NOTE: This is an unpublished Technical Report Dated 2012. This work has since been extended and submitted for publication in 2015 (with additional authors)

Using soft evidence to model mutually exclusive causes in Bayesian networks

Norman Fenton, Martin Neil and David Lagnado

Abstract—Bayesian network (BNs) are especially well suited for reasoning about the impact of uncertain evidence in many domains, especially in legal arguments. Yet, standard BN modelling techniques and algorithms do not capture the correct intuitive reasoning in situations in which different possible causes of some effect are ***mutually exclusive***. Despite this being a very common scenario, it has never been adequately addressed in the literature and none of the 'standard' proposed solutions work properly. We demonstrate a novel and simple solution to the problem. It is achieved by introducing a special type of 'constraint' node, which enforces mutual exclusivity between causes, and the use of specific 'soft evidence' to preserve prior probabilities of causes as well as ensuring that impossible states do not occur. The solution can be implemented in standard BN tools and is applicable to a wide range of common reasoning problems.

Index Terms— Bayesian networks, mutually exclusive causes, legal arguments, uncertain reasoning.

1 Introduction

A Bayesian network (BN) is a graphical probabilistic model that is especially well-suited in decision-making scenarios that require us to consider multiple pieces of uncertain evidence involving causal relationships. A BN consists of a set of nodes (that represent uncertain variables) and directed edges between those nodes for which there is a causal or influential relationship. Every node has an associated *node probability table* (NPT); for any node without parents the NPT specifies the prior probabilities of each of the node states, while for any node with parents the NPT captures the prior probability of each node state conditioned on each combination of states of the parent nodes. A BN enables us to visualise the relationship between different hypotheses and pieces of evidence in a complex argument. In addition to its powerful visual appeal, a BN has an underlying calculus based on Bayes Theorem that determines the revised probability beliefs of all uncertain variables when any piece of new evidence is presented. This process is called ***evidence propagation***. There are widely available BN tools that implement standard propagation algorithms, and hence enable non-specialist users to easily build and run BN models. Propagation enables a BN to be used for both prognostic and diagnostic types of reasoning. In prognostic reasoning we enter evidence about causes to reason about effects (we also refer to this as 'forward inference') whereas in diagnostic reasoning we enter evidence about effects to reason about causes (we also refer to this as 'backward inference').

What we are interested in here is the special case, which is common in legal and other arguments, where different possible causes of some effect are necessarily ***mutually exclusive*** (meaning that at most one of the causes can be true for any instance of the problem). In Section 2 we introduce a set of necessary axioms that need to be satisfied by any BN that attempts to model mutually exlusive causes and show that standard BN modelling techniques and algorithms do not satisfy these axioms. Although a number of studies have touched on the problem [14],[19],[26],[29],[33] the problem has never been stated explicitly or adequately resolved. In Section 3 we explain why none of the proposed 'solutions' (most of which are implicit) are adequate. Hence, in Section 4 we provide a novel and simple solution to the problem that is easily implemented practically. Section 5 provides guidelines on where it is appropriate to use the proposed solution.

Executable versions of all of the BN models described in the paper are freely available for inspection and use at:
www.eecs.qmul.ac.uk/~norman/Models/legal_models.html

2 The problem

One of the most powerful features of BN reasoning is the concept of 'explaining away'. In its simplest form explaining away can be cast in the BN example of Fig. 1, with two possible causes of some event.

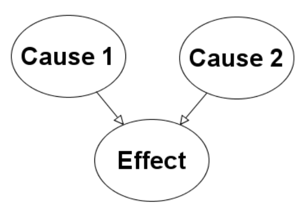

Fig. 1. Generic 'explaining away' pattern

In the classic example cited in [19] the *event* is "wet grass" and the two *causes* are "sprinkler on" and "rain" respectively. For simplicity we assume that all three nodes are Boolean (i.e. hav-

- *Norman Fenton is with the School of Electronic Engineering and Computer Science, Queen Mary University of London, E1 4NS, UK and Agena Ltd 11 Main Street, Caldecote, Cambridge, CB23 7NU. E-mail: norman@eecs.qmul.ac.uk.*
- *Martin Neil is with with the School of Electronic Engineering and Computer Science, Queen Mary University of London, E1 4NS, UK and Agena Ltd E-mail: martin@eecs.qmul.ac.uk.*
- *David Lagnado is with the Dept of Cognitive,Perceptual and Brain Sciences, University College London, London WC1E 6BT, UK. E-mail: d.lagnado@ucl.ac.uk.*

ing just two possible states True and False). A typical implementation of this BN, showing its initial probabilities is shown in Fig. 2.

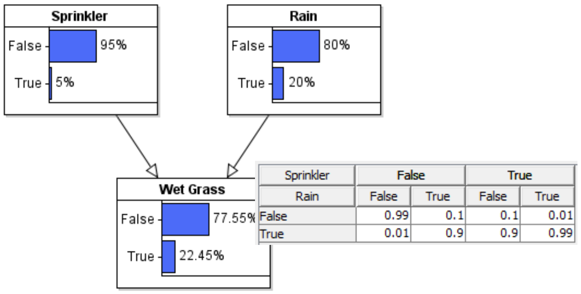

Fig 2 Implementation of BN showing priors and NPT for 'Wet Grass'

If we observe "Wet Grass" is true then the result of the backward inference is shown in Figure 3(a). After this observation we are fairly confident that there must be rain. However, if we discover that the sprinkler is on, as shown in Figure 3(b), our subsequent confidence in rain drops significantly. In other words our belief in rain has been 'explained away' by the sprinkler evidence.

