

Modelling Operational Risk in Financial Institutions using Hybrid Dynamic Bayesian Networks

Authors:

Professor Martin Neil

Department of Computer Science, Queen Mary University of London, Mile End Road,
London, E1 4NS, United Kingdom

Phone: +44 (0) 207 882 5221

E-mail: martin@dcs.qmul.ac.uk

Agena Ltd., Hatton Garden, London EC1N 8DL, United Kingdom

E-mail: martin@agenaco.uk

Dr. Lasse B. Andersen

Department of Industrial Economics, Risk Management and Planning, University of
Stavanger, 4036 Stavanger, Norway

Phone: +47 51 83 16 70

E-mail: lasse.b.andersen@uis.no

David Häger, PhD Student

Department of Industrial Economics, Risk Management and Planning, University of
Stavanger, 4036 Stavanger, Norway

E-mail: david.hager@uis.no

Phone: +47 51 83 14 48, mob: +47 98 88 00 50

Executive Summary

This paper describes the use of Hybrid Dynamic Bayesian Networks (HDBNs) to model operational risk in an AMA context. The approach focuses on cause-effect modelling including interactions between failure modes and controls. Value at Risk is calculated by applying a new state-of-the-art HDBN algorithm that approximates continuous loss distributions and aggregates across loss types. In order to illustrate the natural match between the model and the underlying process, including the causal complexity underlying known and possible severe operational risk losses, we apply the generalised model to a financial trading example — rogue trading. We conclude that the statistical properties of the model have the potential to explain recent large scale loss events and offer improved means of loss prediction.

Keywords: Operational Risk Management, Bayesian Networks, Causal Models, Dynamic Discretization, Basel II, Advanced Measurement Approach, Rogue Trading.

Modelling Operational Risk in Financial Institutions using Hybrid Dynamic Bayesian Networks

David Häger[‡], Martin Neil[†] and Lasse B. Andersen[‡]

[‡] University of Stavanger

[‡] Bayes Risk Management AS

[†] Department of Computer Science, Queen Mary, University of London

[†] Agena Ltd

Abstract

This paper describes the use of Hybrid Dynamic Bayesian Networks (HDBNs) to model the operational risk faced by financial institutions in terms of economic capital. It describes a methodology for modelling financial losses resulting from intentional or accidental events and characterises these by their ability to evade controls and ultimately lead to increasingly severe financial consequences. The approach presented focuses on modelling the causes and effects of loss events using a Dynamic Bayesian Network model based on interactions between failure modes and controls. To calculate the Value at Risk (VaR) for total losses we apply a new state-of-the-art Hybrid Bayesian Network algorithm, called dynamic discretization. The algorithm approximates the continuous loss distribution functions required for each loss event at each point in time and is used to aggregate across loss types. In order to illustrate the natural match between the model and the underlying process, including the causal complexity underlying known and possible severe operational risk losses, we apply the generalised model to a financial trading example — rogue trading. We conclude that the statistical properties of the model have the potential to explain recent large scale loss events and offer improved means of loss prediction.

1. Introduction

The Basel Committee on Banking Supervision, in reaction to a number of well-publicised financial disasters (e.g. Barings bank 1995, Daiwa bank 1995, and Allied Irish Bank 2002) has mandated a system of regulation addressing the issue of Operational Risk (OpRisk) and its assessment and management (Basel 2006). Key to the regulatory process is the need for financial institutions to model their operational risk, in terms of a variety of loss event types (including unauthorized trading) in order to effectively manage risk and establish an appropriate regulatory capital charge.

Basel II defines three approaches to establishing regulatory capital for OpRisk: The Basic Indicator Approach (BIS), the Standardized Approach (SA) and the Advanced Measurement Approach (AMA). AMA is by far the most demanding approach, but also the most rewarding in terms of potentially reducing the OpRisk regulatory capital charge by as much as 20-40% compared to BIS and SA. An AMA model requires approval by the local Financial Supervisory Authority and must at least contain the following important model attributes: quantification of expected/unexpected losses at a given

confidence level; a one year holding period; include scenario analyses; reflect day to day OpRisk management practice.

Of course, operational risk problems are not peculiar to the financial sector and operational risk is not a new topic. In his book, James Reason argues that operational risk is faced by all organisations and he uses examples from the Financial, Rail Transport, Civil Aviation and Nuclear power sectors (Reason 1997) and he concludes that accidents are not solely the result of human fallibility but are supported by organisational features that fail to defend against all-too-human mistakes, slips and (in the case of fraud) malicious acts. Human error might be seen as an onset of a catastrophic event, but without latent weaknesses within the organisation the event would not reach catastrophic proportions. From this we conclude that operational risk prediction is inextricably entwined with good management practices and that measurement of operational risk can only meaningfully be done if the effectiveness of organisational specific risk management and control processes is regularly assessed and included in the modelling. This contrasts sharply with a view that OpRisk modelling solely involves the investigation of statistical phenomena that characterises many of today's data driven AMA models. Moreover, because of this we find that management of OpRisk within the banks business processes are too often detached from the models developed to quantify regulatory capital.

Leaning on more than half a century of experience with OpRisk in safety critical industries we find that modelling for the analysis of operational risk always should strive to support decisions by providing answers to the following questions:

- Is the risk high or low, is it acceptable?
- Which causal factors are most critical?
- What are the differences in risk considering alternative solutions?
- What risk reducing effect can be achieved considering implementation of alternative risk reducing measures?

It is considered fundamental in OpRisk management to be able to create quantitative risk models that constitute a sound basis for fruitful discussions on organisation specific risk. Thus, detailed causal modelling at business process level is required for the model to absorb organisation specific input, highlight the criticality of causal factors (risk drivers), identify the potential lack of controls/barriers, and quantify the overall risk level in terms of expected and unexpected losses. Such a model obviously creates more value than a risk model based solely on actuarial techniques. Various modelling tools have been developed and applied for this purpose in the safety critical industries, the most famous ones being Fault Trees (Haasl 1965) and Event Trees (Nielsen 1971). Even though these tools have dominated quantitative risk assessment modelling for several decades we acknowledge various serious shortcomings in their use: common causes, time dynamics, subjective probability and large number of dependent parameters and multi-state variables rather than binary-state variables. This makes them unsuitable for detailed cause-consequence modelling of large and complex phenomena. Thus, the last decade an increasing number of professionals in the safety critical industry have adopted Bayesian Networks (BNs) to handle such modelling challenges [Ale *et al* 2007, Neil *et al* 2003; Fenton *et al* 2004; Langseth 2002; Røed *et al* 2007]. Moreover, due to significantly increased output value, numerous researchers have suggested causal modelling by means of BN rather than the traditional actuarial techniques [Alexander 2003, Adusei-Poku 2005, Cowell *et al* 2007, Neil *et al* 2005, Mittnik and Starobinskaya 2007].

It is argued that a major benefit of BNs is that they provide a structured and scientifically sound way of combining statistical analysis with other sources of empirical knowledge (e.g. knowledge on business processes and/or near misses) using Bayesian methods (Neil *et al*, 2005). The ability of BNs to account for causal dependencies enables risk analysts to link the operational conditions of the bank, including control environment, directly to the probability of losses occurring, as well as the severity of the losses. In other words quantitatively assess the effect of risk management decisions on risk exposure.

However, research on the application of BN models for operational risk has so far failed to address three fundamental modelling challenges. These are:

- Illustrate applicability to a complex banking process. Examples so far have been simplified and applied only to fragments of a banking processes thus not illustrating wide applicability for the full complexity of operational risk in banks
- Account for time dynamics in operational risk loss events. Operational losses evolve through series of sequential events in time (e.g. sequential checks performed by a set of controls that may, or may not be functional) which has not been captured by any BN model so far
- Implement continuous variables to assess loss severities. Use of BN software tools (e.g. (Hugin 2008) or (Netica 2008)) using static discretization (uniform discretization in fixed intervals) for continuous variables limit the application of continuous variables and thus also the possibilities to model loss severities accurately

If any large scale application of BN methodology is to be achievable in the context of operational risk these three challenges have to be overcome. Operational risk practitioners are of the view that BN models are suitable as a tool for risk management, but not as a tool for establishing economic capital for regulatory purposes (Adusei-Poku 2005). We are assuming that the reluctance to use BNs for quantification of regulatory economic capital is the result of the inability to effectively and accurately handle continuous variables. However, if regulatory capital is to reflect the individual risk exposure (as stated by Basel II as a qualifying criteria for the AMA) it follows that this measure has to change in accordance with efforts made to influence risk i.e. risk management efforts and quantification of regulatory capital must be linked. Thus we see no reason not to use BN for establishing a regulatory capital charge once the above challenges have been properly addressed.

In this paper we suggest solutions to these challenges to the application of BNs in operational risk modelling, using new state-of-the-art algorithms. We aim to use the financial trading process as an example to illustrate how BNs can be used and the background to financial trading is given in Section 2. Section 3 introduces BNs, how they are defined and the algorithms needed to solve them, including approaches to deal with dynamic behaviour and hybrid mixtures of variables. In Section 4 we specify a generalised Hybrid Dynamic Bayesian Network (HDBN) model to account for time dependences in the chain of events affecting operational loss. This model has three layers: a loss event model, a loss severity model and a loss aggregation model. Section 5 uses the generalised HDBN tailored for the trading process and uses causal loss factors derived from actual case data, coupled with fictitious probability and loss severity estimates, for two distinct scenarios, to produce an aggregated loss distribution for trading losses and accompanying Value at Risk (VaR) estimates. Finally, in Section 6 we offer some conclusions.

