to CPT parameters can be made. It is these adjustments that constitute the parameter tuning process. A detailed, mathematical discussion of tuning, which is based on Russell *et al* (1997), Castillo *et al* (1996) and Jensen (1999), can be found in Jensen and Nielsen (2007, pp 218–222).

It is not unusual to elicit marginal distributions from a domain expert in order to parameterize a model. It is this approach to elicitation that forms the background to the technique of probabilistic inversion as described by Bedford and Cooke (2001, Ch.16). In this situation, the marginal probability distributions are elicited from experts

⁶The background transaction volumes (ie 1000 transactions) used in the elicitation questions exceeded the current transaction volume history of SFO. The motivation for designing the questions around such large hypothetical loan transaction volumes was to encourage domain experts to base their responses on their beliefs on the long-term averages within the domain.

in order to parameterize a probability model. The elicited marginal distributions then become input to an inversion algorithm, which determines the corresponding joint probability distribution over the model's parameters. Using this technique, the model is then able to recreate the marginals as specified by the experts, as well as other probabilistic outputs consistent with the model and the expert's judgements. Although parameter tuning is not strictly probabilistic inversion, it is similar in that adjustments to model parameters are identified in order to obtain an expert's marginal judgement. The approach to model parameterization using probabilistic inversion also demonstrates the legitimacy of using marginals for parameterization. In the following discussion, we relate the parameter tuning of the model's payment failure node.

After completing the CPT parameter entries the model's payment failure node registered a marginal probability of a payment failure of 0.026 for a single transaction. In terms of an expected or long-term average, this constituted on average eight payments failures per 300 processed transactions.⁷ This exceeded the expectations of the risk manager and other SFO domain experts who considered five payment failures per 300 processed transactions to be a more reasonable assessment. To replicate such an expectation with the Bayesian network model, an *a priori* marginal probability of a payment failure would need to be closer to a probability of 0.017 for a single transaction.⁸ On the basis of the domain expert's judgement, it was decided to re-align the Payment Failure node marginal distribution by adjusting CPT parameter values within the network.

In order to remain as faithful as possible to the original elicited CPT parameter values, it was decided to adjust only parameters that could be done so without exceeding either their upper or lower bounds, as specified by the domain experts, or a one standard deviation measure from their initial setting. If it was not possible to identify a parameter satisfying these constraints, then a compromise was to be made between the marginal distribution target, and the parameter adjustments applied. A parameter's standard deviation was calculated using the Dirichlet variance expression described in Equation (10).⁹

The Bayesian network development tool, used in the creation of the model, included a parameter tuning facility

⁷At the time of the models development, 300 transactions was the risk manager's estimate of the number of transactions processed by SFO during its operational period.

⁸Care should be taken when interpreting perceived deviations of the model's marginal probabilities from any observed domain events, as any apparent deviations may be due purely to random variation rather than from any misspecification of the node. This important point was highlighted by an anonymous referee in his/her comments on an earlier version of the paper.

⁹The rationale for the Dirichlet distribution is discussed in the parameter adaption section.

Table 4 Candidate parameter changes for payment failure parameter tuning

Node CPT parameter	Current parameter value	Suggested parameter value	Absolute change	Log odds ratio
$p(\text{System Error} = \text{true})$	0.0500	0.0194	0.0306	0.9780
$p(\text{Systems Error} = > \text{Payment Failure} = \text{true} \mid \text{System Error} = \text{true})$	0.0300	0.1165	0.1835	1.1792
	Variance	Standard deviation	Upper bound	Lower bound
$p(\text{System Error} = \text{true})$	0.0009	0.0305	0.0805	0.0195
$p(\text{Systems Error} = > \text{Payment Failure} = \text{true} \mid \text{System Error} = \text{true})$	0.0021	0.0456	0.3456	0.2544

that allowed a model developer to identify a set of CPT parameters and their required adjustments, to achieve a targeted marginal distribution. This facility automatically tests node CPT parameter adjustments to determine if the target node constraints can be achieved, and then reports those CPT parameters that were successful.¹⁰ After entering the target constraint of $\Pr(\text{Payment Failure} = \text{true}) = 0.017$ into the tuning facility, two candidate CPT parameters and their required adjustments were reported by the facility. The two successful CPT parameters reported are listed in Table 4.

Of the two candidate CPT parameters identified, neither could be adjusted without exceeding the one standard deviation constraint criteria imposed. Adjustment to the $p(\text{Systems Error} = \text{true})$ parameter however exceeded the constraint only slightly. The $p(\text{Systems Error} = \text{true})$ parameter was therefore adjusted from 0.05 to 0.02, an adjustment within one standard deviation, achieving a final marginal of $\Pr(\text{Payment Failure} = \text{true}) = 0.0172$. At the completion of the probability elicitation process and parameter tuning, parameter sensitivity analysis was performed as described in the following section.

Validation and sensitivity analysis

Parameter sensitivity analysis investigates the sensitivity of both the *a priori* $\Pr(X)$ and *a posteriori* $\Pr(X|e)$ distributions to changes in the underlying CPT parameters θ located throughout a network. Within the network, any CPT parameter that is d-connected to any other node can influence the distribution of that node.¹¹ Parameter sensitivity analysis can be performed using one-, two- or *n*-way analysis, whereby up to one, two or *n* CPT parameters, anywhere in the network, are varied simultaneously. Their total effect on the target node's marginal distribution is then measured to

¹⁰Owing to the proprietary nature of the parameter tuning facility, underlying implementation details are not available.

¹¹The term 'd-connected' refers to the relation between nodes within a DAG. If two or more nodes are d-connected, then changes in one nodes state will affect the marginal probabilities of states for the other nodes. In this situation we are referring to changes in parameter values. In this context d-connected means that if we consider the parameters to the node, to be actual auxiliary parent nodes, then those auxiliary nodes must be d-connected.

determine how sensitive the node is to variation in those parameters. One-way analysis approach, the most common and easiest to interpret, was used in our model development.

The parameter sensitivity analysis takes advantage of a convenient mathematical property associated with discrete Bayesian networks, which states that for any joint probability distribution, $\Pr(X_o|e)$, over output nodes X_o and evidence e , the distribution is a linear function of a single CPT parameter, θ_{ijk} , located anywhere in the network, provided that the parameter is d-connected with any of the nodes over which the distribution is defined. This linear function can be written as

$$\Pr_{ijk}(X_o \wedge e) = f(\theta_{ijk}) = c_1 \times \theta_{ijk} + c_2 \quad (2)$$

The subscript *ijk* indicates that the joint probability distribution $\Pr_{ijk}(X_o|e)$ is a function of the single CPT parameter, θ_{ijk} . This parameter is identified as the conditional probability parameter, for node *i*, in state *j*, whose parents are in their *k*th configuration. If the parameter θ_{ijk} is varied, the joint distribution $\Pr_{ijk}(X_o|e)$ will change at a rate of c_1 . From this basic property, the posterior distribution, $\Pr_{ijk}(X_o|e)$, of the output node X_o conditional on evidence e is defined as the ratio of two linear functions

$$\Pr_{ijk}(X_o|e) = \frac{\Pr_{ijk}(X_o \wedge e)}{\Pr_{ijk}(e)} = f(\theta_{ijk}) = \frac{c_1 \times \theta_{ijk} + c_2}{c_3 \times \theta_{ijk} + c_4} \quad (3)$$

Equation (3) is commonly referred to as the sensitivity function. This function comprises four coefficients, c_1 , c_2 , c_3 , and c_4 , which can be readily estimated from four separate simultaneous equations, two based on one setting for θ_{ijk} , and the remaining two based on an alternative setting for θ_{ijk} . Jensen and Nielsen (2007, p 188) discusses in more detail the estimation of c_1 , c_2 , c_3 , and c_4 .¹²

The results described by Equations (2) and (3) will only hold if two further conditions are also satisfied. The first is that the other parameters in the *j*-dimensional vector

¹²The proofs for Equations (2) and (3) can be found in Castillo *et al* (1995), with a detailed discussion of their application in Castillo *et al* (1997), van der Gaag *et al* (2007) and Kjerulff and Madsen (2008). A systematic approach to sensitivity analysis is also described in Bednarski *et al* (2004).

$(\theta_1, \theta_2, \dots, \theta_{ik})$ must be co-varied proportionally according to the expression

$$\left\{ \theta_{r \neq j} = \theta_{r \neq j}^0 \times \frac{1 - \theta_j}{1 - \theta_j^0} \right\}_{ik} \quad (4)$$

The parameters $\theta_{r \neq j}^0$ and θ_j^0 in Equation (4) represent the original values of the parameters, while $\theta_{r \neq j}$ and θ_j represent the adjusted parameter settings. The second condition requires that the parameter θ_{ijk} be non-extreme, located in the open range (0,1). The proportional co-variation constraint equation (4) assumes that the parameters of node X_i are varied independently of the parameters of all other nodes, $X_{j \neq i}$, (ie global independence), and independently of all other parameters in X_i 's CPT, which are conditional on all other non- k th parent configurations (ie local independence).

For the distribution $\Pr_{ijk}(X_o|e)$, the parameter sensitivity measure, S_{ijk}^o , is determined by the absolute value of the first derivative of Equation (3) at θ_{ijk}

$$S_{ijk}^o = \left| \frac{d \Pr_{ijk}(X_o|e)}{d \theta_{ijk}} \right| = \left| \frac{c_1 \times c_4 - c_2 \times c_3}{(c_3 \times \theta_{ijk} + c_4)^2} \right| \quad (5)$$

The more sensitive $\Pr_{ijk}(X_o|e)$ is to changes in θ_{ijk} , the larger the value of S_{ijk}^o . The more sensitive $\Pr_{ijk}(X_o|e)$ is to changes in θ_{ijk} , the more elicitation and validation resources that should be expended to ensure that the θ_{ijk} 's current setting is as close to its true value as can be accurately or subjectively determined. Because of the higher sensitivity, any small deviation in θ_{ijk} from its 'true' value will result in a comparatively larger deviation for the marginal, $\Pr_{ijk}(X_o|e)$.

As can be seen from Equation (3), when the network contains evidence e , the function is a ratio of two linear functions of θ_{ijk} . This means that the resultant function takes the form of a rectangular hyperbola (van der Gaag *et al*, 2007). The function can therefore have regions of very high and very long parameter sensitivity. Care must be taken to ensure that a correct interpretation of parameter sensitivity is carried out. High sensitivity regions may exist for a node, but these regions are only problematic if they occur in an area that is most likely to contain the true value of θ_{ijk} .