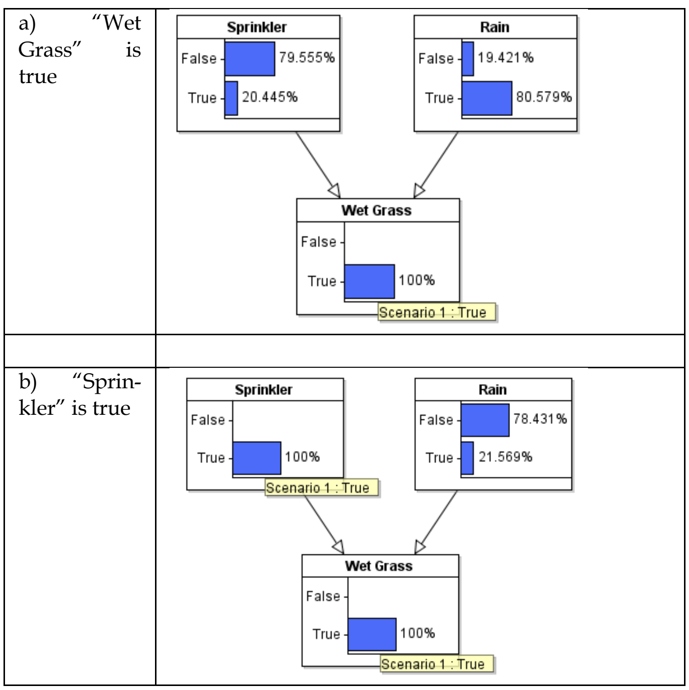

Fig 3 Explaining away evidence

This type of uncertain reasoning is what makes BNs such a powerful and useful tool. Indeed, in the context of legal reasoning Hepler et al [18] considered 'explaining away' as an explicit 'pattern of reasoning'.

However, it turns out that explaining away does not work in a very important class of situations that are especially relevant for reasoning about evidence (notably in legal cases). These are the situations where the causes are ***mutually exclusive***, i.e. if one of them is true then the other must be false, and ***exhaustive***, i.e. where there are no other possible causes. For example, a bruise can either be caused by a criminal attack or be self-inflicted (to keep the example as simple as possible we ignore a third possible cause – accident – but the method proposed in this paper applies to situations where there are an arbitrary number of causes). People regularly report to police stations claiming to have bruises that were the result of a criminal attack. In such situations the police have to determine both whether there is bruising (to merit any investigation at all) and then the most likely cause. This situation is shown in Figure 4.

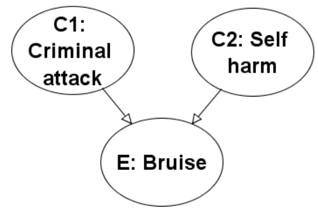

Fig 4 BN with mutually exclusive causes

However the NPTs of the model in Fig. 4 are defined the resulting model should satisfy the following basic axioms:

- Axiom 1 (basic mutual exclusivity): If one of the causes is known to be true then the other must be false.
- Axiom 2 (backward inference for exhaustivity): If E is known to be true then the individual probabilities of the causes being true must sum to 1.
- Axiom 3 (encoding of the priors): Any prior probabilities we have for the cause nodes and the effect node given the causes must be encoded in the model. Also, in its intial state (i.e. before evidence has been entered) the marginal probabilities for each cause node should match its priors.

Axiom 1 immediately rules out the exclusive OR construct as being a solution (in this construct E is defined to be true with probability one when exactly one of the parents is true and false with probability one otherwise). Exclusive OR also fails to satisfy axiom 3 if there is any prior uncertainty about the probability of a cause leading to E.

In fact, it turns out that there is no way to define the NPT of the node E in Fig.4 to satisfy the first two axioms. To see why, let us make the following simple prior assumptions (if we cannot define a solution in this simple case then we certainly cannot for any general case):

1. The prior probabilities of C1 and C2 being true are 0.7 and 0.3 respectively. In other words criminal attacks account for 70% of the cases where people report bruises and self-harm account for the rest.
2. There is a probability of 0.2 that a criminal attack on a person will result in a bruise.
3. There is a probability of 0.9 that an attempt by a person to self-inflict harm will result in a bruise.

From axiom 3 the NPTs for the nodes C1 and C2 must be defined as:

P(C1 = true) = 0.7; P(C1 = false) = 0.3;
P(C2 = true) = 0.3; P(C2 = false) = 0.7;

The NPT for the node E is less obvious. Table 1 shows the completed entries based on the prior information in assumptions 2 and 3.

TABLE 1 NPT FOR NODE E?

C1	False		True	
C2	False	True	False	True

E False	1	0.8	0.1	?
E True	0	0.2	0.9	?

But how are we to define the values in column 4, i.e. when both causes are true? Because of mutual exclusivity this column is meaningless since it represents a situation that is impossible. Unfortunately, there is no mechanism for excluding it from the BN model. In theory we should be able to assign *any* probability values for E in this case without affecting the reasoning. However, it turns out that no matter how we assign them we cannot satisfy the axioms. Consider, for example, the following possibilities

Case 1: There is an argument that the most 'logical' assignment for column 4 of the NPT in Table 1 is to set the probability E being true to be one. After all, the probability is one if either one of the causes is true so it seems sensible that the probability is one if both are true. However, as demonstrated in Fig. 5, the model resulting from this assignment fails to satisfy axioms 1 and 2.

(a) Axiom 1 fails

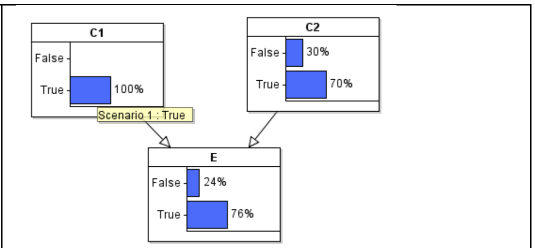

(b) Axiom 2 fails

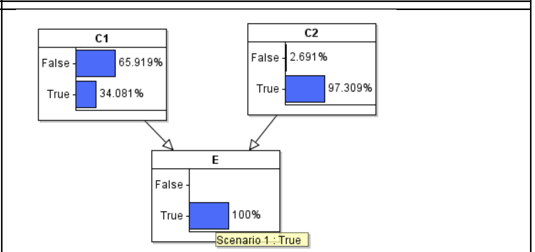

Fig 5 Case 1 does not work

Case 2: Since it is impossible that both causes are true at the same time there is an argument for assigning the probability E being true to be zero in column 4 of the NPT in Table 1. However, although axiom 2 is satisfied in this case, as shown in Fig. 6, axiom 1 is not.