2. Rogue Trading

The background information for the modelling example presented in this paper is gathered from case studies of severe, trading loss events within the “Trading and Sales Business Line” (BIS, 2006), more colloquially known as “rogue trading”. In particular, we have consulted reports on the Barings Bank collapse (Bank of England 1995, IMD International 2002, Burke 2004), the Daiwa bank losses incurred by trader Toshihide Iguchi (Sungard Bancware Erisk 2001 with references), the Allied Irish Bank loss (Wachtell *et al* 2002) and the more recent Société Générale scandal (Progress report of the Special Committee of the Board of Directors of Société Générale, 2008). We find that the wisdom of hindsight provided by the investigations of these events give ample knowledge to construct a model of a generalised trading process in a bank, including the most important controls and control failures. To account for learning within the financial industry we have also consulted best practice documents like the Model Code – The international Code of Conduct and Practice for the Financial Markets (ACI – The Financial Markets Association, 2000) and Management of Operational Risk in Foreign Exchange (The Foreign Exchange Committee, 2004). The trading process and the major losses incurred as a result of this activity suits our purpose of illustrating application of BN models to complex banking

processes. We have however limited the scope of the model only to address the risk of unauthorised proprietary trading occurring in the trading process.

2.1 The Trading Process

The process of trading can, more or less independent of the product traded, be exhaustively described by the following steps; Trade request – Conduct Trade – Registration of Trade – Reconciliation check, Settlement and Netting – Entering of Trade on Trading Books. Any bank involved in trading also conducts continuous monitoring of results and positions, i.e. monitoring of market risk exposure (every 1, 10 or 30 days depending on the product and the bank). Obviously there are individual differences between banks, the products they trade and systems employed to carry out the trading but the basic general structure of the process is essentially the same across banks. Receipt of trade request, conducting trade and registering the trade in the trading system are done by the front office, i.e. the traders. The back office performs the confirmation, reconciliation and settlement check which involve checking counterparty information, authorisations, as well as check that the trade is within the terms of counterparty contract (for non-proprietary trading). The checks performed by the back office have the potential to uncover limit breaches by a trader but are not designed specifically to serve this purpose. Positions and results monitoring is done by the middle office which contains the risk control functions. Here the total market risk exposure is monitored and controlled against the risk limits set by the bank.

One of the major risks identified within the trading process is that of unauthorised trading. Unauthorised trading is essentially trading outside limits specified by the bank and the result is, simply put, an overexposure to market risk. However, trading outside the limits may incur catastrophic losses and involves hiding losses through fictive trades, fabricating trades and manipulating market positions monitoring, as well as falsification of documents, and in other ways influencing the control structure (see e.g. Wachtell *et al* 2002). In other words these are very complex events; however the complexity of the events also provides several opportunities to discover the transgression assuming the bank has sufficient and functional controls in place.

2.2 The Controls Process

The potential for high losses in trading activities, illustrated by events like Barings, and more recently confirmed by the Société Générale £3.7 bn loss, has resulted in banks implementing a complex control structure to manage their trading. Controls are implemented to prevent the traders (and/or other employees) to conduct trades that violate established regulations, limits and contracts. Should errors occur, there are also controls designed to uncover any irregularities and thus avoid or limit the potential loss. Controls can be viewed as direct or indirect; direct controls being specific checks performed to ensure the trade is conducted in a correct manner (e.g. reconciliation check), and indirect controls being operational elements having an indirect impact on the probability of incorrect trades being made and the performance of the direct controls (e.g. organisational culture, risk appetite, incentives structure). The controls perform checks at discrete times in a sequential order following the trade progression from beginning to end in the process.

Firstly there are controls in the front office intended to monitor and restrict the trading activity so that it is kept within the limits set by the bank. An example of such a control is instrumented checks in the trading IT system, monitoring the trades being entered and providing warning should specified constraints be violated. Another example is a requirement for traders to take at least two weeks vacation and during that time hand over their trading portfolio to another trader to enable discovery of unauthorised trading (if any). Several other controls are implemented on several levels and some best practices (not only regarding controls but also routines) are provided by e.g. “The international code of conduct and practice for the financial market” (ACI, the Financial Market Association, 2000). Some of the controls mentioned are direct results of previous unauthorised trading events (e.g. the vacation rule).

Secondly, once the trade has been processed and entered in the trading system by front office personnel, a number of consecutive checks are performed by the back office. Initial checks performed in the back office are done per trade and include checks such as reconciliation check/settlement check

and netting (see e.g. The Foreign Exchange Committee, 2004). In addition the middle office performs continuous positions and results monitoring including independent price checks and calculation of market VaR. This is done periodically varying from daily to every ten days to every month depending on the product and the bank. If unauthorised positions are acquired by a trader these will show up in the market VaR check providing it is not successfully hidden by the trader, through e.g. manipulation of price information. Manipulation of price information to change the value of positions may be uncovered through an independent price check conducted in the middle office.

Finally there is also the control of periodical or random audits of the trading operation. The audit is directed at the entire trading operation including processes and procedures, segregation of duties, etc. and also includes checking a sample of trades to check that there are not false, unauthorised or otherwise illegitimate trades present on the banks trading books.

Based on the reviewed case studies we have arrived at the following controls as being the most important for preventing and uncovering unauthorised trading here presented in sequential order:

- Front office control environment (the control environment affects the probability of unauthorised trading)
- Back office reconciliation checks (performed per trade)
- Market positions and results monitoring, VaR calculation (periodical)
- Audit checks (periodical but not as often as the market checks)

We are aware that these controls are not exhaustive with regard to the trading process; we have for example chosen to exclude the specific control of instrumented checks in the front office, and only include the reconciliation check in the back office. This is a conscious limitation of the model in a necessary trade off between complexity and reader friendliness. However, the modelling examples in section 4 illustrate that any number of controls can easily be included in the model. As we are focussing on proprietary trading no client is included in the control structure. For non-proprietary trades we could view the client as a potential control against unauthorised trades as it is likely that the client could discover unauthorised trades and make this known to the bank.

2.3 Loss Severity

The loss severity of unauthorised trades, being either the result of deliberate acts or mishaps, is a complex phenomenon. Dependent on the time of discovery and market movements a bank can suffer significant losses. The major unauthorised trading events are the result of simple unhedged directional bets, very vulnerable to market movements. It is a combination of these directional bets, volatile markets and commitment to large positions that give rise to the tail events of the loss severity distribution. The point is that the loss severity is the result of a complex set of events in itself. We shall in this paper only demonstrate that BNs can handle continuous variables to an extent necessary to model complex loss severity distributions; further research is needed to establish a methodology for loss severity modelling using BNs to ensure that the distributions reflect the characteristics of operational risk.

3. Bayesian Networks

A BN is a directed acyclic graph, such as the one shown in **Figure 1**, whose nodes represent the variables (with each variable there is associated uncertainty with the state of the variable) of interest and whose edges are the causal or influential links between the variables. Associated with each node is a node probability table (NPT), a statistical distribution or parameterised function. In the case of an NPT the relationship between nodes is governed by a set of conditional probability values that express the relationship between the node and its parents together with any uncertainty that is present in that relationship.

BNs enable reasoning under uncertainty and combine the advantages of an intuitive visual representation with a sound mathematical basis in Bayesian probability. With BNs, it is possible to

articulate dependencies between different variables and to propagate consistently the impact of evidence (observations) on the probabilities of uncertain outcomes.

The underlying theory of BNs combines Bayesian probability theory and uses conditional independence to represent dependencies between variables (Pearl 1986), (Speigelhalter and Cowell 1992). To date BNs have proven useful in many areas of application such as medical expert systems, diagnosis of failures, pattern matching, speech recognition and, more relevantly for the operational risk community, risk assessment of complex systems in high stakes environments, [Neil *et al* 2001, Ale *et al* 2007, Neil *et al* 2003; Fenton *et al* 2004; Langseth 2002; Røed *et al* 2007], including financial institutions [Alexander 2003; Neil *et al* 2005; Adusei- Poku 2005; Cowell *et al* 2007].

BNs were originally developed to model uncertainties between discrete variables with fixed labels. Typically, continuous variables are represented in a BN by a fixed set of discrete intervals and approximated by a piecewise constant function (i.e. a histogram) and this discretisation is then fixed throughout the inference process.

3.1 Definition and Algorithm

Formally, a BN is a probabilistic graphical model that represents a set of **discrete** valued variables and their probabilistic dependencies. BNs are directed acyclic graphs (DAGs), whose nodes represent variables in a probability distribution, and whose arcs encode conditional dependencies between the variables. Each variable X_i with parents $pa(X_i)$, has an associated probability table (also called potential, conditional probability table or node probability table (NPT)), $p(X_i | pa(X_i))$. A BN over a collection of variables $\mathbf{X} = \{X_1, X_2, \dots, X_n\}$, has a joint probability distribution $P(\mathbf{X})$ which is the product of all conditional probabilities, as specified by the given BN:

$$p(\mathbf{X}) = \prod_i p(X_i | pa(X_i)).$$

A *potential* is a real-valued table over domain of finite variables. Here $\phi_{\mathbf{X}}$ is used to denote the potential on the domain of a set of variables \mathbf{X} . There are two basic operations on potentials:

1. Combination: If ϕ is a potential over \mathbf{X} , ψ is a potential over \mathbf{Y} , then $\phi \otimes \psi$ is a potential for $\mathbf{X} \cup \mathbf{Y}$, where \otimes denotes combination.
2. Marginalization: If X_1, X_2 are marginalized out of potential ϕ over set of variables $\mathbf{X} = \{X_1, X_2, \dots, X_n\}$, then: $p(X_3, \dots, X_n) = \sum_{X_1, X_2} p(\mathbf{X})$.