Separate parameter sensitivity analyses were performed for each of the model's main output nodes, payment failure, exposure management failure, and regulatory/tax/legal failure. The result of each analysis is detailed in Tables 5 and 6. For each analysis, the initial step was to take the sensitivity measure of the output node without evidence present in the network. This provided a 'first cut' measure to identify the most influential CPT parameters for the output node. With no evidence present, Equation (2) is the relevant sensitivity function and c_1 is the equivalent sensitivity measure S_{ijk}^o . The results for this initial step are shown in Table 5, which does not identify

Table 5 Parameter sensitivity (no evidence)

ith Node	$X_o = \text{Payment failure}$		$X_o = \text{Exposure management failure}$		$X_o = \text{Regulatory/legal/tax failure}$	
	max S_{ijk}^o	ith Node	max S_{ijk}^o	ith Node	max S_{ijk}^o	ith Node
System Error	0.29	System Error	0.10	System Error	0.05	System Error
Transaction Type	0.05	Transaction Type	0.03	Transaction Type	0.05	Transaction Type
System Error = > Pot. Payment Fail.	0.02	System Error = > Pot. Expose. Fail.	0.02	System Error = > Pot. Regulatory Fail.	0.02	System Error = > Pot. Regulatory Fail.
Data Integrity Error = > Pot. Payment Fail.	0.02	Data Integrity Error = > Pot. Expose. Fail.	0.02	Data Integrity Error = > Pot. Regulatory Fail.	0.02	Data Integrity Error = > Pot. Regulatory Fail.
Mistake	0.02	Mistake	0.02	Lapse	0.02	Lapse
Slip	0.02			Mode Error	0.02	Mode Error
Lapse	0.02			Slip	0.02	Slip
Mode Error	0.02			Mistake	0.02	Mistake

Parameter Sensitivity measure is calculated using Equation (5). The c_1, c_2, c_3, c_4 coefficients in Equation (5) are not shown but estimated by the Bayesian network tool when calculating the parameter sensitivity measure. The parameter sensitivity measure value is shown under **max S_{ijk}^o** which is the maximum sensitivity measure found in the i th node's CPT, for each X_o as the output node of interest. The maximum sensitivity measure **max S_{ijk}^o** value relates to the θ_{ijk} parameter in the i th node's CPT.

the specific parameter θ_{ijk} that has the maximum sensitivity measure, but instead the relevant i th node who's CPT contains the parameter with the maximum sensitivity measure. We have denoted the value of the sensitivity measure as $\max S_{ijk}^o$. For each output node, a listing from the highest to the lowest $\max S_{ijk}^o$ is provided.

The results displayed in Table 5 show that the System Error CPT parameters are the most influential on the Payment Failure node distribution. A sensitivity measure value of 0.29 means that for an absolute change in the value of parameter $\theta_1 = p(\text{System Error} = \text{True})$ of $|\Delta\theta_1| = 0.01$, there would be a resultant absolute change in the marginal probability of the Payment Failure node of $|\Delta\Pr(\text{Payment Failure} = \text{True})| = 0.0029$. This is verified with the model. By adjusting $p(\text{System Error} = \text{True}) = 0.02$ to 0.03, we observe that the marginal $\Pr(\text{Payment Failure} = \text{True}) = 0.0172$, changes to 0.0201, a delta of 0.0029. Table 5 shows that all target nodes are particularly sensitive to the System Error CPT parameters. Each target node's sensitivity set is generally consistent with the other target nodes. A comparatively low sensitivity is recorded for each transaction type CPT parameter, as well as with the CPT parameters of each target node's most immediate, antecedent nodes.

To measure parameter sensitivities of the output node's *a posteriori* distribution, $\Pr_{ijk}(X_o|e)$, evidence e was instantiated into the network. Owing to the large set of candidate nodes available for instantiation, evidence was restricted to those nodes considered the main causal antecedents to the output failure nodes. These included the transaction type, human error and error type nodes. A manual evidence generation exercise performed by the knowledge engineer resulted in 110 evidence sets e_n , $n = 1, \dots, 110$. Each set contained different combinations of transaction type, human error and error event type states.

For each output node X_o , and the evidence set e_n entered into the model, the maximum parameter sensitivity, $S_{ijk}^{o,n}$, was recorded. This resulted in 110 maximum parameter sensitivity measure recordings and the associated ijk parameter, for each output node.

To calculate an overall parameter sensitivity set for each output node X_o , each of the 110 sensitivity measure, $S_{ijk}^{o,n}$, were weighted by the evidence set e_n probability. For example, the two evidence sets, $e_8 = (\text{Transaction Type} = \text{'XIRS'}, \text{Mistake} = \text{TRUE})$ and $e_{62} = (\text{Transaction Type} = \text{'EXC-Derive'}, \text{System Error} = \text{TRUE})$, have according to the model, joint probabilities of $\Pr(e_8) = 0.00325$ and $\Pr(e_{62}) = 0.00001941$. Under evidence set e_{62} , the 'System Error' node's CPT parameters are found to be the most influential, with a maximum absolute sensitivity measure $S_{ijk}^{o,62} = 0.94$. Under evidence set e_8 , the 'Data Integrity Error=>Potential Payment Failure' node's CPT parameters are found to be the most influential, with a maximum absolute sensitivity measure $S_{ijk}^{o,8} = 0.46$. The 'Data Integrity Error=>Potential Payment Failure' node's sensitivity measure is less than the 'System Error'

node. Weighting each of these nodes by the joint probability of the evidence, $\Pr(e_8) \times S_p = 0.00325 \times 0.46 = 0.0015$ versus $\Pr(e_{62}) \times S_p = 0.00001941 \times 0.94 = 0.00001825$, demonstrates that it is better to direct resources to the node 'Data Integrity Error=> Potential Payment Failure', which is the most influential, given the probability of observing the evidence.

The results from the parameter sensitivity analysis involving evidence are shown in Table 6. It should be emphasized that the final listing of nodes, whose CPT parameters were most influential, is ultimately a function of the evidence used in the evidence sets. We are confident, however, that given the breadth of the evidence used, the identification of those nodes and their CPT parameters most influential in an instantiated network is valid. Given that many of the same nodes and CPT parameters were repeatedly identified, in order to find an overall score, the probability weighed sensitivity measures were added together for those nodes and CPT parameters with multiply recordings. Those nodes and CPT parameters with the highest score were then selected as the most influential, given all evidence sets. The ordering of influential nodes shown in Table 6 reflects this where, for example, we see that the CPT parameters for nodes 'System Error=> Potential Exposure Fail' and 'System Error=>Potential Regulatory Failure' are considered to be less influential, despite having a higher $\max S_{ijk}^o$.

During the validation phase, those CPT parameters identified as particularly influential were revisited in discussion with the domain experts responsible for their elicitations. Following the initial elicitation and parameter tuning exercises however, the validation exercise tended to re-confirm the expert's elicitation responses rather than challenge them. Figure 2 provides a screen shot of the completed model following elicitation and validation of the network.

Organizational learning through model adaption

After having elicited, calibrated and validated the central model's probabilities, the remaining important issue was the design of a formal process whereby the central model could be maintained, adapted and improved over time. Although the use of subjective expert judgement was unavoidable in the initial construction of the model, in moving to a post-implementation phase, the arrival of new event information and the need for ongoing maintenance, the adaption and improvement of the model could be more efficiently handled by employing automated adaption techniques.¹³

¹³The adaption of the network model's CPT parameters involves using the sequential adaptation algorithm to incorporate new event data as required. A database was also developed to support the adaption process, capturing operational events relevant to the model's adaption. A document detailing the database design is available from the corresponding author.

Table 6 Parameter sensitivity (with evidence)

ith Node	$X_o = \text{Payment failure}$			$X_o = \text{Exposure management failure}$			$X_o = \text{Regulatory/legal/tax failure}$		
	$\max S_{ijk}^o$	ith Node	$\max S_{ijk}^o$	$\max S_{ijk}^o$	ith Node	$\max S_{ijk}^o$	$\max S_{ijk}^o$	ith Node	$\max S_{ijk}^o$
Data Integrity Error = > Pot. Payment Fail.	0.86	Data Integrity Error = > Pot. Expose. Fail.	0.89	Data Integrity Error = > Pot. Regulatory Fail.	0.83				
Trans. Impl. Error = > Pot. Payment Fail	0.80	Trans. Impl. Error = > Pot. Expose. Fail	0.84	Trans. Impl. Error = > Pot. Regulatory Fail	0.83				
System Error = > Pot. Payment Fail	0.79	System Error = > Pot. Expose. Fail	0.99	System Error = > Pot. Regulatory Fail	0.98				
Oversight Control Error = > Pot. Payment Fail	0.79	Oversight Control Err. = > Pot.Expose.Fail	0.79	Oversight Control Error = > Pot. Regul. Fail	0.80				

Parameter Sensitivity measure is calculated using Equation (5). The c_1, c_2, c_3, c_4 coefficients in Equation (5) are not shown but estimated by the Bayesian network tool when calculating the parameter sensitivity measure. The parameter sensitivity measure value is shown under $\max S_{ijk}^o$ which is the maximum sensitivity measure found in the i th node's CPT, for each X_o as the output node of interest. The maximum sensitivity measure $\max S_{ijk}^o$ value relates to the θ_{ijk} parameter in the i th node's CPT.

Parameter adaption

Although a number of parameter adaption algorithms are now available for Bayesian networks, ‘incremental updating’, an algorithm developed originally by Spiegelhalter and Lauritzen (1990), is the one implemented in the tool’s central Bayesian network model. The conceptual and statistical basis of the incremental updating algorithm is Bayesian inference (see Gelman *et al*, 2004; Lancaster, 2004), which we will first discuss before looking at the incremental updating algorithm in more detail.

Bayesian inference can be characterized as the systematic application of Bayes Theorem to infer from observed data, new information on unobserved or unobservable variables, where such variables may include latent factors or parameters of a probability model. Bayes Theorem can be expressed for discrete probability models in the form,

$$\begin{aligned}
 \Pr(\theta|e) &= \frac{\Pr(e, \theta)}{\sum_{\theta} \Pr(e, \theta)} = \frac{\Pr(e|\theta) \Pr(\theta)}{\sum_{\theta} \Pr(e|\theta) \Pr(\theta)} \\
 &= \frac{\Pr(e|\theta) \Pr(\theta)}{\Pr(e)} = \frac{L(\theta|e) \Pr(\theta)}{\Pr(e)} \\
 &\propto \Pr(e|\theta)P(\theta) = L(\theta|e) \Pr(\theta) \tag{6}
 \end{aligned}$$

In the Bayesian theorem representation of Equation (6), θ is the parameter or parameters of the probability model, while e represents the observations generated from the data-generating process, $\Pr(e|\theta)$. The data generating process can be substituted with the likelihood function, $L(\theta;e)$. The denominator, referred to as the marginal likelihood $\Pr(e)$, represents the probability of observing the data e , after integrating over the model parameter(s) θ . As the marginal likelihood is not a function of θ , it acts as a normalizing constant that ensures the posterior distribution $\Pr(\theta|e)$ integrates to one.

The introduction of prior beliefs over a model’s parameters θ distinguishes Bayesian inference from classical statistical inference. The prior beliefs represent the information pertaining to θ , prior to observing any new information e . These prior beliefs are typically described using a probability distribution function, $\Pr(\theta)$. The prior can be sourced from an investigator’s subjective judgements, from a previously inferred posterior distribution, or from theoretical considerations. Introduction of a prior in the model adaption process is of great benefit, as it provides a facility that allows for the representation of second-order uncertainty over the CPT parameters.