(a) Axiom 1 fails

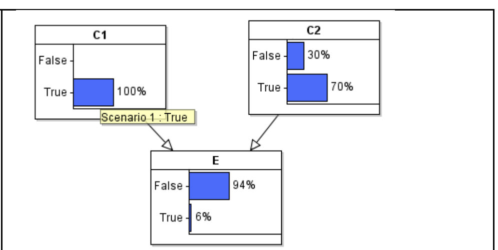

(b) Axiom 2 works

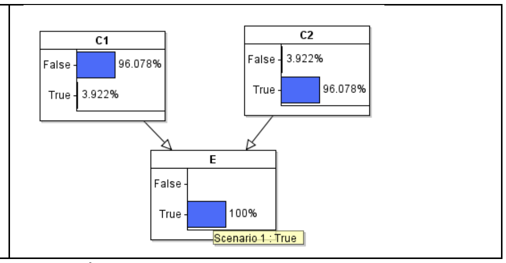

Fig 6 Case 2 does not work

In fact no possible assignment for column 4 of the NPT in Table 1 will result in Axiom 1 being satisfied; this is because if one of the causes is true there will be a non-zero probability that E is true, and by Bayesian inference, this propagates to a non-zero probability that the other cause is true.

3 Previously Proposed Solutions – and Why They Are Inadequate

The fundamental problem is that the notion of mutual exclusivity in BNs is normally encoded via the states of a node rather than by distinct nodes. Indeed, by definition, the states of a BN node represent mutually exclusive and exhaustive possibilities. It follows that the natural way to solve the problem is to collapse causes into a single 'cause' node whose states correspond to the mutually exclusive causes, as shown in Figure 7. Since, by definition, all states of a BN are mutually exclusive this seems a natural and very simple solution.

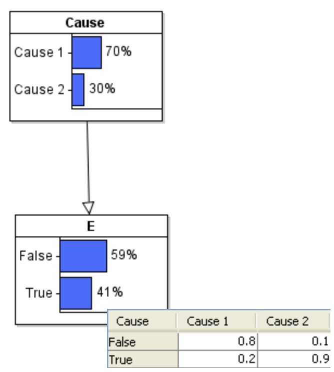

Fig. 7 'Standard' proposed solution with alternative causes modelled as mutually exclusive states of a single node

While Fig. 7 shows that the standard proposed solution satisfies axiom 3, Fig. 8 shows it satisfies the axioms 1 and 2.

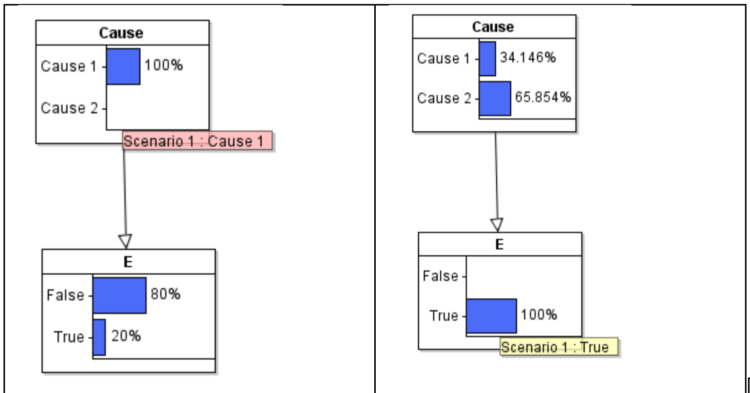

(a) Axiom 1 works | (b) Axiom 2 works

Fig. 8 Solution 1 satisfies the two axioms

So why does this standard solution not work in practice? Because in general there will be:

- separate causal pathways (which may be quite complex) that lead to the different causes
- separate observable pieces of evidence about the separate causes

Schematically the real situation is shown in Figure 9.

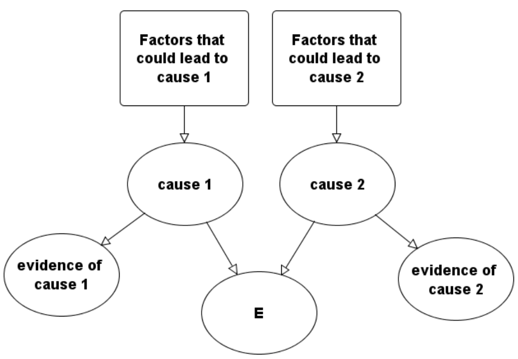

Fig. 9 Separate causal pathways and separate evidence for the mutually exclusive causes

For example, the cause 'self-harm' may itself have been caused by a variety of factors, such as a 'cry for attention', and there may be separate evidence (such as witness testimonies) in support of, or against, that hypothesis.

Now if we apply solution one in this case we would end up with the model shown in Fig.10.

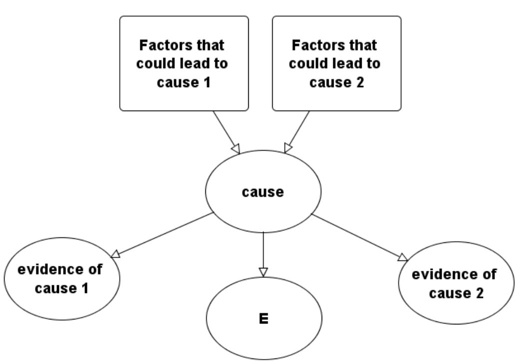

Fig. 10 Collapsing into a single cause node

There is a fundamental practical problem with such a model. The NPTs for both the cause node and the evidence nodes require us to consider meaningless state combinations of their parents. For example, in the simplest possible case there is just a single Factor 1 that could lead to cause 1 and a single Factor 2 that could lead to cause 2. Since the impact of Factor 1 on cause 1 is independent of Factor 2 we should only need to consider the probabilities relating to cause 1 given Factor 1. However, the NPT for the cause node has the structure shown in Table 2.