Once a BN is designed, it can be executed using an appropriate propagation algorithm, such as the Junction Tree algorithm (Jensen 1996). This involves calculating the joint probability and all marginal probabilities for the model from the BN's conditional probability structure in a computationally efficient manner. This is achieved by automatically deriving from the BN an intermediate graph theoretic representation of the BN, called the Junction Tree (JT). The JT allows localised, modular computations to be executed using a message-passing algorithm. This is, in essence, an elaborate form of use of Bayes' theorem. For full details see (Jensen 1996), (Lauritzen and Speigelhalter 1988), (Pearl 1986) or (Speigelhalter and Cowell 1992).

Given a BN over a set of variables $\mathbf{X} = \{X_1, \dots, X_n\}$, we can define the JT, which contains clusters and separators. A *cluster* is a set which contains one or more variables and is represented by a single node. The *separator* is the intersection of two neighbouring clusters. A junction tree over \mathbf{X} is a tree of clusters of variables from \mathbf{X} , such that for each pair of nodes, \mathbf{V} , \mathbf{W} , all nodes on the path between \mathbf{V} and \mathbf{W} contain the intersection $\mathbf{V} \cap \mathbf{W}$.

The junction tree associated with the original BN is constructed in the following way:

1. Form a family of nodes such that for each variable $X_i \in \mathbf{X}$, there is at least one node V such that $\{X_i\} \cup pa(X_i) \in V$.
2. Organize the nodes as a tree with each edge labelled with *separators*. Each node V and separator S of the cluster tree is associated with a factor or potential ϕ_V , (ϕ_S respectively) over its variable set.
3. Given all nodes V and separators S assign a table of ones to ϕ_V and ϕ_S
4. For each variable $X_i \in \mathbf{X}$ choose exactly one node V containing $\{X_i\} \cup pa(X_i) \in V$ and combine ϕ_V using $p(X_i | pa(X_i))$

Clearly the product of all the potentials in the cluster tree is the product of all conditional probability

tables in the BN, so the potential based representation of a BN is $p(\mathbf{X}) = \prod_{k=1}^m \phi_k \equiv \phi_{V_1} \otimes \dots \otimes \phi_{V_m}$.

Propagation involves designating a root of the junction tree first, and then, based on a message-passing scheme, every leaf cluster sends a message towards the root. Once the root has received all the messages it then propagates messages outward to the leafs.

Let V and W be neighbours in a junction tree. Let $S = V \cap W$ be their separator, and ϕ_V , ϕ_W and ϕ_S be their potentials. A message pass from V to W occurs in two major steps, called projection and absorption:

- *Projection*: $\sum_{V \cap W} \phi_V$ computes the projection of the potential of V onto the separator S , i.e., the marginalisation of the potential ϕ_V with respect to S , storing the message in the separator's potential: $\phi_S = \sum_{V \cap W} \phi_V$.
- *Absorption*: The message to cluster W is given by the marginalisation of the potential ϕ_V onto the domain $S = V \cap W$. Update the potential of W with the message from V , by storing the message: $\phi_W^* = \left\{ \phi_W, \sum_{(V \cap W)} \phi_V \right\}$. We can think of absorptions as messages passed between the nodes in the tree. That is, a node V sends a message to its neighbour W when W absorbs from V . The message itself is the projection of the potential of V onto the separator S : $\phi_S = \sum_{V \cap W} \phi_V$.

After all nodes in the JT have sent and received messages from all neighbours the JT is said to be globally consistent and marginal distributions of interest can be reliably obtained.

The JT algorithm is entirely automatic and, in a tool like AgenaRisk (AgenaRisk 2008) or Hugin, is hidden from the domain expert. When the BN is executed the effects of data entered into one or more nodes can be propagated throughout the BN, in any direction, and the marginal distributions of all nodes updated. This makes it ideal for “what if?” and scenario analysis.

Clearly, the key to the successful design of a BN model is the meaningful decomposition of a problem domain into a set of causal or conditional propositions about the domain. Rather than ask an expert for the full joint probability distribution of all the variables of interest, which is obviously a very difficult task, we can apply a “divide and conquer” approach and ask for partial specifications of the model that are themselves meaningful in the experts’ domain. In our case, for operational risk the structure is an obvious artefact derivable from the operational process, as will become evident in later discussion. Next, we need to model the conditional probability tables for each variable (node): this can either be done using historical data (including, for example, using standard Bayesian parameter learning approaches or Monte Carlo simulations), or by simply asking the expert to provide a series of

subjective estimates. The experts consulted have to support credibility in the assessments made to motivate that the estimates are based on experience and knowledge rather than blind guesswork.

3.2 Hybrid Bayesian Networks

An extended notion of a BN is a hybrid BN (HBN) which contains both discrete and continuous variables and an inference algorithm with special operations to deal with continuous deterministic and statistical functions. In HBNs, local exact computations can be performed only under the assumption of Conditional Gaussian (CG) distributions (Lauritzen and Jensen, 2001). The advantages and drawbacks of using Conditional Gaussian distributions are well known. They are useful to model mixtures of Gaussian variables conditioned on discrete variables and weighted combinations of CG parents but they are much too inflexible to support general-purpose inference over hybrid models containing mixtures of discrete labelled, integer and continuous types and non-Gaussian distributions. Most real applications demand non-standard high dimensional statistical models with intermixed continuous and discrete variables, where exact inference becomes computationally intractable.

The present generation of BN software tools such as, [Hugin, 2008, Netica, 2008], attempt to model continuous nodes by numerical approximations using static discretization. Although discretization allows approximate inference in a hybrid BN without limitations on relationships among continuous and discrete variables, current software implementations require users to define a uniform discretization of the states of any numeric node (whether it is continuous or discrete) as a sequence of pre-defined intervals, which remain *static* throughout all subsequent stages of Bayesian inference regardless of any new conditioning evidence. The more intervals you define, the more accuracy you can achieve, but at a heavy cost of computational complexity. This is made worse by the fact that you do not necessarily know in advance where the posterior marginal distribution will lie on the continuum for all nodes and which ranges require the finer intervals. It follows that where a model contains numerical nodes having a potentially large range, results are necessarily only crude approximations.

We can embed continuous and discrete statistical distributions within the HBN model, as NPTs, and generate values for these NPTs by approximation methods, including Monte Carlo simulation. Until very recently, BN tools were unable to handle non-Gaussian continuous variables, and so such variables had to be discretized manually, with inevitable loss of accuracy. However, a breakthrough dynamic discretization algorithm presented in (Neil *et al* 2007) has now been implemented in a software tool, (AgenaRisk 2007). This allows the approximate solution of classical Bayesian statistical problems, involving continuous variables, as well as hybrid problems involving both discrete and continuous variables. This algorithm overcomes most of the problems inherent in the static or uniform discretization approaches. In particular, problems related to computational inefficiency caused by supporting too many states to represent the domain, high level of inaccuracy in posterior estimates for continuous variables, problems in instantiating evidence in areas of the domain that are grossly under sampled that lead to inconsistency and error. Full details of the dynamic discretization algorithm are described in (Neil *et al* 2007) but here we provide an overview of the algorithm.

Let X be a continuous (or integer valued) random node in the BN. The range of X is denoted by Ω_x , and the probability density function (PDF) of X , with support Ω_x , is denoted by f_x . The idea of discretization is to approximate f_x as follows:

1. Partition Ω_x into a set of interval $\Psi_x = \{w_j\}$, and
2. Define a locally constant function \tilde{f}_x on the partitioning intervals.

Discretization operates in much the same way when X takes integer values but in this paper we will focus on the case where X is continuous. As in Kozlov and Koller (1997), we use an upper bound of the Kullback-Leibler (KL) metric between two density functions f and g :

$$D(f \parallel g) = \int_S f(x) \log \frac{f(x)}{g(x)} dx$$

as an estimate of the relative entropy error induced by the discretized function. Under the KL metric, the optimal value for the discretized function \tilde{f} is given by the mean of the function f in each of the intervals of the discretized domain. The main task reduces then to finding an optimal discretization set $\Psi_x = \{\omega_j\}$.

Our approach to dynamic discretization searches Ω_x for the most accurate specification of the high-density regions given the model and the evidence, calculating a sequence of discretization intervals in Ω_x iteratively. At each stage in the iterative process, a candidate discretization, Ψ_x , is tested to determine whether the relative entropy error of the resulting discretized probability density \tilde{f}_x is below a given threshold, defined according to the trade off between the acceptable degree of precision and computation time.

By dynamically discretizing the model we achieve more accuracy in the regions that matter and incur less storage space over static discretizations. Moreover, we can adjust the discretization any time in response to new evidence to achieve greater accuracy. In outline, dynamic discretization follows these steps:

1. Convert the BN to a junction tree (JT) and choose an initial discretization for all continuous variables.
2. Calculate the NPT of each node given the current discretization
3. Enter evidence and perform global propagation on the junction tree, using standard JT algorithms.
4. Query the BN to get posterior marginals for each node, compute the approximate relative entropy error, and check if it satisfies the convergence criteria.
5. If not, create a new discretization for the node by splitting those intervals with highest entropy error.
6. Repeat the process by recalculating the NPTs and propagating the BN, and then querying to get the marginals and then split intervals with highest entropy error.
7. Continue to iterate until the model converges to an acceptable level of accuracy.