In eliciting a model’s parameter values, domain experts often prefer to specify a range over which they believe the true value is located, rather than a point estimate. A range providing some facility to indicate the uncertainty associated with the expert’s beliefs. Narrow ranges implying greater

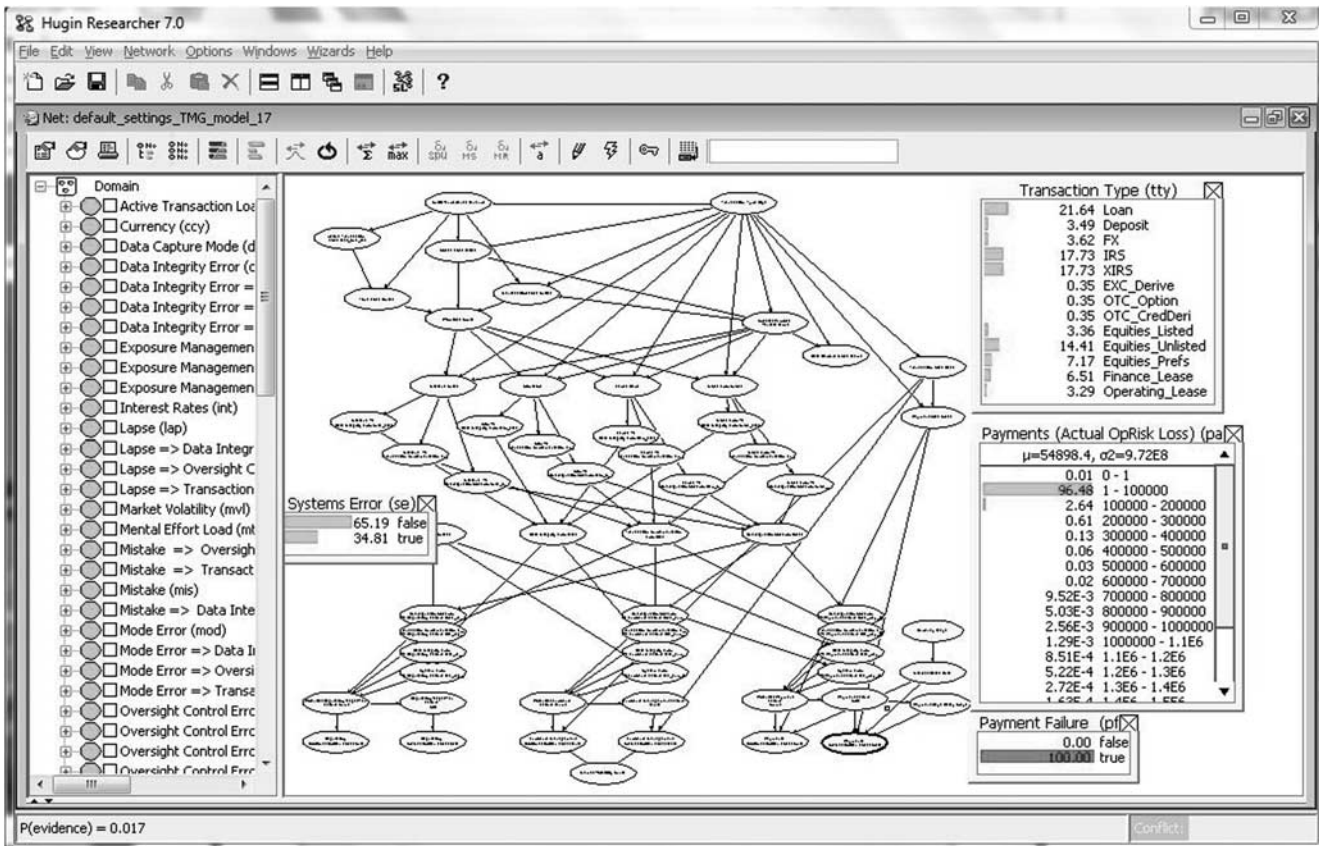


Figure 2 Bayesian network model following completion of probability, elicitations, parameter tuning and sensitivity analysis.

certainty whereas broad ranges implying less certainty. Without the facility to capture uncertainty, all point values of the CPT parameters appear to be specified with complete certainty. In the Bayesian network model developed, however, such a measure of uncertainty can be captured through the specification of a prior distribution over the CPT parameters in conjunction with the experience parameter ζ .

The experience parameter ζ_i for node X_i consists of the vector $\zeta_i = \{\zeta_{i1}, \zeta_{i2}, \dots, \zeta_{ik}\}$. Each element of the experience vector ζ_i is a positive real number. The experience value can be interpreted as the count of observations on which the current CPT parameters are based for the given parent configuration k . Despite this interpretation, however, it should be noted that in situations where new evidence is only partially observed, or considered highly unlikely, the value of ζ_i can actually decrease due to the greater parameter uncertainty introduced by the new observations.

To adapt or infer new CPT parameters when new data are observed using Bayesian inference it is necessary first to establish a prior over each of the CPT parameters and to define the data-generating process $\Pr(e|\theta)$. As this is a discrete Bayesian network model, the parameters contained within a node's $(j \times k)$ CPT matrix Θ_i can be

considered to represent the parameters of k distinct, j -dimensional, multinomial distributions,

$$p_{ik}(e|\theta) \propto \left\{ \prod_{m=1}^j \theta_m^{e_m} \right\}_{ik} \quad (7)$$

In the un-normalized representation of a j -dimensional multinomial represented by Equation (7), the observation e is a $j \times 1$ vector of counts, where $e_{ik} = \{e_1, e_2, \dots, e_j\}$. The vector e_{ik} contains the number of times that X_i has been in its j th state, whenever its parent nodes were in their k th configuration. Although not explicit, the multinomial distribution of e_{ik} in Equation (7) is also conditional on the total number of trials or observations, which is equal to the sum of all the counts in e_{ik} .

The prior over each of the CPT parameters takes the form of a Dirichlet distribution, which is the multivariate version of a Beta distribution. The Dirichlet distribution has two properties that make it especially efficient for use in Bayesian inference and parameter adaption within a discrete Bayesian network. The first property is that it has sufficient statistics, which makes the storage of information from past observations very efficient, as all information

contained within a data sample e can be summarized as a function of the observed data without any information loss. The sufficient statistic for a j -dimensional Dirichlet distribution consists of the counts of the outcomes for each of the j possible states observed. The Dirichlet prior over node X_i 's CPT parameters, having j possible states, and a k th parent configuration, is denoted as *Dirichlet* _{ik} ($\theta_1, \dots, \theta_j | \alpha_1, \alpha_2, \dots, \alpha_j$). The Dirichlet prior is parameterized by the hyperparameters $\alpha_{ik} = \{\alpha_1, \alpha_2, \dots, \alpha_j\}_{ik}$. These are referred to as hyperparameters to distinguish them from the model parameters, θ_j . A Dirichlet distribution takes the form

$$p_{ik}(\theta_1, \dots, \theta_j | \alpha_1, \dots, \alpha_j) \propto \left\{ \prod_{m=1}^j \theta_m^{\alpha_m - 1} \right\}_{ik} \quad (8)$$

In this un-normalized Dirichlet distribution represented by Equation (8), each $\theta_m \geq 0$ and the sum of all θ_m must be equal to one. The first and second moments of a Dirichlet random variable, which will be later shown to play an important role in incremental updating, are calculated as follows:

$$E_{ik}(\theta_j) = \left(\frac{\alpha_j}{\zeta_k} \right)_i \quad (9)$$

$$Var_{ik}(\theta_j) = \left(\frac{\alpha_j(\zeta_k - \alpha_j)}{\zeta_k^2(\zeta_k + 1)} \right)_i = \frac{E_{ik}(\theta_j)[1 - E_{ik}(\theta_j)]}{(\zeta_k + 1)} \quad (10)$$

where

$$(\zeta_k)_i = \left(\sum_{m=1}^j \alpha_j \right)_{ik}$$

The ik subscripts indicate that the moments are conditional on the k th parent node configuration for node X_i .

The second valuable property of the Dirichlet distribution is that it is a conjugate prior to the multinomial distribution. To be a conjugate prior means that the inferred posterior is of the same distributional form as the prior, that is Dirichlet. This allows the updated posterior Dirichlet to be used as the new prior, when subsequent parameter updates are required. This can be seen in Equation (11), where the hyperparameters of a Dirichlet prior distribution are updated to become the hyperparameters of the Dirichlet posterior via the likelihood function.

$$\left\{ \prod_{m=1}^j \theta_m^{\alpha_m + e_m - 1} \right\}_{ik} = \left\{ \prod_{m=1}^j \theta_m^{e_m} \times \prod_{m=1}^j \theta_m^{\alpha_m - 1} \right\}_{ik} \quad (11)$$

posterior
likelihood
prior

In Equation (11), the hyperparameters of the posterior Dirichlet are calculated by simply adding the observation counts to the existing prior Dirichlet parameters, where the

posterior hyperparameters are $\alpha + e = (\alpha_1 + e_1, \alpha_2 + e_2, \dots, \alpha_j + e_j)_{ik}$.

Incremental updating algorithm

In parameter adaption, because all node parameters are updated together, CPT parameter dependencies must be taken into account. The adaption process makes two assumptions regarding parameter dependencies. These are referred to as 'global independence' and 'local independence' assumptions. With global independence, the assumption is that each X_i 's CPT parameter can be adapted independently of all other node's CPT parameters. With local independence, the assumption is that the parameters within a single node's CPT can be updated independently for a given parent configuration k .

Parameter adaption is relatively straightforward when the evidence used contains observations of all nodes. It becomes more complex when this is not the case, as the absence of observed node states creates dependencies between the nodes and their CPT parameters, making parameter adaption impossible without reference to other parameters within the network (Jensen and Nielsen, 2007, p 207).

The following details of the incremental updating algorithm are discussed in the context of updating a single node X_i , having j possible states and k parent node configurations. The CPT parameters to be updated, whose current value settings are derived from previous observations on X_i , are denoted as $p_{ik}(X_i | Pa_k(X_i)) = (\theta_1, \theta_2, \dots, \theta_j)_{ik}$. The functional relationship between the current CPT parameter values $(\theta_1, \theta_2, \dots, \theta_j)_{ik}$, the hyperparameters $\alpha_{ik} = (\alpha_1, \alpha_2, \dots, \alpha_j)$ of the Dirichlet prior and the experience parameter ζ_k is

$$(\theta_1, \theta_2, \dots, \theta_j)_{ik} = \left(\frac{\alpha_1}{\zeta_k}, \frac{\alpha_2}{\zeta_k}, \dots, \frac{\alpha_j}{\zeta_k} \right)_{ik}$$

Each hyperparameter α_{ijk} does not need to be stored explicitly, but can be recovered from the current CPT parameter values and their experience value, $\alpha_{ik} = (\alpha_1, \alpha_2, \dots, \alpha_j)_{ik} = (\zeta_k \times \theta_1, \zeta_k \times \theta_2, \dots, \zeta_k \times \theta_j)_{ik}$. The posterior Dirichlet distribution over node X_i 's parameters is updated using the rule

$$p_{ik}(\theta_1, \theta_2, \dots, \theta_j | \alpha_1 + e_1, \alpha_2 + e_2, \dots, \alpha_j + e_j) \quad (12)$$

The CPT parameter point values are then found from the posterior using the expectation function as shown in Equation (9).

In describing the incremental updating algorithm, one must consider the various ways that data e can be observed. This is important as it impacts the manner in which the incremental updating algorithm adapts the CPT parameters. The four scenarios under which e can be observed

include (i) both the state of X_i and the configuration of $\text{Pa}(X_i)$ are observed, (ii) the state of X_i is unobserved while the configuration of $\text{Pa}(X_i)$ is observed, (iii) the state of X_i is observed while the configuration of $\text{Pa}(X_i)$ is unobserved, and (iv) the state of X_i and the configuration of $\text{Pa}(X_i)$ are both unobserved.

In the first scenario, in which all nodes are observed, the updating is relatively straight forward and only involves updating the prior to the posterior as given in Equation (12). However, when e has incomplete information on nodes, as in scenarios (ii)–(iv), the states of the unobserved nodes must be ‘filled-in’ and weighted by their probabilities. When this is the case, the posterior distribution can no longer be represented as a single Dirichlet distribution as shown in Equation (12), but instead it must be represented as a mixture of Dirichlet distributions. This mixture becomes more complex as the number of separate incomplete cases used in updating increases. To avoid this added complexity, the incremental updating algorithm approximates the mixture distribution by replacing it with a single Dirichlet distribution whose mean and variance match those of the mixture.