TABLE 2 NPT FOR CAUSE NODE

Factor 1	True		False	
Factor 2	True	False	True	False
cause 1				
cause 2				

To complete this NPT we have to consider cause 1 given *both* Factor 1 and Factor 2. So, if we know that

P(cause 1 | Factor 1 is true) = 0.2

then we would be forced to enter both

P(cause 1 | Factor 1 is true, Factor 2 is true) = 0.2

P(cause 1 | Factor 1 is true, Factor 2 is false) = 0.2

to capture this information in the NPT. So, not only are the entries of the NPT confusing but half of them are redundant.

The NPT for the node 'evidence of cause 1' (whose structure is shown in Table 3) has a related problem.

TABLE 3 NPT FOR 'EVIDENCE OF CAUSE 1' NODE

Cause	cause 1	cause 2
True		
False		

In this case the entire shaded column is redundant, since the evidence of cause 1 given cause 2 is meaningless.

In summary, merging the causal pathways, into a single node may detract from the semantics of the model and make elicitation and communication difficult. It also goes against the whole spirit of BNs, which seek to provide a compact representation of the joint probability distribution to aid representation and inference. This is achieved by separating out independent causal pathways, so the collapsing of alternative causes into a single causal variable undermines a key motivation for the BN framework.

Hence, we introduce a fourth axiom:

> Axiom 4: (Separate causal pathways). The mutually exlusive causes must be represented as separate nodes.

In addition to the above 'standard' proposed solution there are two other textbook 'solutions' to the problem. We will present these in turn, explaining why they are inadequate in practice.

Proposed solution 2 (introduce dependency between the cause nodes):

In this approach we retain the separate cause nodes (to preserve axiom 4), but introduce a dependency between them to ensure mutual exclusivity, as shown in Fig. 11.

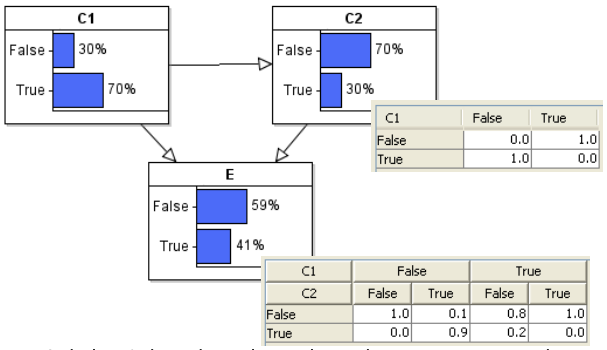

Fig. 11 Solution 2: introduce dependency between cause nodes

One immediate problem with this solution is that again we are forced to define NPT entries for node E for the 'impossible' state combination where both causes are true. However, providing we assign a probability of one to E being False in this case, the model satisfies axioms 1 and 2, as shown in Fig. 12 (Fig. 11 already shows that the model satisfies axiom 3).

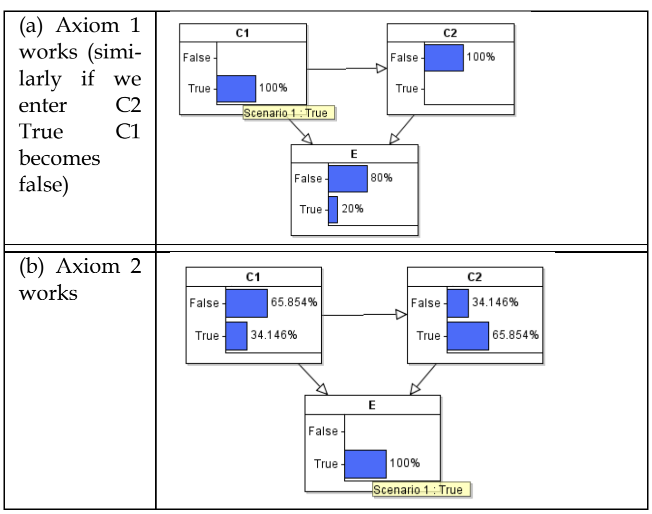

Fig. 12 Solution 2 satisfies the axioms

The primary reason why solution two does not work in practice is because the introduction of a dependency between the causes means that, in the case where there are separate causal pathways for the two causes, cause 1 necessarily becomes part the causal pathway leading to cause 2. This means that, for every possible state of the factors that could lead to cause 2, we have to consider whether cause 1 is true or false. Hence, defining the NPT for cause 2 involves a whole range of meaningless columns and unnecessary complexity. In addition to this fundamental problem there is another unsatisfactory aspect of solution 2: the fact that we had to arbitrarily decide which one of cause 1 and 2 was to be the 'parent' of the other, despite these nodes representing completely alternative and equally likely hypotheses.

Proposed Solution 3: Introducing Boolean constraint node

Jensen and Nielsen [19] propose that mutual exclusivity among different nodes can be ensured by introducing a Boolean constraint node as shown in Figure 13

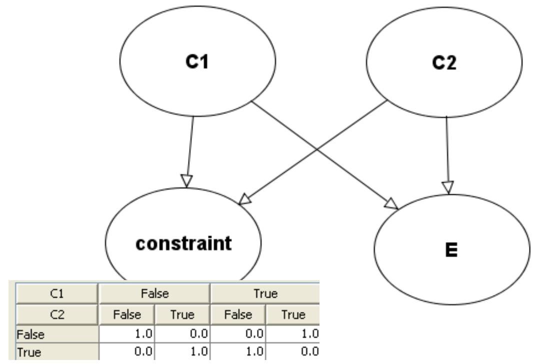

Fig.13 Enforcing mutual exclusivity by introducing a constraint node

As shown, the NPT for the constraint node is defined as true when exactly one of the parents is true and false otherwise. Providing the constraint is always set to be true when the model is run, axioms 1 and 2 are satisfied. However, this proposed solution fails to satisfy axiom 3, because as soon as the constraint is set to true, the priors for the cause nodes change even though no actual evidence about the problem has been entered. This is shown in Fig. 14.