3.3 Hybrid Dynamic Bayesian Networks

Dynamic Bayesian Networks (DBNs) allow us to model temporal dependencies of complex systems as a joined sequence of BNs, each representing a particular object or the entire process at a particular point in time (called a time-slice). Some variable $A = \{a_1, \dots, a_n\}$ may change state as some process, O , progresses through discrete time intervals $t = \{1, \dots, T\}$. At each discrete time interval, t additional information E is also received about the state of A . The additional information may be observations, specific measurements or tasks performed at time t affecting the state of A . The state of A at time t is determined based on the state in the previous time interval, i.e. $t - 1$. The state of A might then be determined as $p(A_t | A_{t-1}, E)$.

DBN models potentially overcome most of the problems inherent in state-space based models and in particular, it avoids the state-space explosion problem of the Markov chains based approaches. Hybrid DBNs are simply DBNs that contain mixtures of discrete and continuous variables together.

The current state-of-the-art in approximate inference techniques on HDBN models with arbitrary probability distributions is primarily based on stochastic algorithms and Markov Chain Monte Carlo methods, (Murphy 2002), (e.g., importance sampling, sequential Monte Carlo, and Rao-Blackwellised Particle Filtering). These methods rely on intensive sampling algorithms that require drawing tens of thousands of dependent samples from, usually, high dimensional probability distributions. This presents two main shortcomings: it makes simulation techniques computationally inefficient and it

requires specialized statistical knowledge. In particular, specialised knowledge is required to ensure convergence of the dependent samples to a reliable result, and, even more difficult, to identify and deal with the special structures within the hidden variables of the HDBN required to make sampling in high-dimensional spaces a feasible task, specifically in Rao-Blackwellised schemes.

4. The Generalised HDBN Model for Operational Risk

Here we describe a generalised HDBN model for operational risk. The generalised HDBN model for operational risk comprises three layers:

- Loss event model
- Loss Severity model
- Aggregated loss model

Each layer is represented by a different BN, DBN or HDBN, as appropriate with interface links between them comprising common parameters. Use of Hybrid Bayesian Networks have also been suggested by Mittink and Starobinskaya (2007) to model dependencies between operational risk classes.

We use the word “generalised” to reflect the property that the layers are general enough to cope with a wide variety of situations, processes and contexts but, in practice, it is necessary to instantiate them by identifying specific variables local and meaningful in the process under study.

4.1 The Loss Model

The first layer is the “loss event model” which models the potential loss events, E_t , and how these dynamically evolve over time as they are influenced by controls, C_t , embedded within the business process. This dynamic time-based evolution of an event given the controls, is modelled by $p(E_t | E_{t-1}, C_t)$ with time periods $t = 1, \dots, T$. The performance of each control, C_t , is modelled as a function of a set of operational failure modes, O_j . These failure modes are in turn influenced by a set of causal factors, F_i , which in isolation or combined initiate the operational failure. Dependence between operational failure modes for different controls is modelled through dependency factors D_k . The dependency factors D_k , are conditioned on the occurrence of some operational failure mode specific for a control, C_t , and they, in turn, influence the probability of occurrence of a failure mode in a secondary control, C_{t+s} , where $s = 1, \dots, T - t$ such that:

$$p(O_{C_{t+s}} | D_k) p(D_k | O_{C_t})$$

An operational failure mode may be modelled as a function of both causal factors, F_i , and dependency factors, D_k . Operational failures that have no dependency relationship with other failure modes are simply modelled as a function of causal factors. The business process is then represented by a sequence of discrete time-dependent events such that we have a Dynamic Bayesian Network as shown in **Figure 1** which is a graph representation of the full joint distribution:

$$p(\mathbf{E}, \mathbf{C}, \mathbf{O}, \mathbf{F}, \mathbf{D}) = \prod_{t=1}^T \prod_{j=1}^m \prod_{i=1}^n \prod_{k=1}^o \prod_{q=1}^r p(E_t | E_{t-1}, C_t) p(C_t | \mathbf{O}_{C_t}) p(O_j | \mathbf{F}_{O_j}, \mathbf{D}_{O_j}) p(D_k | O_q) p(F_i) p(C_0)$$

where $j \neq q$, and $\mathbf{E}, \mathbf{C}, \mathbf{O}, \mathbf{F}, \mathbf{D}$ are sets of loss events, controls, operational failure modes, causal factor variables and dependency factors such that:

$$\mathbf{E} = \{E_0, E_1, \dots, E_t\}, \mathbf{C} = \{C_0, C_1, \dots, C_t\}, \mathbf{O} = \{O_1, O_2, \dots, O_m\}, \mathbf{F} = \{F_1, F_2, \dots, F_n\}, \mathbf{D} = \{D_1, D_2, \dots, D_l\}$$

Furthermore, we assume that the operational failure modes, O_j , can influence any of the control variables within any time period, t , and are thus time independent, using index j instead of t . A similar justification applies for the conditional relationship between the causal factors and the operational failure modes. The use of index q for operational failure modes affecting dependency factors D_k denotes that an operational failure mode cannot have a dependency relationship with itself. By adding E_t and C_t nodes and O_j, F_i , and if applicable, D_k variables we can easily increase the number of controls, assuming a placement on the discrete time continuum can be established.

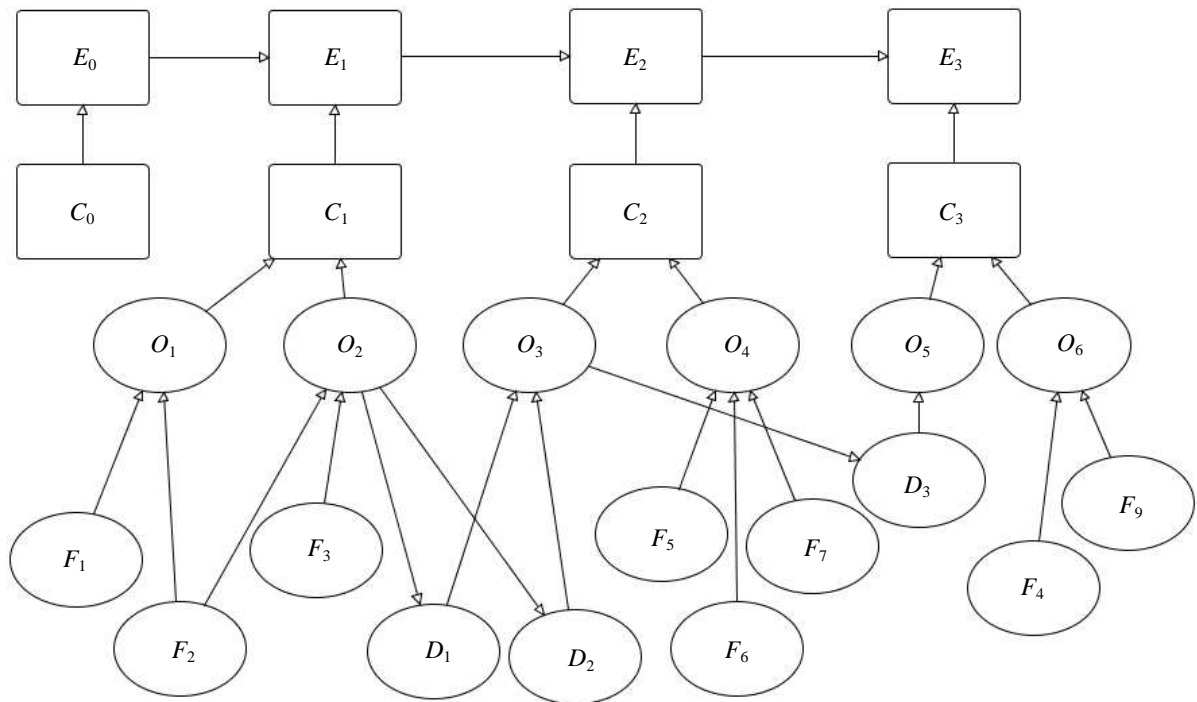


Figure 1 DBN of Loss event model

Each of the state transitions $p(E_t | E_{t-1}, C_t)$ is governed by a discrete node probability table that models the transition probability of the loss event from an undetected (unsafe) to a detected state, dependent on the control state. Should the control variable be operating correctly at time t , the loss event at E_{t-1} would transit to a correct operating state at E_t using the following logical conditional transition probabilities:

$$\begin{aligned}
p(E_t = fail \mid E_{t-1} = fail, C_t = fail) &= 1 \\
p(E_t = fail \mid E_{t-1} = fail, C_t = OK) &= 0 \\
p(E_t = fail \mid E_{t-1} = OK, C_t = fail) &= 0 \\
p(E_t = fail \mid E_{t-1} = OK, C_t = OK) &= 0 \\
p(E_t = OK \mid E_{t-1} = fail, C_t = fail) &= 0 \\
p(E_t = OK \mid E_{t-1} = fail, C_t = OK) &= 1 \\
p(E_t = OK \mid E_{t-1} = OK, C_t = fail) &= 1 \\
p(E_t = OK \mid E_{t-1} = OK, C_t = OK) &= 1
\end{aligned}$$

It is easy to see how we could generalise this to cope with more loss event states, including those showing a severity scale.