Using scenario (iv) as the most general case of incomplete data, the posterior distribution of the parameter vector can be shown to be

$$\Pr_{ik}(\theta_1, \theta_2, \dots, \theta_j | e) = \sum_{q=1}^k \sum_{r=1}^j \Pr(\theta | X_i(r), Pa_q(X_i), e) \times \Pr(X_i(r) | Pa_q(X_i), e) \times \Pr(Pa_q(X_i) | e) \quad (13)$$

Here the posterior of the parameters $(\theta_1, \theta_2, \dots, \theta_j)_{ik}$, having observed e , is calculated by summing over and eliminating, or integrating out, the unobserved X_i and $\text{Pa}_k(X_i)$ node states. In Equation (13), $X_i(r)$ implies that X_i is in state r where $r = 1, \dots, j$. Expanding Equation (13) shows this mixture of Dirichlet distributions as

$$\begin{aligned} \Pr_{ik}(\theta_1, \theta_2, \dots, \theta_j | e) = & \text{Dirichlet}(\theta_1, \theta_2, \dots, \theta_j | \alpha_1 + 1, \alpha_2, \dots, \alpha_j) \\ & \times \Pr(X_i(1) | Pa_k(X_i), e) \times \Pr(Pa_k(X_i) | e) \\ & + \text{Dirichlet}(\theta_1, \theta_2, \dots, \theta_j | \alpha_1, \alpha_2 + 1, \dots, \alpha_j) \\ & \times \Pr(X_i(2) | Pa_k(X_i), e) \times \Pr(Pa_k(X_i) | e) \\ & \vdots \\ & + \text{Dirichlet}(\theta_1, \theta_2, \dots, \theta_j | \alpha_1, \alpha_2, \dots, \alpha_j + 1) \\ & \times \Pr(X_i(j) | Pa_k(X_i), e) \times \Pr(Pa_k(X_i) | e) \\ & + \text{Dirichlet}(\theta_1, \theta_2, \dots, \theta_j | \alpha_1, \alpha_2, \dots, \alpha_j) \\ & \times \sum_{q \neq k} \sum_j \{ \Pr(X_i(j) | Pa_q(X_i), e) \times \Pr(Pa_q(X_i) | e) \} \quad (14) \end{aligned}$$

Equation (14) can be simplified further by recognizing that the double summation represents the probability of not observing the parent configuration k , given evidence e ,

which is

$$\begin{aligned} \sum_{q \neq k} \sum_j \{ \Pr(X_i(j) | Pa_q(X_i), e) \times \Pr(Pa_q(X_i) | e) \} \\ = 1 - \Pr(Pa_k(X_i) | e) \end{aligned} \quad (15)$$

The posterior represented by Equation (14) then becomes

$$\begin{aligned} \Pr_{ik}(\theta_1, \theta_2, \dots, \theta_j | e) = \sum_{r=1}^j \{ \text{Dirichlet}(\theta_1, \theta_2, \dots, \theta_j | \alpha_1, \dots, \alpha_r + 1, \dots, \alpha_j) \\ \times \Pr(X_i(r) | Pa_k(X_i), e) \times \Pr(Pa_k(X_i) | e) \\ + \text{Dirichlet}(\theta_1, \theta_2, \dots, \theta_j | \alpha_1, \alpha_2, \dots, \alpha_j) \\ \times (1 - \Pr(Pa_k(X_i) | e)) \} \quad (16) \end{aligned}$$

To reduce complexity, Spiegelhalter and Lauritzen (1990) suggested an approximation where the posterior mixture in Equation (16) is replaced with a single approximating Dirichlet distribution.

The incremental updating algorithm uses a moment matching technique so that the means and average variance of the single approximating Dirichlet are equal to those of the Dirichlet mixture. Spiegelhalter and Lauritzen (1990) recommend calculating the mean of the j th variable for the posterior mixture as

$$\begin{aligned} \bar{\theta}_j = \sum_{r=1}^j \{ \bar{\theta}_{rj} \times \Pr(X_i(r) | Pa_k(X_i), e) \times \Pr(Pa_k(X_i) | e) \} \\ + \bar{\theta}_j \times (1 - \Pr(Pa_k(X_i) | e)) \quad (17) \end{aligned}$$

Here $\bar{\theta}_{rj}$ is the expected value or mean of the j th parameter in the r th mixture component of Equation (16), while $\bar{\theta}_j$ is the expected value or mean of the j th parameter in the last component of Equation (16). The mean or expected values can be calculated using Equation (9). In a similar manner, the average variance of the j th parameter of the posterior mixture is given by

$$\begin{aligned} \bar{v}_j = \sum_{r=1}^j \left\{ \left(v_{rj} + (\bar{\theta}_{rj} - \bar{\theta}_j)^2 \right) \right. \\ \left. \times \Pr(X_i(r) | Pa_k(X_i), e) \times \Pr(Pa_k(X_i) | e) \right\} \\ + \left(v_{0j} + (\bar{\theta}_{0j} - \bar{\theta}_j)^2 \right) \times (1 - \Pr(Pa_k(X_i) | e)) \quad (18) \end{aligned}$$

where v_{rj} is the variance of the j th parameter of the r th mixture component of Equation (16), which can be calculated using Equation (10). The average variance of the variables resulting from the posterior mixture is given by the probability weighting

$$v_{ik} = \left(\sum_{r=1}^j \bar{\theta}_j \times \bar{v}_j \right)_{ik} \quad (19)$$

The variance of the single approximating Dirichlet is given by the probability weighting of Equation (10):

$$(\tilde{v})_{ik} = \left(\sum_{r=1}^j \bar{\theta}_j \times \frac{\bar{\theta}_j(1 - \bar{\theta}_j)}{\zeta_k + 1} \right)_i \quad (20)$$

By equating the average variance of the posterior mixture, with the variance of the single approximating Dirichlet, the corresponding posterior value for the new experience parameter value can be calculated as

$$\begin{aligned} \Rightarrow v_{ik} &= (\tilde{v})_{ik} \\ \Rightarrow \left(\sum_{r=1}^j \bar{\theta}_j \times \frac{\bar{\theta}_j(1-\bar{\theta}_j)}{\zeta_k+1} \right)_i &= \left(\sum_{r=1}^j \bar{\theta}_j \times \bar{v}_j \right)_{ik} \\ \Rightarrow \zeta_{ik} &= \left(\sum_{r=1}^j \frac{\bar{\theta}_j^2(1-\bar{\theta}_j)}{\bar{\theta}_j \times \bar{v}_j} - 1 \right)_{ik} \end{aligned} \quad (21)$$

The CPT values for the k th parent configuration are then updated with the $(j \times 1)$ vector of mean parameters, as calculated in Equation (17), while the experience value is updated with that calculated in Equation (21). Further discussions with regards to adaption algorithms can be found in Cowell (1999, pp 198–223), Jensen and Nielsen (2007, pp 207–214) and in Kjærulff and Madsen (2008, pp 209–211).

Experience and fading

Prior to any first parameter adaption, the experience parameters must be set with their initial values, which were drawn from the domain expert’s uncertainty levels, captured during the elicitation process.¹⁴ These experience levels ordered the nodes, states and conditional probabilities in terms of their judged uncertainty, whereby in invoking the law of large numbers, a lesser degree of uncertainty warranted a larger initial experience value, while a greater degree of uncertainty warranted a smaller initial experience value.

The final parameter associated with the parameter adaption process is the fading parameter $\eta_i = \{\eta_{i1}, \eta_{i2}, \dots, \eta_{ik}\}^T$, where each component has a domain in the closed interval [0, 1]. The fading variable is a $(k \times 1)$ vector, with one component corresponding to each of node X_i ’s parent configurations, k . The fading values η_i determine the long-run upper bound on the experience value. Bounding the experience values limits the influence that earlier event observations have in updating CPT parameters relative to more recent events. For example, assuming two fading values, $\eta_{ik} = 0.98$ and $\eta_{ik} = 0.998$, the long-run experience, ζ_k^* , has an upper bound of

$$(\zeta_k^*)_i = \left(\frac{1}{1 - \eta_k} \right)_i = \frac{1}{1 - 0.98} = 50 \quad (22)$$

¹⁴The uncertainty levels are ‘very high’, ‘high’, ‘medium’, and ‘low’ uncertainty, to which the knowledge engineer assigned equivalent experience values of 50, 100, 200, and 400, respectively.

and

$$(\zeta_k^*)_i = \frac{1}{1 - 0.998} = 500 \quad (23)$$

In these two cases, the experience value will not exceed the long-run value of 50 or 500, respectively. Hence, the adapted CPT parameters for node X_i , and parent configuration $\text{Pa}_k(X_i)$, would capture the influence of only the last 50 or 500 case events, respectively. In effect, the fading parameter provides a ‘forgetting’ facility for the adaption process. As the fading value is reduced, more recent domain events carry greater weight relative to past historical events. Having lower experience means that the variance of the prior Dirichlet is larger, as can be seen in Equation (10). As the experience parameter ζ_k becomes larger the variance of the probability parameter becomes smaller. When updating the posterior through new observations, if the variance of the prior is larger, the prior will be less concentrated and therefore will have less influence over the determination of the new posterior relative to the information contained within the observed data.

In selecting a preferred long-run experience value, ζ_k^* , the required fading parameter value can be calculated by rearranging Equation (22) to give

$$\eta_{ik} = \left(\frac{\zeta_k^* - 1}{\zeta_k^*} \right)_i \quad (24)$$

In the final Bayesian network model, the fading parameters were set to either 0.99 or 0.995, corresponding to a long-run experience of either 100 or 200 respectively. The choice was based on the subjective assessment of each node’s stationarity or stability and the uncertainty associated with their probability settings. Those nodes considered comparatively stable and certain were set at the higher fading and experience values, while those considered less stable and more uncertain were set at lower values. Although the fading parameters were set to these initial values, the risk manager still has the discretion to adjust them as required. Where the structural dynamics of the environment are changing, the risk manager can reduce the fading parameter to allow the model to adapt more quickly or, if the environment is stable, reset them to a higher value to ensure that the parameters adapt to signal rather than noise.

Forecasting loss distributions and operational value-at-risk

A key output of the tool’s central Bayesian network model is the forecasting of operational loss distributions over a specified time horizon. From these forecasts a percentile value known as the OpVaR (which represents the operational loss value not likely to be exceeded with a given probability) can be calculated. To specify the forecast horizon a forward projection of transaction volumes (the

driver of operational loss events) is required. If the risk manager required an OpVaR estimate at a 99% confidence level for the next 12 months, then a forward estimate of the transaction volume for those 12 months would be required. Using the transaction volume information, they would then use the Bayesian network model to simulate a set of alternative event histories, which could then be used to construct a loss distribution. From the constructed loss distribution, the corresponding 99th percentile (or OpVaR) could then be identified.

An advantage of using simulation is that any uncertainties associated with the forward transaction volume estimates, or model parameters for example, can be incorporated into the outputs. Event frequency and severity distributions generated from the transaction volume histories can also be useful in further validating the Bayesian network model by comparing predicted frequencies and severities with their observed values.

Before an operational loss distribution can be constructed, an appropriate distribution must be identified. For discrete frequency distributions, the common candidates include binominal, geometric and Poisson distributions, while for severity, continuous distributions (such as the exponential, lognormal, Weibull, gamma, beta, Pareto and Burr) are often used (see Chernobai *et al*, 2007). Although a more rigorous approach to distribution selection would require formal selection tests, such as a goodness-of-fit measure for each distribution candidate, a lognormal distribution is adopted to allow a forecasting demonstration of operational losses.

To be consistent with the parameter adaption process, Bayesian inference is used to estimate the operational loss distribution. Conjugate priors are used to keep the computation of parameter posteriors fast, efficient and comparatively straightforward. Therefore inference from priors to posteriors uses a set of simple updating rules that combine the prior's hyperparameters with the likelihood function's sufficient statistics.