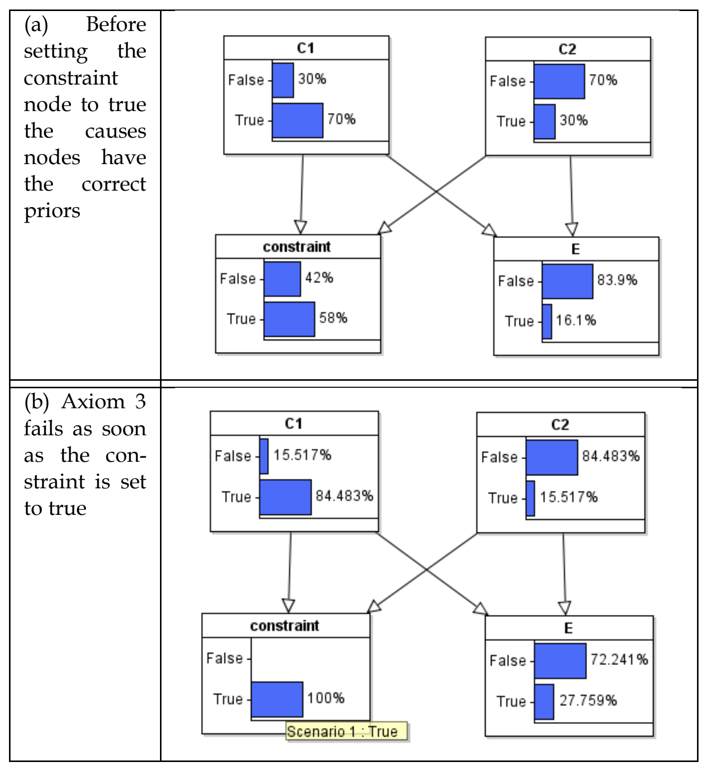

Fig. 14 Solution 3 fails to satisfy axiom 3

It is important to note that in the examples in [19] the priors for the mutually exclusive nodes were assumed to be uniform. In this very special case axiom 3 is satisfied.

4 Proposed Solution

It turns out that there is a satisfactory solution based on introducing a constraint node, but we need a very different type of constraint node to the one proposed in [19].

Our solution is completely general - it works for any number of mutually exclusive and exhaustive causes $C_1, C_2, \ldots, C_n$ - and it works irrespective of the particular priors for the C_i's and for any assignment of probabilities for E given C_i. However, before providing the general proof, it is helpful to illustrate the solution for the simple two-cause case above.

The trick, as shown in Fig. 14, is to introduce a constraint node that has states corresponding to each of the causes (we label these states using lower case letters c1, c2, …to distinguish them from the associated cause nodes) plus a state "NA" representing the impossible state.

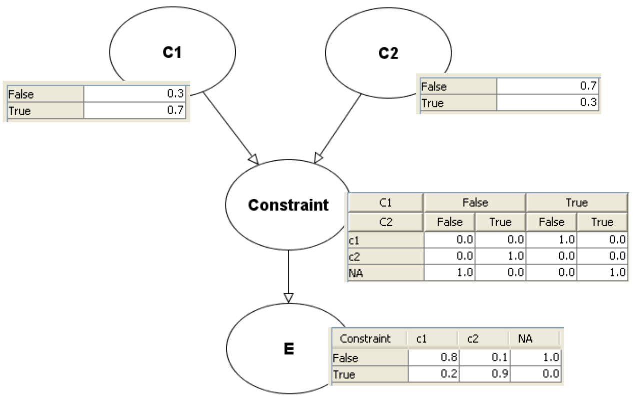

Fig. 15 Solution with special constraint node

We define the NPT for E in such a way that NA is true for any impossible combination of states of the cause nodes.

So the constraint state is equal to:

- c_1 if and only if cause C_1 is true and cause C_2 is false.
- c_2 if and only if cause C_1 is false and cause C_2 is true.
- NA for all other combinations of states for the cause nodes.

It follows that, if we can set as evidence the probability of the state NA as zero, then axioms 1 and 2 will be satisfied. We use the technique of soft evidence to do this. For reasons that will become clear once we provide the proof in the general case, the soft evidence we have to set for the constraint node (in the case where we are assuming as before that P(C_1 is true) = 0.7 and P(C_2 is true) = 0.3) is:

$$c_1 : \frac{1}{0.7}$$

$$c_2 : \frac{1}{0.3}$$

$$NA : 0$$

This results in the marginal values being equal to the priors (to satisfy axiom 3) as shown in Fig. 16.

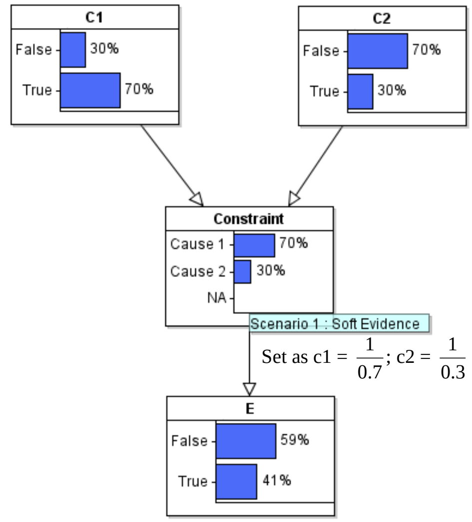

Fig. 16 Constraint model with marginals shown

Fig. 17 shows this solution satisfies axioms 1 and 2 (the results in Fig. 16 and 17 should be compared with the results in Fig. 7 and 8).