So, if the initiating loss event has not occurred, or has been detected by a previous control at a previous time, the question of whether this or later controls operate correctly become irrelevant. Thus the chance of detection and transition to a correct (safe) operating state is determined by the number of implemented controls as well as their individual performance (reliability).

Each control variable, C_t , is dependent on a collection of operational failure variables, \mathbf{O}_{C_t} , containing a subset of the operational failure modes deemed to prevent the control from performing its intended function. Thus the performance of each control is modelled by $p(C_t \mid \mathbf{O}_{C_t})$. Likewise the operational failure modes, O_j , are conditionally dependent on a subset of the causal factors, F_i , and dependency factors D_k relevant to the occurrence of the individual failure modes, and this is modelled by $p(O_j \mid \mathbf{F}_{O_j}, \mathbf{D}_{O_j})$.

For simplicity we have chosen to represent all control, operational failure, causal factor and dependency factor variables using Boolean nodes, for which we can model $p(C_t \mid \mathbf{O}_{C_t})$ and $p(O_j \mid \mathbf{F}_{O_j}, \mathbf{D}_{O_j})$ using Boolean logic to represent the type and form of failure. This modelling choice allows use of Boolean operators, such as AND, OR, XOR, NOT etc within a discrete BN model. For example an OR operator would be declared as:

$$p(C_1 = fail \mid O_1, O_2) = \begin{cases} 1 & \text{if } O_1 \cup O_2 = fail \\ 0 & \text{otherwise} \end{cases}$$

The NPT for node C_1 with parents O_1 and O_2 is then generated according to the above expression. Alternatively, we can manually declare an NPT to represent the form of dependencies that best matches the process, such as:

$$\begin{aligned}
p(C_1 = fail \mid O_1 = fail) &= 0.8 \\
p(C_1 = OK \mid O_1 = fail) &= 0.2 \\
p(C_1 = fail \mid O_1 = OK) &= 0 \\
p(C_1 = OK \mid O_1 = OK) &= 1
\end{aligned}$$

This is equivalent to saying we would expect the failure of control, C_1 , to occur with chance 80% when operational failure mode, O_1 , has occurred. This contrasts with the use of the Boolean OR operator which assumes a 100% chance of the control failing.

The presented methodology is by no means restricted to the use of Boolean variables. On the contrary one of the benefits of using BNs is that the modeller may declare customized variables and state spaces to fit the problem domain. Choice of variable types and state space is ultimately left to the modeller in the process of developing a model best suited for the process being analysed.

Armed with this model specification we can calculate the marginal probability of occurrence for each loss event at any discrete time step in the DBN, thus:

$$p(\mathbf{E}) = \sum_{\mathbf{C}, \mathbf{O}, \mathbf{F}, \mathbf{D}} \prod_{t=1}^T \prod_{j=1}^m \prod_{i=1}^n \prod_{k=1}^o \prod_{q=1}^r p(E_t | E_{t-1}, C_t) p(C_t | \mathbf{O}_{C_t}) p(O_j | \mathbf{F}_{O_j}, \mathbf{D}_{O_j}) p(D_k | O_q) p(F_i) p(C_0)$$

Inference here can be carried out using the JT algorithm described in section 3.1, and the model will typically execute within a few seconds with dozens of nodes.

The advantages of this model over a Monte Carlo simulation model are:

- Unlike in simulation, all computations are exact so it is eminently suitable for calculating ultra low probability events.
- Can deal with multiple, interacting, common causes.
- In addition to strong mathematical foundations the graph representation of the BN model maps neatly onto the underlying process and clearly shows how losses propagate.
- We can generate a set of loss event probabilities for different severities of events at different stages of propagation through the process and use this to calculate our severity and loss models.

Note that some clarification of what we mean by “causes” and how we identify them is needed here since there is much understandable dispute and confusion of how these terms might be productively applied to operational risk and other problems. We recognise that there are an infinite number of actual and possible interactions between unique events and that it is therefore impossible to represent these “atomic level” interactions in any model. Instead we seek to model broad classes of events that have characteristics, behaviour and outcomes that are similar enough to be considered homogenous and which can reasonably be expected to interact in a similar way. Given this any distinction between classes of causal events is a practical decision as much informed by process understanding, documentation, data collection standards and budget as any refined philosophical position. For instance in a practical situation we might class operational IT failure as a single causal class covering “security lapses, systems failures and performance degradation” simply because the consequential effect on a control that depends on the IT infrastructure would be similar in all cases.

4.2 The Loss Severity Model

Here we use the probabilities generated by the loss event model to predict the total losses by severity class, given the severity distribution and a measure of volume to scale the losses. We assume that it is a subset of the states in $p(\mathbf{E})$ we are interested in here, specifically those that incur a loss or are “unsafe”: $p(E_t = e_t^L)$, and of course there may be more than one loss event for each discrete time step, t , hence the superscript, L . For brevity we let $p(E_t = e_t^L) = p_{e_t^L}$.

We can model the total losses using a conditional dependency model, again represented as a DBN, where the total number of loss events within each time step, $N_{e_t^L}$, is binomially distributed, given a volume measure, e.g. the total number of trades conducted weekly, monthly or annually, V :

$$N_{e_t^L} \sim \text{Bin}(W_{e_t^L}, p_{e_t^L})$$

, where $W_{e_t^L} = V - \sum_{i=0}^{t-1} N_{e_i^L}$ is the volume of transactions minus the cumulative number of loss events predicted for all previous time steps, i.e. $t-1$ in the DBN. Graphically this model is shown in **Figure 2**.

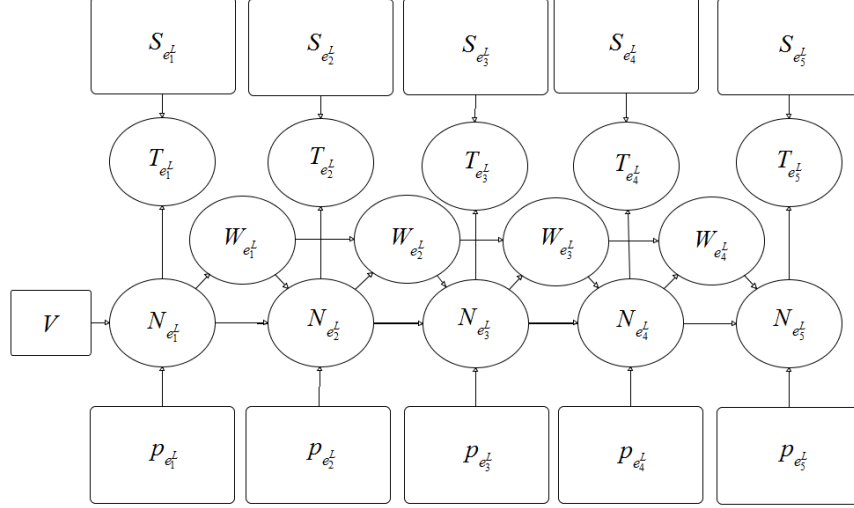


Figure 2 Loss severity model

Associated with each e_t^L is a severity distribution $f(S_{e_t^L})$, representing the distribution of financial losses, rework costs or other penalties. We assume that the severity distribution can take any form and that the conditionally deterministic variable, $T_{e_t^L}$, representing the total loss for a single event is:

$$f(T_{e_t^L} | S_{e_t^L}, N_{e_t^L}) = S_{e_t^L} \times N_{e_t^L}$$

where $S_{e_t^L} \geq 0$, i.e. we only consider losses and not gains.

Furthermore we assume that the severity of events that are discovered later in time are more serious since they typically lead to larger expected losses or have larger uncertainties attached. Earlier, “near miss” losses will tend to be of a more predictable nature since they are likely to be encountered more often. Also early discovery by a control will presumably provide greater opportunity to limit the potential loss of an unauthorised trade. Therefore, one might expect to constrain the model such that: $E(S_{e_1^L}) < E(S_{e_2^L}) < \dots < E(S_{e_T^L})$.

The marginal distribution for the total losses for each loss event is therefore:

$$f(T_{e_t^L}) = \sum_{S_{e_t^L}, N_{e_t^L}, W_{e_t^L}} f(T_{e_t^L} | S_{e_t^L}, N_{e_t^L}) f(N_{e_t^L} | W_{e_t^L}, p_{e_t^L}) f(S_{e_t^L})$$

4.3 The Aggregated Loss Model

Now we have a set of total loss variables, $T_{e_t^L}$, for each loss event, e_t^L and wish to calculate the total aggregated losses, A , as the sum of total losses associated with each event in each time period. This aggregated sum is simply the deterministic function, $A = \sum_{t=0}^T T_{e_t^L}$. This can either be solved by

convolution or sampling, however, in our BN framework we instead use the dynamic discretization algorithm described in Section 3.2.

The total aggregated loss distribution is obtained by marginalising:

$$f(A) = \sum_{\mathbf{T}_{e_t}} f(A | \mathbf{T}_{e_t}) f(\mathbf{T}_{e_t})$$

, where \mathbf{T}_{e_t} is the collection of total losses for each loss event.

From $f(A)$ we can now calculate the Value at Risk (VaR); under the Basel regulations the VaR statistic is the 99.9th percentile of this marginal distribution: $f_{\alpha=99.9\%}(A)$.