A lognormal probability density function (pdf) can be parameterized as,

$$p(y|m, q) = \frac{\sqrt{q}}{y\sqrt{2\pi}} \exp\left(-\frac{q}{2}(\log y - m)^2\right) \quad (25)$$

Here $y > 0$ and the parameters m and q are also greater than zero. The conjugate prior over the parameters m and q is a gamma-normal joint distribution, which takes the form

$$p(m, q|\alpha, \beta, \tau, \mu) = \frac{q^{\alpha-1} \exp\left(-\frac{q}{\beta}\right)}{\Gamma(\alpha)\beta^\alpha} \times \sqrt{\frac{q\tau}{2\pi}} \exp\left(-\frac{q\tau}{2}(m - \mu)^2\right) \quad (26)$$

The hyperparameters of Equation (26) are $\alpha > 0$, $\beta > 0$, $\tau > 0$, and μ . The updating rules from prior to posterior

where posterior hyperparameters are identified by prime superscripts, as described in Fink (1997), are

$$\begin{aligned} \alpha' &= \alpha + \frac{n}{2}, & \beta' &= \left(\frac{1}{\beta} + \frac{1}{2}SS + \frac{\tau n(\bar{y} - \mu)^2}{2(\tau + n)}\right)^{-1} \\ \mu' &= \frac{\tau\mu + n\bar{y}}{\tau + n}, & \tau' &= \tau + n \end{aligned} \quad (27)$$

where

$$\begin{aligned} \bar{y} &= \frac{\sum_i \ln y_i}{n} \\ SS &= \sum_n (\ln y_i - \bar{y})^2 \end{aligned}$$

y_i in Equation (27) represents observed data, which are used to update the posteriors. When updating the parameters of the prior loss distribution, simulated y_i observations from the Bayesian network model are used to ensure that the probabilistic information stored within the network's structure and CPTs is reflected in the posterior loss distribution. The n is the sample size used in the parameter updating, and SS represents the sum of squared $\ln y_i$ deviations. Updating rules are quite standard, with further details available in Raiffa and Schlaifer (1961), DeGroot (1970) and Fink (1997).

To construct the lognormal operational loss distribution, the risk manager starts by generating a set of alternative transaction histories from the Bayesian network model. Each alternative history comprises a set of transactions. The number of transactions in each history corresponds with the forward estimate for the number of transactions SFO will process over the specified time period. The actual number of alternative histories generated is at the discretion of the risk manager, but the larger the number, the more information that is available for updating the loss distribution's parameter posteriors. The total dollar loss experienced in one alternative history is the sum of the individual losses simulated from the Actual OpRisk Loss nodes, shown in Figure 1 and Table 1. For one alternative history, the value of y_i is the total dollar loss experienced. Therefore, if one thousand alternative transaction histories were to be generated, the number of y_i observations would be one thousand, $i = 1, \dots, 1000$, and n in Equation (27) would equal one thousand. Depending on the loss distribution required by SFO management, the total dollar loss experienced in an alternative history will be the sum of individual Actual OpRisk Loss dollar amounts related to payment failures, exposure management failures, regulatory failures, a combination of these, or all of them. However, the individual loss types used for the total losses in each alternative history must be consistent. Following the generation of the simulated y_i data, and using the updating rules described in Equation (27), the posterior gamma-normal distribution over m and q can then be inferred.

The approach for generating the loss distribution and OpVaR represents a convolution of the frequency of operational loss events with their loss severity. This is true because each transaction in an alternative transaction history will comprise either no operational event, with of course no loss, or a single operational loss event and its associated loss amount. For the one entire alternative transaction history, the total number of events and the total loss will represent one transaction history sample, which is a mixture of loss events and loss amounts or severities. By generating a large number of alternative transaction histories this convolution can be captured in the predictive loss distribution.

A posterior predictive lognormal operational loss distribution is then constructed using the parameter values sampled from the gamma-normal posterior. Sampling from the posterior can be done first by sampling from the posterior gamma marginal distribution, $q|\alpha', \beta' \sim \text{Gamma}(\alpha', \beta')$, and then sampling m from the conditional posterior normal distribution $m|q, \mu', \tau' \sim \text{Normal}(\mu', 1/q)$. Generating a posterior predictive loss distribution means that the uncertainties associated with the lognormal parameters can be incorporated. As far as outputs are concerned, instead of producing a single OpVaR measure, a distribution over the measure can be generated that captures the uncertainty associated with it.

In using a Bayesian analysis to generate the operational loss distribution and OpVaR, priors over the data generating process parameters must be specified. In Bayesian inference, hyperparameter values of the prior are often selected so that as little prior information is included with regards to the model parameters as possible. This is done to ensure that the information contained within the observed data, not the information contained within the prior, dominates the updating of the posterior. However, the information content of the prior is less important when a large set of observations are available, so the information contained within the data overwhelms the information contained within the prior. In the next section we will discuss in more detail the selection of priors over the model parameters.

Construction of the frequency distribution, which shows the distribution of the number of loss events over a prescribed period, is not demonstrated despite being relatively straightforward. Using a Binominal distribution to represent the frequency distribution of operational loss events, the parameter of the binominal model can be estimated using output from the simulated alternative histories described. The observations in this case are not the loss total, but instead the counts of operational loss events within any given history. The set of observed counts, generated from the alternative histories, can be used to update a Beta-Binominal model, where the Beta distribution is the conjugate prior of the Binominal. The posterior Beta can then be used to sample parameters for the

posterior-predictive Binominal frequency distribution. Alternatively, instead of sampling parameters from the posterior Beta, the mean of this distribution can be used instead as the representative parameter of the posterior-predictive Binominal frequency distribution.

Despite the straightforward method described above to generate loss distributions and OpVaR values, the approach does have some disadvantages. The Bayesian network model is now detached or separated from the loss distribution model and OpVaR determination, even though the probabilistic information stored in the Bayesian network model is transmitted and summarized by the loss distribution. The disadvantage of this separation is that diagnostic inference cannot be performed, implying that the user cannot specify an observation on a child node and then infer the probable states of its parents and remaining nodes within the network that are d-connected. It is this inference, which is one of the strengths of Bayesian network models, as they are able to perform not only predictive inference from parent to child, but also diagnostic inference from child to parent. With the separation of the operational loss distribution model from the Bayesian network, only conditional and unconditional predictive inference can be performed. Retaining diagnostic inference would allow the risk manager to identify the most probable marginal distributions of causal antecedent nodes for a given loss distribution, as specified by the risk manager. This information would more easily assist the risk manager in identifying variables within the domain that would warrant operational improvement. Neil *et al* (2009) describe an operational risk model that allows the underlying Bayesian network model to incorporate the loss distribution directly. Neil *et al* (2009) use a Bayesian network that allows for dynamic discretization, which provides for greater flexibility in approximating continuous distributions, such as a loss distribution, through automated discretization.

Model application: a hypothetical case study

The following hypothetical case study describes in more detail the application of the Bayesian network model in generating a predictive operational loss distribution and the OpVaR. How the model adapts to a changing operational environment will also be illustrated. The SFO case study uses the following hypothetical scenario:

Operational Scenario: Due to a sudden increase in staff turnover in SFO, two new staff members, an operational manager and analyst, are appointed. The introduction of less experienced staff members into SFO leads to an increase in the incidence of transaction implementation errors, which leads to an increase in payment failures. On investigation, the risk manager finds these error events to have been in part the result of a decline in the effectiveness of oversight control.

During the operational period in which the scenario takes place, the SFO risk manager records the operational information related to SFO transaction processing. The operation event information is stored within the event capture database tool, which captures 116 transaction records.¹⁵ Of these 116 transactions, 61 were loans, nine deposits, two FX, 15 unlisted equities, four listed equities, five interest rate swaps, 11 preference equities and nine finance leases.

To compare the predictive performance of various adapted models in the case study scenario, the logarithmic score (LS) metric (Cowell *et al.*, 2007) is applied. The LS metric is defined as,

$$LS_M = \sum_m -\ln P_{(m)}(\varepsilon_m) \quad (28)$$

The LS metric in Equation (28) is the sum of the negative logarithm of the probability $P_{(m)}$ of observing the state ε_m of the m th event transaction. Observing events that have a low probability generates a higher score than events that are considered more likely. A better performing model will have less ‘surprise’ events and a lower LS score. Cowell *et al.* (2007) provide further details on model assessment, including absolute standardization and model monitors.

In order to gauge the impact of model adaption on the LS metric and the aggregate loss distributions produced, three separate adaption processes were performed producing three adapted Bayesian network models. Each of the three adaption processes begins from the same initial baseline model, B_0 , which is the complete, fully parameterised, Bayesian network model described in Figure 1 and Table 1.¹⁶

For each adaption process the 116 scenario event records are entered into the B_0 model sequentially, one record at a time. The LS and cumulative LS metrics is record for each record. In the first process, no adaption was used to provide a benchmark cumulative LS metric score. This model was designated as B_1 , which was identical to the B_0 model. In the second process, adaption was used, resulting in the model B_2 . The third process resulted in the adapted model B_3 . B_3 differed from B_2 in that the fading parameter values of the baseline model B_0 were adjusted downwards, creating upper bounds on the experience values that were 50% of those in the B_2 model. A node in B_2 that had a maximum experience value of 200 would have a maxi-

imum experience value in B_3 of 100. This adjustment was done to provide a demonstration of the fading factor’s impact on the adaption performance.

The cumulative LS_M for each adaption process were recorded at 45.4 for B_1 (no adaption), 44.5 for B_2 (adaption), and 43.5 for B_3 (adaption with (w/-) 50% effective sample size). The lower cumulative LS_M demonstrated the improved performance of the two adapted models, B_2 and B_3 , over the static model B_1 , registering fewer ‘surprise’ observations. The B_3 model, having the smallest LS_M metric, demonstrated an improved adaptive performance after its upper bounds on the experience value were constrained further.

After the adaption exercises, aggregate lognormal operational loss distributions were constructed. To estimate aggregate loss distributions, data were first simulated from each of the three models, B_1 , B_2 and B_3 . For each model, 300 distinct 12-month-long event histories were simulated, containing 100 transaction events.¹⁷ Using 300 histories was arbitrary, although motivated by the desire to ensure a large representation of ‘possible SFO worlds’. The number of histories used is ultimately at the discretion of the risk manager.

Using the simulated event histories, the parameters of the aggregate loss distribution were estimated. Because Bayesian inference is used, the risk manager must provide priors over the parameters of the lognormal model of the aggregate losses. Here we have assumed that the risk manager has an *a priori* set of beliefs regarding the parameters values, based on their previous aggregate loss experience within SFO.

The gamma-normal conjugate prior distribution is used to represent the risk manager’s prior beliefs over the parameters of the lognormal aggregate loss model. These prior distributions may be set directly by the risk manager, or instead, be based on the posterior of the previous period’s loss model.

The prior and posterior hyperparameter values for the gamma-normal distribution and the sufficient statistics used to update the prior settings to their posterior values are displayed in Table 7. The sufficient statistics in Table 7 are derived from B_1 ’s simulate data. Because B_1 was not adapted during the processing of the 116 transaction event records, it contains the *a priori* beliefs for SFO’s risk profile prior to the new transaction event observations. For this reason, data simulated from B_1 are used to construct the initial gamma-normal prior used to represent the risk manager’s prior beliefs. Two posteriors are inferred, one using data simulated from B_2 and the other from B_3 .

To provide further confirmation to the risk manager that the information contained within the prior reflects their subjective beliefs, a prior predictive aggregate loss

¹⁵The data used in the hypothetical case study were produced by simulating event records from the initial *a priori* Bayesian network model after evidence on transaction characteristics, human errors, error types, and operational event failures had been entered into the network. Subsequently, to reflect the changing risk profile of SFO, 10 of the generated loan transactions were modified to record payment failure events (caused by transaction implementation errors and oversight control errors) thus resulting in an apparent increase in these events.

¹⁶It was also model B_0 that was used to generate the hypothetical 116 event records for use in the hypothetical case study.

¹⁷The number of 100 transaction events per history approximated the expected count over a single year of SFO loan transaction activity.