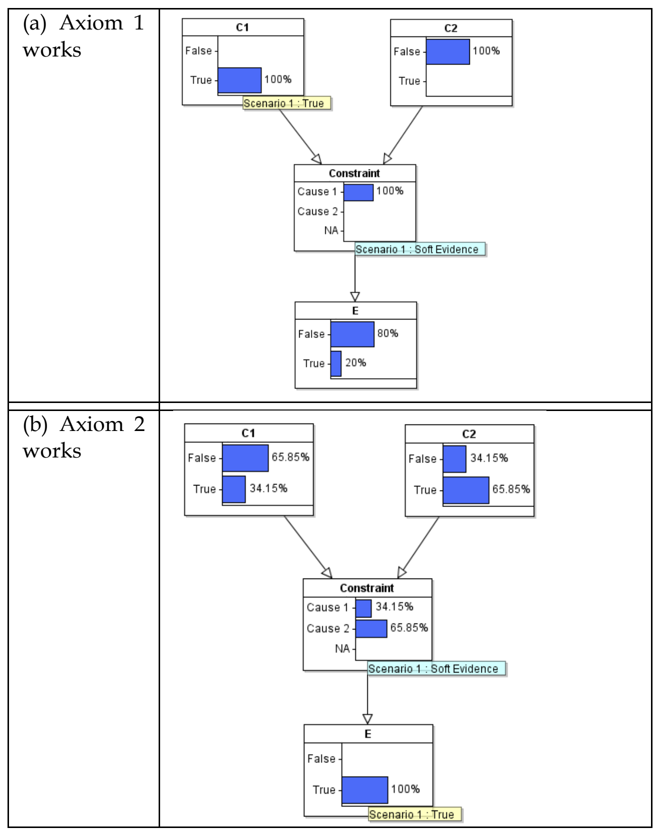

Fig. 17 Constraint solution works

It is important to note that the soft evidence required for the constraint node must be exactly (up to constant multiples) as stated. It is not sufficient to assume that the NA state has zero probability. For example, if we simply give equal weights to c1 and c2 in the constraint node, axiom 3 is not satisfied, as is shown in Fig. 18.

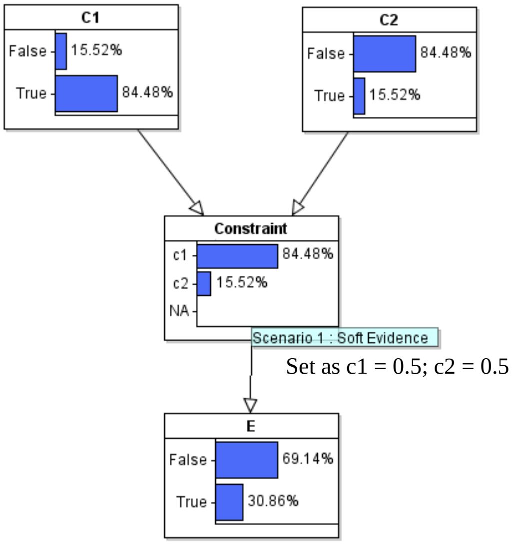

Fig. 18 Axiom 3 is not satisfied when incorrect soft evidence is set

The general solution model is shown in Fig. 19.

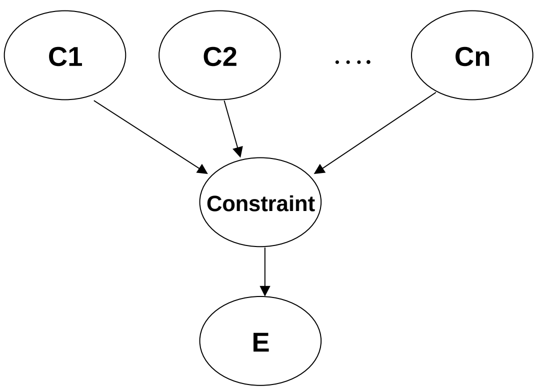

Fig. 19 General solution

We will assume that, for each i=1 to n, the prior probability that node C_i is true is x_i.

The NPT for the constraint node is defined as in Table 2.

TABLE 2 NPT FOR CONSTRAINT NODE

C1	False				True			
C2	False		True		False		True	
…								
Cn	False	True	False	True	False	True	False	True
c1	0	0	0	0	1	0	0	0
c2	0	0	1	0	0	0	0	0
…								
cn	0	1	0	0	0	0	0	0
NA	1	0	0	1	0	1	1	1

So the constraint state is equal to c_i if and only if Cause C_i is true and the other causes are false. For all other combinations of states for the cause nodes, the constraint state is NA.

It follows then that the marginal probabilities for the states of the constraint node are:

$$P(c_1) = P(C_1 = true, C_2 = false, ..., C_n = false) = x_1(1-x_2)...(1-x_n)$$

$$P(c_2) = P(C_1 = false, C_2 = true, C_3 = false) = (1-x_1)x_2...(1-x_n)$$

…

so, in general for each i:

$$P(c_i) = P(C_1 = true, C_2 = false, ..., C_n = false) = x_i \prod_{j \neq i}^{n} (1-x_j)$$

and for the state NA by the complement rule:

$$P(NA) = 1 - (\sum_{i=1}^{n} P(c_i)) = 1 - \sum_{i=1}^{n} \left[x_i \prod_{j \neq i}^{n} (1-x_j) \right]$$

Before establishing what soft evidence weights are required, note that the calculation of the marginals given soft evidence weights of w_1, w_2, …, w_n, w_{NA} respectively for the states of the constraint node, are:

$$P(c_i \mid \text{soft evidence}) = \frac{w_i x_i \prod_{j \neq i}^{n} (1-x_j)}{D} \quad \text{for each } i=1,..,n$$

and

$$P(NA \mid \text{soft evidence}) = \frac{w_{NA} \left[1 - \sum_{i=1}^{n} \left[x_i \prod_{j \neq i}^{n} (1-x_j) \right] \right]}{D}$$

where

$$D = \sum_{i=1}^{n} \left[w_i x_i \prod_{j \neq i}^{n} (1-x_j) \right] + w_{NA} \left[1 - \sum_{i=1}^{n} \left[x_i \prod_{j \neq i}^{n} (1-x_j) \right] \right]$$

But we know that, for axiom 3:

$$P(c_i \mid \text{soft evidence}) = x_i \quad \text{for each } i \text{ and}$$

$$P(NA \mid \text{soft evidence}) = 0$$

Hence, for each i=1,..,n we have:

$$x_i = \frac{w_i x_i \prod_{j \neq i}^{n} (1-x_j)}{D}$$

and so

$$w_i = \frac{D}{\prod_{j \neq i}^{n} (1-x_j)}$$

and $w_{NA} = 0$.