5. An HDBN for the Rogue Trading Process

5.1 The Loss Event Model

In the context of the example model in this paper unauthorised trading is defined as a trader (intentionally or unintentionally) exceeding his authorised trading limit. The main controls in place to prevent such an event, and detect it should it occur are:

- Front office control environment
- Reconciliation check carried out by the back office
- Position and results monitoring (market risk monitoring) carried out by the middle office
- Audit (periodically or random)

Any trade conducted is assumed to be in one of a finite set of mutually exclusive states corresponding to potential loss events, E_i . Each of the implemented controls performs a check that may change or confirm (i.e. the state remains as before) the state of the trade given its true state. We assign the following definition to the loss event variables:

E_0 : State of trade when entered in the trading system by the trader, C_0

E_1 : Loss event discovered during reconciliation check, C_1

E_2 : Loss event discovered during market risk VaR check, C_2

E_3 : Loss event discovered by audit, C_3 , or which has escaped the control system

The states assigned to each variable are:

$E_0 = \{ \text{Authorised, Accidental, Illegal Fame, Illegal fraud} \}$

$E_1 = \{ \text{Discovered, OK, Accidental, Illegal Fame, Illegal fraud} \}$

$E_2 = \{ \text{Discovered, OK, Accidental, Illegal Fame, Illegal fraud} \}$

$E_3 = \{ \text{Discovered, OK, Accidental, Illegal Fame, Illegal fraud} \}$

The meaning of each of the states are as follows:

- *Authorised/Ok trade.* An authorised trade permitted and approved by management beforehand.
- *Accidental unauthorised trade.* Trades that are accidental due to mistakes by the trader, i.e. trades that were not intended to be unauthorised and of course not intended to incur losses for either the bank nor its clients
- *Illegal Fame trade.* Unauthorised trades with the intent to further the trader's career, status within the bank and/or size of bonuses. For this category of trades the trader intends to make money for the bank and subsequently also provide benefits for himself. The important aspect is that the trade is not intended to directly benefit the trader but indirectly through providing success for the bank
- *Illegal Fraud trade.* Unauthorised trades with the sole intent of benefiting the perpetrator, if the bank or anyone else suffers losses as a result of the trade is irrelevant as long as the perpetrator makes money off the trade
- *Discovered.* A trade of any unauthorised category is revealed by a control and is thus a discovered unauthorised trade.

We could use only two states for the event variable, E_0 ; *Normal* and *Unauthorised* but that will not properly reflect the possible severity of the event. Likewise we distinguish between illegal fame and fraud states because we are assuming different severities, and wish to assign different severity distributions to them. Alternatively we could design three individual models using the presented methodology, each addressing one of the unsafe trade states separately (i.e. accidental, fame and fraud trades). Such an approach could potentially allow a greater degree of tailoring the model to the operational failure modes and influencing factors relevant specifically to a particular loss event. However such an approach would also result in greater model complexity and size.

The trade has a true loss event state when entered in the trading system by the trader. At $t = 0$ the loss event state is unknown to us and as the trade progresses through the controls process and checks are performed the state of the trade may change to "discovered" if an unauthorised trade is uncovered, otherwise it will remain in the "OK" state as the trade is assumed authorised. There is also a possibility that an unauthorised trade will not be discovered and continue to be in an "unknown" unauthorised state and escape the control system and our model of the process. Such loss events typically have higher severity and their impact felt much later when subsequent frauds are discovered and forensic analysis throws up unauthorised past trades. We may also use escaped unauthorised trades to assess the severity of loss where severity is assumed dependent on the occurrence of previous unauthorised trades, this will be addressed in more detail in further research.

The described process can be modelled using our HDBN approach and the corresponding loss event model is presented in **Figure 3**.

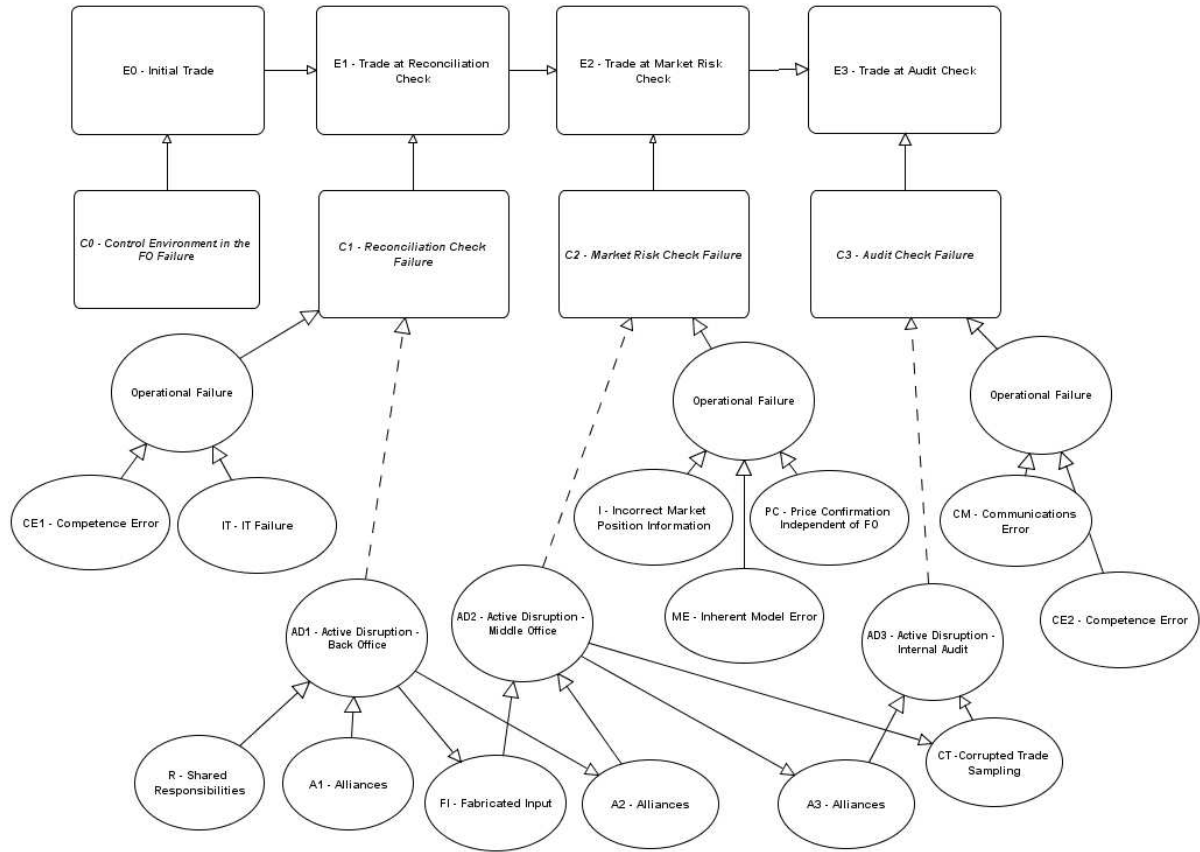


Figure 3 DBN Loss event model for rogue trading process

In **Figure 3** each control is influenced by a series of operational failure nodes, \mathbf{O}_{C_t} , where each operational failure influence the performance of one or more controls. Here the operational failures are particular to each control with the exception of the operational failure “Active Disruption” (AD) which is shared by all of the controls, except the front office. However since the controls normally are separated from each other, different actions, and alliances are needed by the trader to disable a complete set of controls. Thus we have included an AD node, O_{AD_t} for each control, C_t except C_0 . Also the probability of active disruption in one control is dependent on whether or not there has been active disruption in the preceding control, assuming a trader that has successfully disrupted one control is likely to attempt to disrupt the next. Hence given active disruption in the back office there is an increased probability of active disruption in the remaining controls. Such dependency is modelled using dependency factors, D_k . This is a useful feature and an important advantage of the BN approach in that common causes of failure that undermine the whole process can be accounted for. In the case of “Active disruption” the dependency factors, \mathbf{D}_{O_j} , influencing controls failure include failed segregation of duties (between front and back office), fabricated input data (in the positions monitoring), corrupted trade sampling (in the audit process) and alliances between staff, within and out with, the institution. The fact that this operational failure affects all back office, middle office and audit controls means a rogue trader could negate or reduce the effectiveness of all controls across the board. Also, modelling the active disruption of controls as suggested enables the analyst to account for the complexity in the deliberate acts to negate or reduce the efficiency of controls. In contrast with “Active Disruption” each of the other controls are modelled using causal factors, such as competence failure, IT failure and position/pricing mistakes.

To illustrate how our approach works we have assigned probability values to each of the causal factors as well as the front office control, for two scenarios, as shown in **Table 1**. The first scenario

represents status quo on a trading floor whilst the second represents the case where the control environment in the front office has failed and there is active disruption in the back office.