Table 7 Gamma-normal prior to posterior parameter values

Prior				Sufficient statistics			Posterior			
α	β	μ	τ	\bar{y}	SS	n	α	β	μ	τ
10	1	11	50	11.276676	62.108010	235	127.5	0.029734	11.22814	285
100	0.01	11	200	11.276676	62.108010	235	217.5	0.007397	11.14947	435
10	1	12	5	11.276676	62.108010	235	127.5	0.029999	11.29175	240
10	0.5	10	5	11.276676	62.108010	235	127.5	0.026995	11.25008	240

Posterior parameters updated using simulated data from non-adapted B_1 model. The sufficient statistics include \bar{y} = average of $\ln y_i$ over sample observations, SS = sum of squared $\ln y_i$ deviations, and n = sample size.

Table 8 Prior predictive aggregate loss distribution and OpVaR (99%)

Prior				Aggregate loss distribution summary					OpVaR (99%)	
α	β	μ	τ	Mean	Standard deviation	25%	50%	99%	Mean	Standard deviation
10	1	11	50	\$63,481	\$21,660	\$48,049	\$60,007	\$130,489	\$130,626	\$19,985
100	0.01	11	200	\$99,565	\$131,245	\$30,547	\$60,053	\$622,724	\$623,517	\$86,081
10	1	12	5	\$172,915	\$58,561	\$131,219	\$163,571	\$353,904	\$353,959	\$91,757
10	0.5	10	5	\$25,130	\$12,505	\$16,420	\$22,463	\$67,446	\$67,572	\$21,092

Table 9 Posterior predictive aggregate loss distribution and OpVaR (99%) for non-adapted B_1 model

Posterior				Aggregate loss distribution summary					OpVaR (99%)	
α	β	μ	τ	Mean	Standard deviation	25%	50%	99%	Mean	Standard deviation
127.5	0.029734	11.22814	285	\$85,911	\$47,339	\$53,174	\$75,250	\$249,097	\$249,561	\$15,946
217.5	0.007397	11.14947	435	\$95,007	\$88,280	\$40,807	\$69,503	\$436,825	\$437,339	\$32,130
127.5	0.029999	11.29175	240	\$91,581	\$50,274	\$56,790	\$80,248	\$264,898	\$264,918	\$16,027
127.5	0.026995	11.25008	240	\$89,085	\$51,830	\$53,484	\$76,963	\$270,518	\$270,589	\$17,581

distribution is also generated. The prior predictive aggregate loss information is the result of integrating over the joint probability distribution of observed operational loss data, Y , and the parameters of the lognormal loss model, m and q . The prior predictive distribution is the integral

$$\Pr(Y) = \int_{R^+} \int_{R^+} \Pr(Y|m, q) \Pr(m, q) dq dm \quad (29)$$

Here $\Pr(m, q)$ is the gamma-normal prior distribution over the hyperparameters m and q and $\Pr(Y|m, q)$ is the lognormal loss model. The prior predictive distribution, $\Pr(Y)$, is obtained by integrating over the domains of m and q . Sample values of m and q were drawn from the gamma-normal prior $\Pr(m, q)$, followed by the sampling of an operational loss value Y from the lognormal distribution, indexed by the sampled values (m, q) parameters. Ten million losses were simulated for the prior predictive distribution $\Pr(Y)$, which are summarized in Table 8.

The summary information provided in Table 8 shows that each set of priors specifies a different set of beliefs regarding aggregate losses. For example, the prior with

hyperparameters $\alpha = 10$, $\beta = 1$, $\mu = 11$ and $\tau = 50$ produces a prior predictive aggregate loss distribution over a 12-month period with mean losses of \$63,481 and a standard deviation of \$21,660. The prior predictive 99% OpVaR has an expected value of \$130,489 and a standard deviation of \$19,985. Although it is usual for a quoted OpVaR to be a fixed loss amount (because the aggregate loss distribution is simulated over a range of possible parameter values, (m, q)), the parameter uncertainty can be taken into account that results in a distribution for the 99% OpVaR value. The remaining priors given in Table 8 show a range of possible OpVaRs, starting with a minimum of \$67,572 up to a maximum of \$623,517. More importantly, the posteriors results given in Table 9 show a convergence of OpVaRs, after the priors have been updated using the simulated loss data from B_1 .

The posterior predictive aggregate loss and OpVaR distributions are shown in Table 9. Similar to the prior predictive, the posterior predictive losses and OpVaRs are calculated via simulation, this time taking samples of the (m, p) parameters from the posterior gamma-normal distribution. Table 9 shows that most of the posterior

Table 10 Gamma-normal prior to posterior parameter values

Prior				Sufficient statistics			Posterior			
α	β	μ	τ	\bar{y}	SS	n	α	β	μ	τ
127.5	0.029734	11.22814	285	11.383787	67.644414	244	249.5	0.014483	11.29993	529
217.5	0.007397	11.14947	435	11.383787	67.644414	244	339.5	0.005770	11.23367	679
127.5	0.029999	11.29175	240	11.383787	67.644414	244	249.5	0.014778	11.33815	484
127.5	0.026995	11.25008	240	11.383787	67.644414	244	249.5	0.013899	11.31749	484

Posterior parameters updated using simulated data from adapted B_2 model. The sufficient statistics include \bar{y} = average of $\ln y_i$ over sample observations, SS = sum of squared $\ln y_i$ deviations, and n = sample size.

Table 11 Posterior predictive aggregate loss distribution and OpVaR (99%) for adapted B_2 model

Graph	Posterior				Aggregate loss distribution summary					OpVaR (99%)	
	α	β	M	T	Mean	Standard deviation	25%	50%	99%	Mean	Standard deviation
(a)	249.5	0.014483	11.29993	529	\$92,846	\$52,383	\$56,717	\$80,841	\$274,634	\$274,963	\$12,620
(b)	339.5	0.005770	11.23367	679	\$97,644	\$79,625	\$46,714	\$75,639	\$398,043	\$399,195	\$22,195
(c)	249.5	0.014778	11.33815	484	\$96,191	\$53,738	\$59,076	\$83,966	\$282,335	\$282,271	\$12,637
(d)	249.5	0.013899	11.31749	484	\$95,165	\$55,148	\$57,264	\$82,328	\$287,866	\$287,983	\$13,413

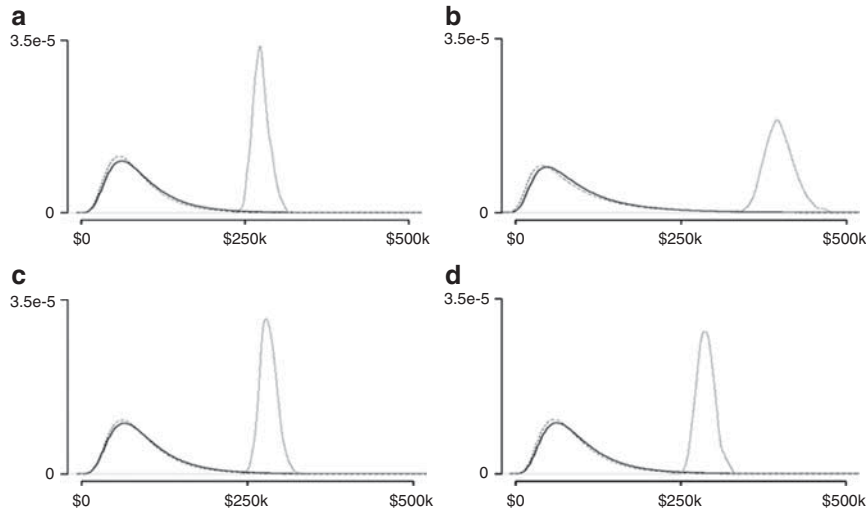


Figure 3 Posterior predictive aggregate loss and OpVaR distributions. Solid line is the posterior predictive aggregate loss distribution; dashed line is the prior predictive loss distribution. Distribution shown in the tail of the aggregate loss distribution is the posterior predictive OpVaR distribution at a confidence level of 99%.

predictive distributions of the OpVaR have now converged towards a range from \$249,561 to \$270,589, while the larger of the OpVaRs has shrunk from \$623,517 to \$437,339. The posterior parameter settings shown in Table 9 are used as the prior settings for updating the aggregate loss distribution based on the data simulated from B_2 and B_3 .

Table 10 displays the posterior hyperparameter values inferred from the prior and sufficient statistics generated from the B_2 model. Summary details of the posterior

predictive aggregate loss and OpVaR are found in Table 11, where it can be seen that all aggregate loss distributions have achieved mean losses in the range \$92,000—\$98,000. The changes in the posterior aggregate loss and OpVaR distributions can be seen graphically in Figure 3. As expected, given the increasing risk profile for SFO under the case study scenario, payment failures during loan transaction events have increased with the result that the mean aggregate losses have all increased upward. The OpVaRs based on the B_1 data have increased

Table 12 Gamma-normal prior to posterior parameter values

Prior				Sufficient statistics			Posterior			
α	β	μ	τ	\bar{y}	SS	n	α	β	μ	τ
127.5	0.029734	11.22814	285	11.620052	91.655200	272	263.5	0.011093	11.41952	557
217.5	0.007397	11.14947	435	11.620052	91.655200	272	353.5	0.005011	11.33051	707
127.5	0.029999	11.29175	240	11.620052	91.655200	272	263.5	0.011623	11.46616	512
127.5	0.026995	11.25008	240	11.620052	91.655200	272	263.5	0.010917	11.44663	512

Posterior parameters updated using simulated data from the adapted B_3 model. The sufficient statistics include \bar{y} = average of $\ln y_i$ over sample observations, SS = sum of squared $\ln y_i$ deviations, and n = sample size.

Table 13 Posterior predictive aggregate loss distribution and OpVaR (99%) for adapted B_3 model

Graph	Posterior				Aggregate loss distribution summary					OpVaR (99%)	
	α	β	M	T	Mean	Standard deviation	25%	50%	99%	Mean	Standard deviation
(e)	263.5	0.011093	11.41952	557	\$108,253	\$69,196	\$61,425	\$91,184	\$356,026	\$356,381	\$17,912
(f)	353.5	0.005011	11.33051	707	\$110,480	\$96,261	\$50,105	\$83,249	\$478,553	\$479,814	\$25,846
(g)	263.5	0.011623	11.46616	512	\$112,456	\$69,934	\$64,932	\$95,484	\$360,773	\$361,590	\$17,496
(h)	263.5	0.010917	11.44663	512	\$111,313	\$71,845	\$62,800	\$93,515	\$368,951	\$369,505	\$18,517




Figure 4 Posterior predictive aggregate loss and OpVaR distributions. Solid line is posterior predictive aggregate loss distribution; dashed line is prior predictive loss distribution. Distribution shown in the tail of the aggregate loss distribution shows the posterior predictive OpVaR distribution at a confidence level of 99%.

from \$249,561 to \$274,963, \$264,918 to 282,271, and \$270,589 to \$287,983, after updating with the B_2 data, while the higher OpVaR has declined from \$437,339 to \$399,195. The standard deviations of OpVaRs have also reduced.

Finally, Tables 12 and 13 show the posterior hyperparameters and posterior predictive aggregate loss and OpVaR distributions inferred from the prior generated from B_1 and the sufficient statistics from B_3 's simulated data. The results in Table 13 show that B_3 , having adapted more readily to the changing SFO environment than B_2 , is now producing posterior predictive aggregate loss distributions whose

central locations have moved further outwards, with mean losses now occupying a range from \$108,000 to \$113,000. The 99% percentile has also extended outwards by nearly another \$80,000, while the OpVaRs increased to a new range of \$356,381 through to \$479,814. Figure 4 shows in graphically the changes in the posterior predictive aggregate losses and OpVaR distributions, which are more extreme than those displayed in Figure 3.