Since D is constant for each of the weights w_1,..,w_n, if follows that we can simply set the weights as:

$$w_i = \frac{1}{\prod_{j \neq i}^{n} (1-x_j)}$$

Example

Here we determine the explicit formulas in the case where we have three causes C1, C2, C3. First, the marginal probabilities for the constraint node states are:

$$P(c_1) = P(C_1 = true, C_2 = false, C_3 = false) = x_1(1-x_2)(1-x_3)$$
$$P(c_2) = P(C_1 = false, C_2 = true, C_3 = false) = (1-x_1)x_2(1-x_3)$$
$$P(c_3) = P(C_1 = false, C_2 = false, C_3 = true) = (1-x_1)(1-x_2)x_3$$
$$P(NA) = 1 - \left[x_1(1-x_2)(1-x_3) + (1-x_1)x_2(1-x_3) + (1-x_1)(1-x_2)x_3\right]$$

The marginals given the soft evidence are:

$$P(c_1 \mid \text{soft evidence}) = \frac{w_1 x_1(1-x_2)(1-x_3)}{D}$$
$$P(c_2 \mid \text{soft evidence}) = \frac{w_2(1-x_1)x_2(1-x_3)}{D}$$
$$P(c_3 \mid \text{soft evidence}) = \frac{w_3(1-x_1)(1-x_2)x_3}{D}$$
$$P(NA) = \frac{w_4\left[x_1(1-x_2)(1-x_3) + (1-x_1)x_2(1-x_3) + (1-x_1)(1-x_2)x_3\right]}{D}$$

where

$$D = [w_1 x_1(1-x_2)(1-x_3)] + [w_2(1-x_1)x_2(1-x_3)] +$$
$$[w_3(1-x_1)(1-x_2)x_3] +$$
$$w_4\left[x_1(1-x_2)(1-x_3) + (1-x_1)x_2(1-x_3) + (1-x_1)(1-x_2)x_3\right]$$

But, since we want

$$P(c_1 \mid \textit{soft evidence}) = x_1$$
$$P(c_2 \mid \textit{soft evidence}) = x_2\,;$$
$$P(c_3 \mid \textit{soft evidence}) = x_3$$
$$P(NA) = 0$$

we have

$$w_1 = \frac{1}{(1-x_2)(1-x_3)}$$
$$w_2 = \frac{1}{(1-x_1)(1-x_3)}$$
$$w_3 = \frac{1}{(1-x_1)(1-x_2)}$$
$$w_4 = 0$$

So, suppose $x_1 = 0.7$; $x_2 = 0.2$; $x_3 = 0.1$, then we simply set the soft evidence as:

$$w_1 = \frac{1}{0.8 \times 0.9} = 1.388888$$
$$w_2 = \frac{1}{0.3 \times 0.9} = 3.7037037$$
$$w_3 = \frac{1}{0.3 \times 0.8} = 4.166666$$
$$w_4 = 0$$

Fig. 20(a) shows how the soft evidence weights are set in the AgenaRisk tool (most BN tools have a similar simple mechanism for setting soft evidence). Fig. 20(b) shows the resulting marginals.