Table 1: Failure probabilities for causal factors in two scenarios

Variable (true = “failed”)	Probabilities for scenario 1	Probabilities for scenario 2
$p(C_0 = true)$	0.01	1.0
$p(F_{CE1} = true)$	0.05	0.005
$p(F_{IT} = true)$	0.01	0.001
$p(F_R = true)$	0.05	1.0
$p(F_{AI} = true)$	0.01	1.0
$p(F_I = true)$	0.01	0.001
$p(F_{PC} = true)$	0.05	0.005
$p(F_{ME} = true)$	0.01	0.001
$p(F_{CM} = true)$	0.01	0.001
$p(F_{CE2} = true)$	0.01	0.001

For the dependency factors resulting in active disruption in the middle office and the internal audit, which are also dependent on previous active disruption $D_{O_{ADi}} | O_{ADi}$, the assigned probabilities are:

$$p(D_{A2} = true | O_{AD2} = false) = 0.001$$

$$p(D_{A2} = true | O_{AD2} = true) = 0.7$$

$$p(D_{FI} = true | O_{AD2} = false) = 0.001$$

$$p(D_{FI} = true | O_{AD2} = true) = 0.8$$

$$p(D_{A3} = true | O_{AD3} = false) = 0.001$$

$$p(D_{A3} = true | O_{AD3} = true) = 0.6$$

$$p(D_{CT} = true | O_{AD3} = false) = 0.001$$

$$p(D_{CT} = true | O_{AD3} = true) = 0.6$$

Note that the probability of occurrence of the dependency factors is assumed to increase dramatically when there is active disruption in the preceding control.

For the initial event state $E_0 | C_0$, the NPT is shown in **Table 2**. Notice that, when front office controls are working the odds for illegal fame and fraud events are in the region of 1:1000 for illegal fame and 1:000000 for illegal fraud. This rises to 1:100 and 1:10000 when controls have failed.

Table 2: NPT probabilities for $p(E_0 | C_0)$

	$C_0 = false$	$C_0 = true$
$E_0 = Authorised$	0.9890011	0.96963763
$E_0 = Accidental$	0.00998991	0.009794319
$E_0 = Illegal Fame$	9.989911E-4	0.019588638
$E_0 = Illegal Fraud$	9.989911E-6	9.794319E-4

When we execute the DBN, inference is carried out on all variables and marginal distributions produced as shown in **Figure 4** for the first scenario.

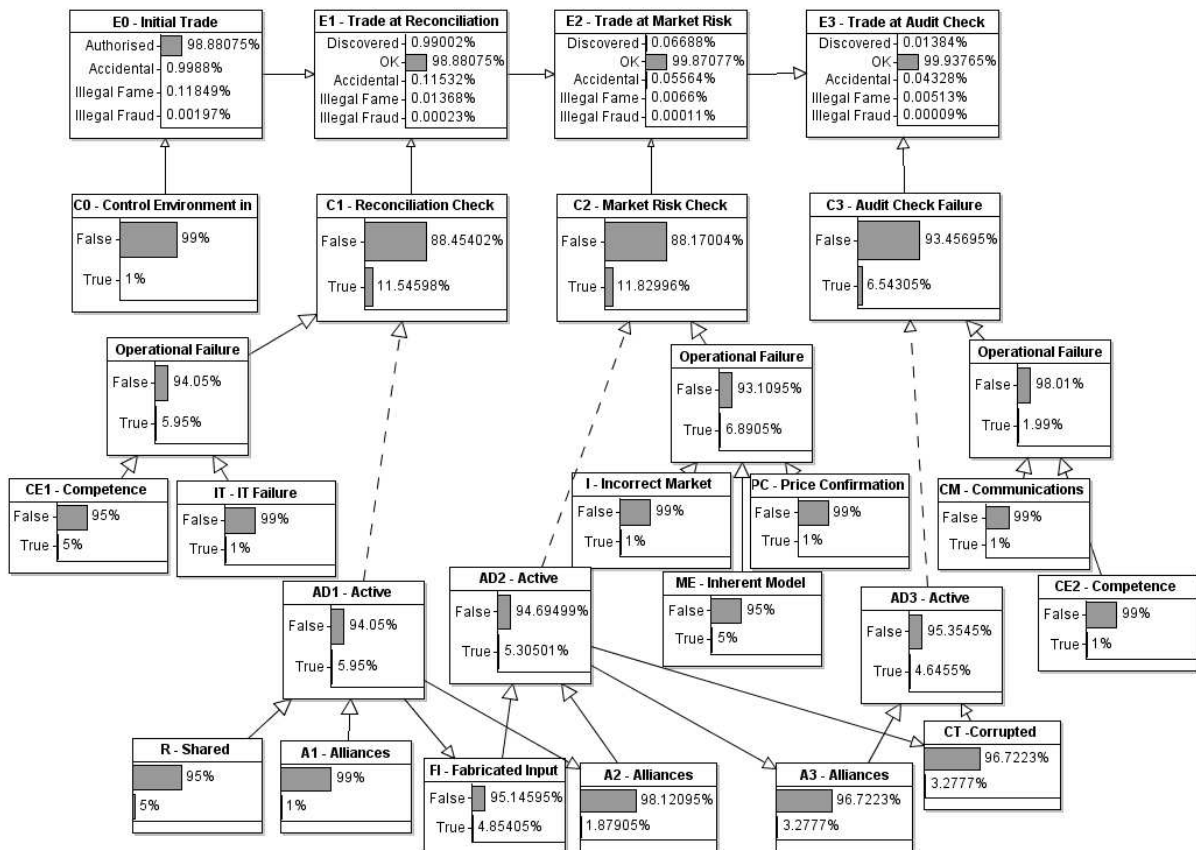


Figure 4 DBN Loss event model for rogue trading process with superimposed marginal probability distributions

From **Figure 4** we obtain the loss event probabilities for scenario 1, and these are listed in **Table 3**. These loss probabilities give the probability that a single trade will belong to that particular loss class. From these we can calculate that the “escape” probability for a loss, i.e. the probability the trade is unauthorised and evades all of the controls is 0.00049 in scenario one and 0.0220233 in scenario 2. The probability of a trade being unauthorised and discovered in scenario one is 0.0107074 and 0.008339 in scenario 2.

Table 3: Loss event probabilities computed for two scenarios

Loss event	Probabilities for scenario 1	Probabilities for scenario 2
$e_1^{discovered}$	0.0099002	0.0
$e_2^{discovered}$	0.0006688	0.0039578
$e_3^{discovered}$	0.0001384	0.0043812
$e_3^{accidental}$	0.0004328	0.0071043
$e_3^{illegal\ fame}$	0.0000513	0.0142086
$e_3^{illegal\ fraud}$	0.0000009	0.0007104

5.2 The Loss Severity Model

Here we assume a trade volume, V , of one million trades over a year and use this to predict the total losses, $T_{e_t^L}$, for each of the loss events. In our example the severity distribution for each loss event, $S_{e_t^L}$, is a left truncated Gaussian distribution with parameters as listed in **Table 4**. Note that the mean loss increases with the severity of loss as does the variance. The assigned mean can be interpreted as representing the expected loss from an unauthorised trade discovered at time t , (or a trade that escapes discovery) and the variance as a representation of uncertainty related to the actual size of the loss.

Table 4: Parameters for loss severity distributions (\$)

Loss severity	μ	σ^2
$S_{e_1^{discovered}}$	100	100,000
$S_{e_2^{discovered}}$	200	100,000
$S_{e_3^{discovered}}$	500	100,000
$S_{e_1^{accidental}}$	800	1,000,000
$S_{e_1^{illegal\ fame}}$	1000	1,000,000
$S_{e_1^{illegal\ fraud}}$	5000	1,000,000

For unauthorised fame trades the marginal total loss distribution, $T_{\epsilon_1^{illegal\ fame}}$, is shown in **Figure 5**. The mean total losses are \$65,549 and the 99th and 99.9th percentile values are \$189,920 and \$251,320 respectively.




Figure 5 Marginal total loss distribution for unauthorised fame trades under scenario one, with 99th and 99.9th percentiles shown

5.3 The Aggregated Loss Model

Finally, we can calculate the marginal aggregate loss distribution, $f(A)$, for each of our two scenarios, using the dynamic discretization algorithm previously discussed. The resulting mean and VaR summary statistics are shown in **Table 5**. Clearly scenario two is substantially worse than scenario one, which is what we would of course expect if we have active disruption in the back office and a failed control environment in the front office.

The aggregated marginal loss distributions for each scenario are shown in **Figure 6** and **Figure 7**. Note that each presents the characteristic long tail shape assumed of loss distributions of unexpected operational losses.

Table 5: Mean and VaR statistics for scenarios (\$)

Scenario	Mean	99 th percentile	99.9 th percentile
Scenario 1	3,798,400	9,973,100	14,073,000
Scenario 2	33,605,000	68,650,000	92,095,000




Figure 6 Marginal aggregate loss distribution for scenario one, with 99th and 99.9th percentiles superimposed




Figure 7 Marginal aggregate loss distribution for scenario two, with 99th and 99.9th percentiles superimpose

6. Concluding Remarks

This paper has described a new method for applying Hybrid Dynamic Bayesian Networks (HDBNs) to model the operational risk to financial institutions. We describe a methodology for modelling financial losses that result from intentional or accidental events that can be characterised by complex interactions of events and their ability to evade controls and ultimately lead to increasingly severe financial consequences. We have focused on modelling the causes and effects of loss events using a Dynamic Bayesian Network model incorporating interactions between failure modes and controls.

We have argued that Bayesian Networks are a natural choice and powerful method for modelling operational problems and have presented a generalised approach using HDBNs. This paper shows that the approach can successfully model dependencies between events and processes in complex environments evolving over time, and we have illustrated this using the financial trading process. We present an example model containing causal factors gleaned from a careful review of published rogue trading events and use fictitious data to produce loss distributions and VaR predictions for two scenarios. Furthermore we show how our approach overcomes the three major limitations arising from the applications of BN technology by other researchers: applicability, time dynamics and continuous variables.