In the case study, the risk manager must decide on which of the aggregate loss distributions reflect SFO's new risk profile. They can check their prior beliefs of SFO's risk profile by generating *a priori* predictive aggregate losses

and OpVaRs based on a set of alternative gamma-normal priors. If they believed that the prior ($\alpha = 127.5$, $\beta = 0.029734$, $\mu = 11.22814$, $\tau = 285$) best reflects their prior beliefs, their likely preferred posterior result would be either (a) in Table 11 (based on B_2) or (e) in Table 13 (based on B_3), depending on their choice of fading factor.

Conclusions and future research

In this research, an operational risk modelling tool has been constructed and its application demonstrated. It is designed to support a business unit level operational risk manager in recording, modelling, analysing, communicating, and predicting operational losses. Central to the tool is a Bayesian network model that encapsulates the probabilistic and causal relationships between the key risk factors and operational losses within the domain. Application of the tool has been demonstrated, with predictive operational loss distributions and OpVaRs being generated from probabilities stored within the Bayesian network model.

The central model underlying the tool is sufficiently modular to allow it to support the inclusion of future, alternative causal antecedents, as deemed necessary by the operational risk experts within SFO, without the need for costly re-modelling. This modularity also supports the application of the operational risk tool to other similar operational environments.

The initial construction of the Bayesian network model was carried out using probability elicitation, parameter tuning and model validation techniques. The initial development placed considerable demands on the SFO staff involved, which prompted a search for opportunities to increase the level of automated model adaption to aid the tool's response to new events and the changing risk dynamics within the SFO. The current automation of model adaption has been restricted to parameter updating only, although automated structural adaption will be investigated in future research for the purpose of improving the tool.

As was mentioned earlier in the description of the model, continuous valued nodes were approximated by static discretization, achieving a satisfactory representation of the environment, rather than an optimal one. A number of authors have suggested alternative techniques for approximating continuous nodes that warrant future investigation. These alternatives include dynamic discretization (Neil *et al.*, 2008), mixtures of truncated exponentials (Cobb and Shenoy, 2008) and variational methods (Murphy, 1999). For any alternative approach, however, given the central importance to the operational risk tool of automated model adaption, a critical selection criterion will be the availability of efficient model adaption algorithms.

Acknowledgements—The authors would like to acknowledge the financial support provided through two research grants awarded by the

Department of Accounting and Finance, Monash University, and the Melbourne Centre for Financial Studies. The second author is also supported by an ARC Discovery grant for which he is grateful.

References

- Adusei-Poku K, Van den Brink GJ and Zucchini W (2007). Implementing a Bayesian network for foreign exchange settlement: A case study in operational risk management. *Journal of Operational Risk* 2(2): 101–107.
- Alexander C (2000). *Bayesian methods for measuring operational risk*. University of Reading, Discussion Papers in Finance, No 2000-02.
- Alexander C (2003). Managing operational risks with Bayesian networks. In: Alexander C (ed) *Operational Risk: Regulation, Analysis and Management*. Financial Times Prentice-Hall: London, pp 285–295.
- Anderson R, Mackoy R, Thompson V and Harrell G (2004). A Bayesian network estimation of the service-profit chain for transport service satisfaction. *Decision Sci* 35(4): 665–689.
- Aquaro V, Bardoscia M, Bellotti R, Consiglio A, De Carlo F and Ferri G (2010). A Bayesian networks approach to operational risk. *Physica A* 389(8): 1721–1728.
- Baron J (2008). *Thinking and Deciding*. 4th edn, Cambridge University Press: New York.
- Basel Committee on Bank Supervision (2004). *International Convergence of Capital Measurement and Capital Standards: A Revised Framework*. Bank for International Settlements: Basel.
- Bedford T and Cooke RM (2001). *Probabilistic Risk Analysis: Foundations and Methods*. Cambridge University Press: Cambridge.
- Bednarski M, Cholewa W and Frid W (2004). Identification of sensitivities in Bayesian networks. *Engineering Applications of Artificial Intelligence* 17(4): 327–335.
- Bobbio A, Portinale L, Minichino M and Ciancamerla E (2001). Improving the analysis of dependable systems by mapping fault trees into Bayesian networks. *Reliability Engineering & System Safety* 71(3): 249–260.
- Bobbio A, Ciancamerla M, Franceschinis G, Gaeta R, Minichino M and Portinale L (2003). Sequential application of heterogeneous models for the safety analysis of a control system: A case study. *Reliability Engineering & System Safety* 81(3): 269–280.
- Bonafede CE and Giudici P (2007). Bayesian networks for enterprise risk assessment. *Physica A* 382(1): 22–28.
- Boudali H and Dugan JB (2005). A discrete-time Bayesian network reliability modeling and analysis framework. *Reliability Engineering & System Safety* 87(3): 337–349.
- Bromley J, Jackson NA, Clymer OJ, Giacomello AM and Jensen FV (2005). The use of hugin[®] to develop Bayesian networks as an aid to integrated water resource planning. *Environmental Modelling & Software* 20(2): 231–242.
- Castillo E, Gutierrez JM and Hadi AS (1995). Parametric structure of probabilities in Bayesian networks. In: Froidevaux C and Kohlas J (eds) *Lectures Notes in Artificial Intelligence: Symbolic and Quantitative Approaches to Reasoning and Uncertainty*. Vol. 946, Springer-Verlag: New York, pp 89–98.
- Castillo E, Gutierrez JM and Hadi AS (1996). A new method for efficient symbolic propagation in discrete Bayesian networks. *Networks* 28(1): 31–43.
- Castillo E, Gutierrez JM and Hadi AS (1997). Sensitivity analysis in discrete Bayesian networks. *IEEE Transactions on Systems, Man and Cybernetics A* 27(4): 412–423.
- Chernobai AS, Rachev ST and Fabozzi FJ (2007). *Operational Risk: A Guide to Basel II Capital Requirements, Models and Analysis*. John Wiley & Sons: Hoboken, NJ.

- Cobb BR and Shenoy PP (2008). Decision making with hybrid influence diagrams using mixtures of truncated exponentials. *European Journal of Operational Research* **186**(1): 261–275.
- Cooke RM and Goossens LHJ (1999). *Procedures for structured expert judgement*. Report Prepared under Contract Number ETNU-CT93-0104-NL Report for the Commission of European Communities Directorate-General XI (Environment and Nuclear Safety) Directorate D (DG 11 CCJH).
- Cooke RM and Goossens LHJ (2004). Expert judgement elicitation for risk assessments of critical infrastructure. *Journal of Risk Research* **7**(6): 643–656.
- Cooper RG (2000). Strategic marketing planning for radically new products. *Journal of Marketing* **64**(1): 1–16.
- Cornalba C (2009). Clinical and operational risk: A Bayesian approach. *Methodology and Computing in Applied Probability* **11**(1): 47–63.
- Cornalba C and Giudici P (2004). Statistical models for operational risk management. *Physica A* **338**(1): 166–172.
- Coupé VMH and van der Gaag LC (2002). Properties of sensitivity analysis of Bayesian belief networks. *Annals of Mathematics and Artificial Intelligence* **36**(4): 323–356.
- Coupé VMH, Peek N, Ottenkamp J and Habbema JDF (1999). Using sensitivity analysis for efficient quantification of a belief network. *Artificial Intelligence in Medicine* **17**(3): 223–247.
- Cowell RG (1999). *Probabilistic Networks and Expert Systems*. Springer: New York.
- Cowell RG, Verrall RJ and Yoon YK (2007). Modeling operational risk with Bayesian networks. *Journal of Risk Insurance* **74**(4): 795–827.
- Davis GA (2003). Bayesian reconstruction of traffic accidents. *Law, Probability and Risk* **2**(2): 69–89.
- DeGroot MH (1970). *Optimal Statistical Decisions*. McGraw-Hill Book Company: New York.
- Fenton N and Neil M (2000). Bayesian belief networks: A causal model for predicting defect rates and resource requirements. *Software Testing and Quality Engineering* **2**(1): 48–53.
- Fenton N, Littlewood B, Neil M, Strigini L, Sutcliffe A and Wright D (1998). Assessing dependability of safety critical systems using diverse evidence. *IEE Proceedings – Software Eng* **145**(1): 35–39.
- Fink D (1997). A compendium of conjugate priors. Available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.157.5540>.
- Granger-Morgan M and Henrion M (1990). *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press: Cambridge.
- Gelman A, Carlin B, Stern HS and Rubin DB (2004). *Bayesian Data Analysis*. 2nd edn, Chapman & Hall/CRC: Boca Raton, FL.
- Groth K, Wang C and Mosleh A (2010). Hybrid causal methodology and software platform for probabilistic risk assessment and safety monitoring of socio-technical systems. *Reliability Engineering & System Safety* **95**(12): 1276–1285.
- Jensen FV (1999). Gradient descent training of Bayesian networks. In: Hunter A and Parsons S (eds) *Proceedings of the Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning under Uncertainty, Lecture Notes in Artificial Intelligence*. Springer-Verlag: New York, pp 190–200.
- Jensen FV and Nielsen TD (2007). *Bayesian Networks and Decision Graphs*. Springer Science + Business Media, LLC: New York.
- Kadane JB and Schum DA (1996). *A Probabilistic Analysis of the Sacco and Vanzetti Evidence*. Wiley: New York.
- Kemmerer B, Mishra S and Shenoy PP (2002). *Bayesian causal maps as decision aids in venture capital decision making: Methods and applications*. Working Paper, University of Kansas.
- Khodakarami V, Fenton N and Neil M (2007). Project scheduling: Improved approach to incorporate uncertainty using Bayesian networks. *Project Management Journal* **38**(2): 39–49.
- Kjerulf U and Madsen A (2008). *Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis*. Springer: New York.
- Koller D and Friedman N (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press: Cambridge, MA.
- Korb KB and Nicholson AE (2004). *Bayesian Artificial Intelligence*. Chapman & Hall/CRC: Boca Raton, FL.
- Lancaster T (2004). *An Introduction to Modern Bayesian Econometrics*. Blackwell Publishing: Padstow (Cornwall).
- Lauritzen SL and Spiegelhalter DJ (1988). Local computations with probabilities on graphical structures and their application to expert systems (with discussion). *Journal of the Royal Statistical Society Series B* **50**(2): 157–224.
- Lucas PJ, van der Gaag LC and Abu-Hanna A (2004). Bayesian networks in biomedicine and health-care. *Artificial Intelligence in Medicine* **30**(3): 201–214.
- Marquez D, Neil M and Fenton N (2010). Improved reliability modeling using Bayesian networks and dynamic discretization. *Reliability Engineering & System Safety* **95**(4): 412–425.
- Mittnik S and Starobinskaya I (2010). Modeling dependencies in operational risk with hybrid Bayesian networks. *Methodology and Computing in Appl* **12**(3): 379–390.
- Moosa IA (2008). *Quantification of Operational Risk under Basel II: The Good, Bad and Ugly*. Palgrave Macmillan: London.
- Moosa IA (2010). Basel II as a casualty of the global financial crisis. *Journal of Banking Regulation* **11**(2): 95–114.
- Moosa IA (2011). Basel II and Basel III: A great leap forward? In: La Brosse JR, Olivares-Caminal R and Singh D (eds) *Managing Risk in the Financial System*. Edward Elgar: Cheltenham.
- Moosa IA (2012). Basel 2.5: A lot of sizzle but little nutritional value. *Journal of Banking Regulation* **13**(4): 320–335.
- Moosa IA and Burns K (2012). Basel III as a regulatory response to the global financial crisis. *International Journal of Applied Business and Economic Research* **10**(1): 31–44.
- Murphy K (1999). A variational approximation for Bayesian networks with discrete and continuous latent variables. In: Laskey K and Prade H (eds) *Uncertainty in Artificial Intelligence*, Vol. 15, Morgan Kaufmann Publishers Inc: Burlington, MA, pp 467–475.
- Neapolitan RE (2004). *Learning Bayesian Networks*. Prentice-Hall: Harlow.
- Neil M, Fenton N and Nielsen L (2000). Building large-scale Bayesian networks. *Knowledge Engineering Review* **15**(3): 257–284.
- Neil M, Fenton N, Forey S and Harris R (2001). Using Bayesian belief networks to predict the reliability of military vehicles. *Computing and Control Engineering* **12**(1): 11–20.
- Neil M, Fenton N and Tailor M (2005). Using Bayesian networks to model expected and unexpected operational losses. *Risk Analysis* **25**(4): 963–972.
- Neil M, Tailor M, Marquez D, Fenton N and Hearty P (2008). Modelling dependable systems using hybrid Bayesian networks. *Reliability Engineering & System Safety* **93**(7): 933–939.
- Neil M, Häger D and Andersen LB (2009). Modeling operational risk in financial institutions using hybrid dynamic Bayesian networks. *Journal of Operational Risk* **4**(1): 3–33.
- Nordgård DE and Sand K (2010). Application of Bayesian networks for risk analysis of MV air insulated switch operation. *Reliability Engineering & System Safe* **95**(12): 1358–1366.
- Paté-Cornell ME and Guikema S (2002). Probabilistic modeling of terrorist threats: A systems analysis approach to setting priorities among countermeasures. *Military Operations Research* **7**(4): 5–20.