We claim that BNs meet the requirements of the Basel Accord, (Basel 2006), for an Advanced Measurement Approach (AMA) in providing an individual measure of risk exposure using a combination of expert knowledge, data and business environment and internal control factors. In addition the model has the potential for reflecting developments in risk exposure as a result of degraded or improved controls. Adopting a BN based approach should therefore lead to better operational risk governance and provide a foundation for sound risk management with subsequent reduction of regulatory capital charge. The BN approach presented here strongly contrasts with a purely statistical approach based on historical loss data alone. We believe that traditional statistical analysis techniques will neither provide good predictions of future operational risk losses, nor provide a mechanism for controlling and monitoring such losses. This latter goal is obviously of utmost importance in practice.

This paper has also provided an introduction to BN technology and algorithms and its contribution to operational risk modelling. We show how a new state-of-the-art algorithm for HBNs, called dynamic discretization, enables more sophisticated modelling and analysis than can be supported using previous generation of BN algorithms. For example, such tools now support *learning* from loss data in a way that competes with techniques that have hitherto been recognised as state of the art, such as Monte Carlo Markov Chain (MCMC) (see Gelman *et al* 2004). The new approach also make it relatively easy to model more complex dependencies between losses and severities than those covered here including, time-based losses, time to respond, time to failure and complex mixture models.

In addition to applying the approach in practice, we anticipate carrying out further research work in the following areas:

- The current approach to modelling loss severity is not sufficiently explicit in two key areas. Firstly the model does not define the type and scope of losses, which in practice would include penalty costs, rework costs, legal and reputational losses. Secondly, in financial trading unauthorised losses can have an upside risk, whereas this model assumes a downside risk only i.e. the losses are unhedged.
- We have modelled the process as a discrete time process where trade volume is assumed to be a discrete measure. The applicability of Bayesian versions of continuous time processes, where rate of occurrence rather than probability of occurrence is a key parameter, deserves attention.
- Challenges in scaling up the presented model to model the complete array of risks faced by a bank in a complete methodological approach also needs attention. In order for the approach suggested here to be applicable as a framework for an AMA model more comprehensive models need to be constructed and validated.

Acknowledgements

The authors would like to thank David Marquez, Norman Fenton and Jørgen Osenbroch for their helpful comments and feedback. This paper is based in part on work undertaken on the DyFusion: “Towards a Novel Universal tool for Modelling and Reasoning under Uncertainty” project, EPSRC project code EP/E033954/1 and is published as a part of an OpRisk research project funded by the Norwegian banking industry in collaboration with the Norwegian Research council.

References

Agenarisk (2008) Bayesian Network and Simulation software. www.agenarisk.com

Adusei-Poku, K. (2005). Operational Risk Management – Implementing a Bayesian Network for Foreign Exchange and Money Market Settlement, Presented for the degree of Doctor of Philosophy at the Faculty of Economics and Business Administration of the University of Göttingen.

- Ale, B. J. M., Bellamy, L. J., van der Boom, R., Cooper, J., Cooke, R. M., Goossens, L. H. J., Hale, A. R., Kurowicka, D., Morales O. (2007). Further development of a Causal model for Air Transport Safety (CATS); Building the mathematical heart, In Proceedings of European Safety and Reliability Conference (ESREL) 2007.
- Alexander, C.(2003). Managing operational risks with Bayesian Networks, *Operational Risk, Regulation, Analysis and Management*, Prentice Hall, pp 285-295.
- Bank of England (1995). Extract from: Report of the Board of Banking Supervision Inquiry into the Circumstances of the Collapse of Barings, July 1995, <http://www.numa.com/ref/barings/bar00.htm>
- Basel Committee on Banking Supervision (2003). International Convergence of Capital Measurement and Capital Standards. <http://www.bis.org>.
- Burke, S. (2004). Currency Exchange Trading and Rogue Trader John Rusnak, Concept 2004, <http://www.publications.villanova.edu/Concept/2004.html>
- Cowell, R. J., Verral R. J. and Yoon Y. K. (2006). Modelling Operational Risk with Bayesian Networks, *The Journal of Risk and Insurance*, Vol. 74, No. 4, pp 795-827. .
- Fenton, N. E., Marsh, W., Neil, M., Cates, P., Forey, S. and Tailor M. (2004). Making Resource Decisions for Software Projects. In Proceedings of 26th International Conference on Software Engineering (ICSE 2004), (Edinburgh, United Kingdom, May 2004) IEEE Computer Society 2004, ISBN 0-7695-2163-0, 397-406.
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin D. B. (2004). *Bayesian Data Analysis* (2nd Edition), Chapman and Hall, pp. 209 – 302.
- Hugin (2008). www.hugin.com
- Haasl, D.F. (1965): "Advanced Concepts in Fault Tree Analysis", System Safety Symposium, The Boeing Company, Seattle, June 1965.
- IMD International (2002), The Barings Collapse (A): Breakdowns in Organizational Culture and Management, International Institute for Management (IMD), 2002
- IMD International (2002), The Barings Collapse (B): Failures in Control and Information use, International Institute for Management (IMD), 2002
- Jensen, F. V. (1996). *An Introduction to Bayesian Networks*. UCL Press, London, 1996.
- Kozlov, A. V. and Koller, D. (1997). Nonuniform dynamic discretization in hybrid networks, in D. Geiger and P.P. Shenoy (eds.), *Uncertainty in Artificial Intelligence*, 13:314-325.
- Lauritzen, S.L. (1996). *Graphical Models*, Oxford University Press.
- Lauritzen, S.L., and Jensen, F. (2001). Stable local computation with conditional Gaussian Distributions, *Statistics and Computing*, 11, 191–203.
- Lauritzen, S.L., and Spiegelhalter, D.J. (1988). Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems (with discussion), *Journal of the Royal Statistical Society Series B*, Vol. 50, No 2, pp.157-224.
- Lauritzen, S. L. and Spiegelhalter, D. J. (1988). Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems (with discussion)". *Journal of the Royal Statistical Society Series B*, Vol. 50, No 2, pp.157-224.

- Lauritzen, S.L., and Jensen, F. (2001). Stable local computation with conditional Gaussian Distributions, *Statistics and Computing*, 11, 191–203.
- Langseth H. (2002). Bayesian Networks with applications in Reliability Analysis, PhD Thesis, Department of Mathematical Science, Norwegian University of Science and Technology, 2002.
- Murphy, K (2002). Dynamic Bayesian Networks: Representation, Inference and Learning. PhD thesis, Dept. Computer Science, UC Berkeley, 2002.
- Mittnik S. and Starobinskaya I. (2007) Modelling Dependencies in Operational Risk with Hybrid Bayesian Networks, *Methodology and Computing in Applied Probability*, Springer Netherlands, 2007.
- Neil, M, Fenton, N. E., Forey, S. and Harris R. (2001). Using Bayesian Belief Networks to Predict the Reliability of Military Vehicles, *IEE Computing and Control Engineering*, 12(1), 2001, pp. 11-20.
- Neil, M., Malcolm B. and Shaw R. (2003). Modelling an Air Traffic Control Environment Using Bayesian Belief Networks. Presented at the 21st International System Safety Conference, August 4 - 8, 2003, in Ottawa, Ontario, Canada.
- Neil M., Fenton N. and Tailor M. (2005). Using Bayesian Networks to model Expected and Unexpected Operational Losses. *Risk Analysis Journal*, August 2005.
- Neil M., Tailor M., Marquez D. (2007). Inference in Bayesian Networks using Dynamic Discretization , *Statistics and Computing* 17:3 September 2007.
- Neil M., Tailor M. and Marquez, D. (2007). “Inference in Bayesian Networks using dynamic discretization”. *Journal of Statistics and Computing*, Springer Netherlands, Vol. 17, Number 3, September 2007.
- Nielsen, D.S. (1971): "Use of Cause-Consequence Charts in Practical Systems Analysis" Reliability and Fault Tree Analysis, SIAM, 1971.
- Netica (2008). www.norsys.com
- Pearl, J. (1986). Fusion, propagation, and structuring in belief networks, *Artificial Intelligence*, Vol. 29, 1986.
- Rawnsley, J. (1995). *Going for Broke: Nick Leeson and the Collapse of Barings Bank*. HarperCollins, 2002.
- Reason, J. (1997). *Managing the Risks of Organisational Accidents*. Ashgate Publishing Limited, 1997.
- Røed, W., Mosleh, A., Vinnem, J. E., Aven, T. (2007). On the Use of Hybrid Causal Logic Method in Offshore Risk Analysis. Submitted for possible publication.
- Sungard Bancware ERisk (2001). Daiwa Bank, An ERisk.com Case Study, <http://www.erisk.com/Learning/CaseStudies/DaiwaCaseStudy.pdf>
- Spiegelhalter, D. J. and Cowell, R. J. (1992). Learning in Probabilistic Expert Systems, *Bayesian Statistics*, 4, pp. 447-465. Oxford University Press, 1992.
- The Times newspaper “SocGen rogue trader 'ready to talk to police’”, 25 January 2008.
- Venkataraman, Subu S. (1997). Value at risk for a mixture of normal distributions: the use of Quasi-Bayesian estimation techniques. *Economic Perspectives*, Federal Reserve Bank of Chicago.

Wachtell, Lipton, Rosen, Katz & Promontory Financial Group (2002). Report to the Boards of Allied Irish Banks, p.l.c., Allfirst Financial Inc. and Allfirst Bank Concerning Currency Trading Losses March 12.