- Pearl J (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann: San Mateo, CA.
- Raiffa H and Schlaifer R (1961). *Applied Statistical Decision Theory*. Division of Research, Graduate School of Business Administration, Harvard University: Boston, MA.
- Reason J (1990). *Human Error*. Cambridge University Press: Cambridge.
- Reid GB and Nygren TE (1988). The subjective workload assessment technique: A scaling procedure for measuring mental workload. In: Hancock PA and Meshkati N (eds) *Human Mental Workload*. North-Holland: Amsterdam, pp 185–218.
- Renooij S (2001). Probability elicitation for belief networks: Issues to consider. *Knowledge Engineering Review* 16(3): 255–269.
- Russell SJ, Binder J, Koller D and Kanazawa K (1997). Adaptive probabilistic networks with hidden variables. *Machine Learning* 29(2): 213–244.
- Sanford A and Moosa IA (2012). A Bayesian network structure for operational risk modelling in structured finance operations. *Journal of the Operational Research Society* 63(4): 431–444.
- Sigurdsson JH, Walls IA and Quigley JL (2001). Bayesian belief networks for managing expert judgement and modelling reliability. *Quality and Reliability Engineering International* 17(3): 181–190.
- Simon HA (1996). *The Sciences of the Artificial*. 3rd edn, MIT Press: Cambridge, MA.
- Spiegelhalter DJ and Lauritzen SL (1990). Sequential updating of conditional probabilities on directed graphical structures. *Networks* 20(5): 579–605.
- Taroni F, Aitken C, Garbolino P and Biedermann A (2006). *Bayesian Networks and Probabilistic Inference in Forensic Science*. Wiley: New York.
- Trucco P, Cagno E, Ruggeri F and Grande O (2008). A Bayesian belief network modelling of organisational factors in risk analysis: A case study in maritime transportation. *Reliability Engineering & System Safety* 93(6): 823–834.
- Uusitalo L (2007). Advantages and challenges of Bayesian networks in environmental modelling. *Ecological Model* 203(3): 312–318.
- van der Gaag LC, Renooij S, Witteman CLM, Aleman BMP and Taal BG (2002). Probabilities for a probabilistic network: A case study in oesophageal cancer. *Artificial Intelligence in Medicine* 25(2): 123–148.
- van der Gaag LC, Renooij S and Coupé VM (2007). Sensitivity analysis of probabilistic networks. In: *Advances in Probabilistic Graphical Models*. Springer: Berlin, Heidelberg, pp 103–124.
- Weber P and Jouffe L (2006). Complex system reliability modelling with dynamic object oriented Bayesian networks (DOOBN). *Reliability Engineering & System Safety* 91(2): 149–162.
- Werdigier J (2008). Trading Scandal Diverts Attention from Société Générale Subprime Losses. *New York Times*, 29 January.
- Wickens CD and Hollands G (2000). *Engineering Psychology and Human Performance*. 3rd edn, Prentice-Hall: New Jersey.
- Willems A, Janssen M, Verstege C and Bedford T (2005). Expert quantification of uncertainties in a risk analysis for an infrastructure project. *Journal of Risk Research* 8(1): 3–17.
- Wilson AG and Huzurbazar AV (2007). Bayesian networks for multilevel system reliability. *Reliability Engineering & System Safety* 92(10): 1413–1420.
- Woodberry O, Nicholson AE, Korb KB and Pollino C (2004). Parameterising Bayesian networks. In: Webb GI and Xinghuo Yu (eds) *AI 2004 LNAI 3339*. Springer-Verlag Berlin: Heidelberg, pp 1101–1107.

Appendix A

Human error definitions

The following human error categories and definitions have been taken from Reason (1990) and Wickens and Hollands (2000).

Mistake

A mistake is defined as an error that occurs when an operator fails to formulate the correct intentions. Possible causes include failures in perception, memory, and/or cognition. Normally ‘mistakes’ can be further categorized into ‘knowledge’-based mistake, that is, Decision Making, where an incorrect plan of action is developed due to a failure to understand the operational situation. Alternatively, ‘rule’-based mistakes can also occur, for example, when an operator knows or believes they know the situation and invoke a rule or plan of action to deal with it. These action rules usually take the form of an ‘IF-THEN’ logic. The rule-based mistake occurs when a ‘Good’ rule is misapplied, that is, when the IF condition does not apply to the current environment or when a ‘Bad’ rule is learned and subsequently applied.

Slip

Slips are defined as errors that involve the correct intention implemented incorrectly. Capture errors are a common class of slip in which an intended stream of behaviour is ‘captured’ by a similar, well-practiced behaviour pattern. Slips occur under three different scenarios. The first scenario is where the intended action or action sequence involves a slight departure from a routine, or frequently performed action. The second is where some characteristics of the stimulus environment or the action sequence itself are related to the now inappropriate action. And thirdly, the action sequence is relatively automated and therefore not monitored closely by the operator’s attention.

Lapse

A lapse error is defined as the failure to carry out an action. Lapses can be directly related to failures of memory. They are however different from knowledge-based mistakes associated with the overloading of operators working memory, resulting in poor decision making. Important lapses may involve the omission of steps in a procedural sequence. In this situation, an interruption is what often causes the sequence to be stopped,

then restarted again a step or two later than it should have been.

Mode error

A mode error is closely related to a slip, but also has the memory failure characteristics of lapses. A mode error results when a particular action that is highly appropriate in one mode of operation is performed in a different, inappropriate mode because the operator has not correctly remembered the appropriate context.

Appendix B

Workload definitions

The following workload definitions are drawn from the SWAT described in Wickens and Hollands (2000) and Reid and Nygren (1988). The following table is reproduced from Wickens and Holland's Table 11.1 (Wickens and Hollands, 2000, p 467). The SWAT comprises three dimensions, Time Load, Mental Effort Load and Stress Load, each of which has a rating scale 1–3. Combination of each to an overall workload rating is achieved by adding all ratings together.

Time load

Measures the time pressures experienced by the operator during the current task.

Rating

- 1 = Often have spare time. Interruptions or overlap among activities occur infrequently or not at all.
- 2 = Occasionally have spare time. Interruptions or overlap among activities occur frequently.
- 3 = Almost never have spare time. Interruptions or overlap among activities are very frequent, or occur all the time.

Mental effort load

Measures the mental load experienced by the operator involved in the current tasks.

Rating

- 1 = Very little conscious mental effort or concentration required. Activity is almost automatic, requiring little or no attention.
- 2 = Moderate conscious mental effort or concentration required. Complexity of activity is moderately high due to uncertainty, unpredictability, or unfamiliarity. Considerable attention required.
- 3 = Extensive mental effort and concentration necessary. Very complex activity requiring total attention.

Stress load

Measures the psychological load experienced by the operator during the current tasks.

Rating

- 1 = Little confusion, risk, frustration, or anxiety exists and can be easily accommodated.
- 2 = Moderate stress due to confusion, frustration, or anxiety noticeably adds to the workload. Significant compensation is required to maintain adequate performance.
- 3 = High to very intense stress due to confusion, frustration, or anxiety. High to extreme determination and self-control required.

Work load states

The states are defined under the Work Load node and present an overall environmental impression combining the time, mental effort, and stress load factors. The overall work load state is a simple summed combination of the individual factor states. Hence the work load states are from 3 to 9.

<i>State</i>	<i>Description of work load states</i>
3–<5	<p>Lower Bound Working Environment (3):</p> <p>(a) Often have spare time. Interruptions or overlap among activities occur infrequently or not at all. (b) Very little conscious mental effort or concentration required. Activity is almost automatic, requiring little or no attention. (c) Little confusion, risk, frustration, or anxiety exists and can be easily accommodated.</p> <p>Upper Bound Working Environment <5:</p> <p>Replace any of the above with two from</p> <p>(a) Occasionally have spare time. Interruptions or overlap among activities occur frequently. (b) Moderate conscious mental effort or concentration required. Complexity of activity is moderately high due to uncertainty, unpredictability, or unfamiliarity. Considerable attention required. (c) Moderate stress due to confusion, frustration, or anxiety noticeably adds to the workload. Significant compensation is required to maintain adequate performance.</p>
5–<7	<p>Lower Bound Working Environment (5): (a) Occasionally have spare time. Interruptions or overlap among activities occur frequently. (b) Moderate conscious mental effort or concentration required. Complexity of activity</p>

<i>State</i>	<i>Description of work load states</i>
	<p>is moderately high due to uncertainty, unpredictability, or unfamiliarity. Considerable attention required. (c) Moderate stress due to confusion, frustration, or anxiety noticeably adds to the workload. Significant compensation is required to maintain adequate performance.</p> <p>Upper Bound Working Environment (< 7): Replace any of the above with any one from (a) Almost never have spare time. Interruptions or overlap among activities are very frequent, or occur all the time. (b) Extensive mental effort and concentration necessary. Very complex activity requiring total attention. (c) High to very intense stress due to confusion, frustration or anxiety. High to extreme determination and self-control required.</p>
7-9	<p>Lower Bound Working Environment (7): Replace any two of the following with two from below (a) Occasionally have spare time. Interruptions or overlap among activities occur frequently. (b) Moderate conscious mental effort or concentration required. Complexity of activity is moderately high due to uncertainty, unpredictability, or unfamiliarity. Considerable attention required. (c) Moderate stress due to confusion, frustration, or anxiety noticeably adds to the workload. Significant compensation is required to maintain adequate performance.</p> <p>Upper Bond Working Environment (9): (a) Almost never have spare time. Interruptions or overlap among activities are very frequent, or occur all the time. (b) Extensive mental effort and concentration necessary. Very complex activity requiring total attention. (c) High to very intense stress due to confusion, frustration or anxiety. High to extreme determination and self-control required.</p>

Appendix C

TMG model operational error definitions

Data integrity error

Definition: Error involving the data capture of transaction information either into a manual/spread sheet-based system or an automated bank legacy information system.

Transaction implementation error

Definition: Error involving the implementation of a transaction. Note if the error created is due to a data capture error, then this is not an implementation error but a data integrity error.

Oversight control error

Definition: Error involving oversight control. This could include reconciliation control, 'four eyes, six eyes' review, signoffs, limit control, discretion levels etc.

System error

Definition: Error involving the failure of a bank legacy information system. Such an error could involve system availability, program error, hardware fault, power outage etc.

*Received April 2010;
accepted March 2013 after two revisions*

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.