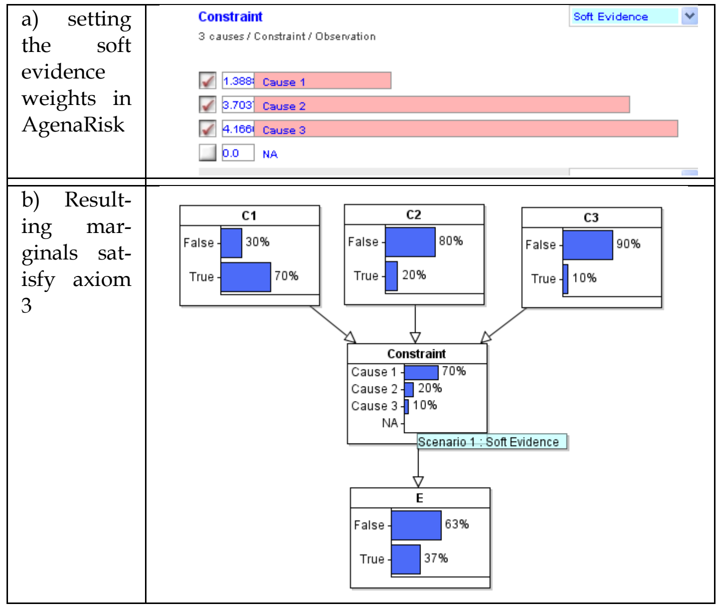

Fig. 20 Soft evidence and resulting marginals

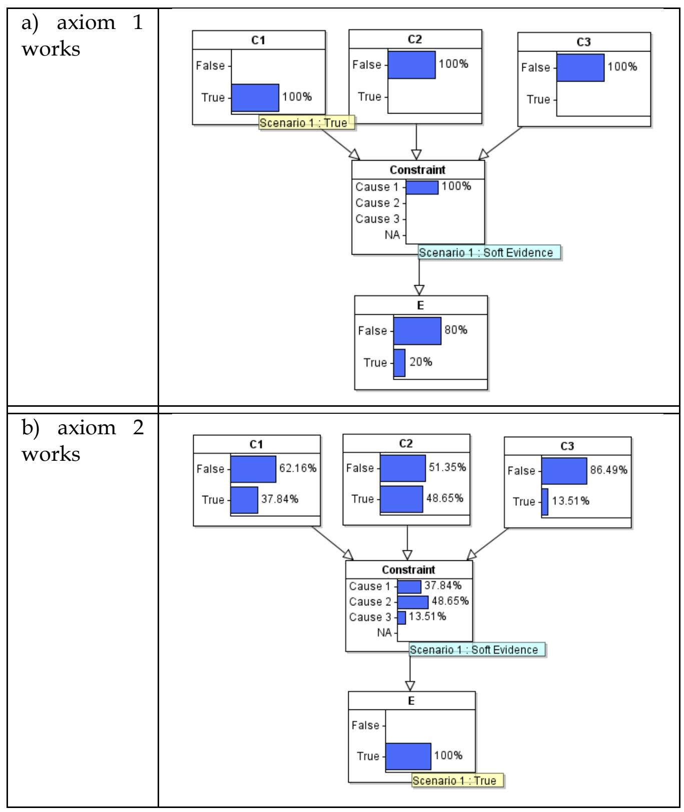

Fig. 21 Constraint solution works

Fig. 21 shows how axioms 1 and 2 are satisfied.

So far, we have made the assumption that the 'true' priors for the cause nodes must sum to one. But even if the causes are mutually exhaustive and exclusive is this a reasonable assumption? The answer is no because the priors for the causes only need to sum to one once we known that the constraint is true (axiom 2). In practice, we may wish to define the 'true' priors based on a wider population of situations other than those where we know one cause must be true. For example, suppose a particular type of electronic unit fails if exactly one of two components A or B fails. In this case the relevant prior probabilities we have will be the failure rates of A and B. Suppose, for example we know that A fails on average once in every 10 uses, while B fails on average twice in every 10 uses. Then the prior probability that A fails is most naturally expressed as 0.1 (and 0.2 for B). Although these do not sum to one, they must sum to one once we know the system has failed, and in that case expect the ratio between the true marginals for A and B should be retained (i.e. the probability A fails should be 1/3 and the probability B fails should be 2/3).

It turns out that the formula for the weights will work exactly as required in such cases. So, in the above example, if we set the soft evidence weights to be 1/0.8 (=1.25) and 1/0.9 (=1.1111) respectively then we get exactly the marginals 1/3 and 2/3 (see Fig. 21).

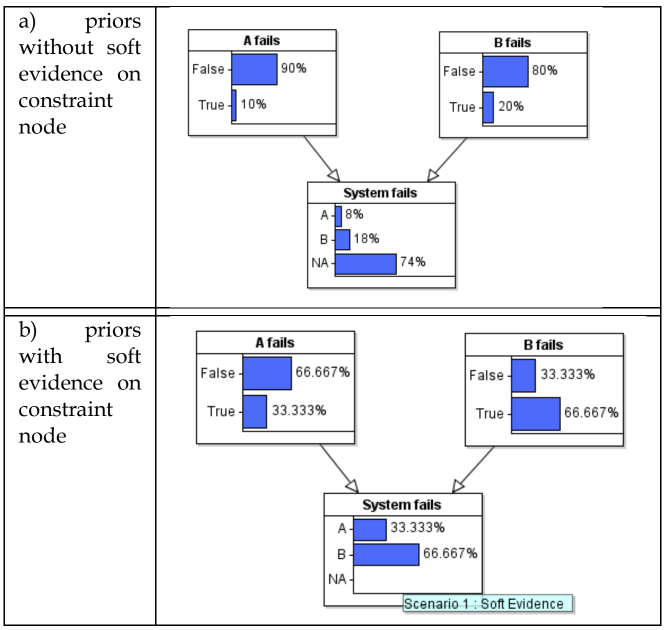

Fig. 22 Dealing with prior 'trues' that do not sum to 1

5 USING THE PROPOSED SOLUTION

Although our proposed solution works in a much broader class of problems than the previously proposed solutions, there are still situations, where care is needed before applying the proposed solution:

Unknown causes

In most realistic situations it will be impossible to identify the set of all potential causes and hence there is a danger in assuming that the set of identified potential causes is exhaustive. This danger is especially pertinent for legal arguments, where the causes represent different hypotheses (for example, the defendant fired a weapon without provocation or fired in self-defence). Strong evidence against one hypothesis here would result in favouring the other. This would be an inappropriate conclusion in the case where an additional potential hypothesis, say 'fire by accident', had been wrongly omitted from the model. The 'fix' to this problem is either to accept that mutual exclusivity and exhaustivity does not apply (and to use probability to deal with unknown causes) or to add a catch-all 'other/unknown' cause to the list of known causes. However, including such a node in a BN model creates different problems, notably that of completing the necessary prior conditional probabilities for the event given a cause that we do not know.

Handling priors and soft evidence weighting in the case where the causes have ancestors

As we have made clear, a major rationale for the proposed solution was to satisfy axiom 4, where the potential causes were themselves dependent on separate complex arguments; i.e. where the causes have ancestor nodes. In such situations any evidence entered on ancestor nodes will change the priors for the cause nodes. For the proposed solution to work properly the soft evidence weights for the constraint node have to therefore be adjusted to take account of any evidence entered for ancestors. No existing BN tools provide a means of making such adjustments automatically. However, there is a practical manual solution as follows:

1. Enter all available evidence about ancestor nodes and run the model.
2. Use the resulting marginals for the cause nodes to set the soft evidence weights based on the same formula as above.
3. Now evidence on descendent nodes can be entered freely as before.
4. If at any stage new evidence about ancestors becomes available, we need to remove any descendant evidence and go to step 2.

6 CONCLUSIONS

BNs have proven to be a very powerful method for reasoning about uncertainty. When there are different possible causes of some event E a simple BN with the structure shown in Fig. 1 should be sufficient to produce the expected reasoning, for both forward and backward inference and 'explaining away'. Yet, in the special but very common situation when the causes are mutually exclusive, it turns out that no formulation of the NPTs in the simple BN structure provides a satisfactory solution.

There have been three proposed solutions to the problem, but we have shown that none of these works in practice. The simplest proposed solution – to replace the separate cause nodes with a single cause node - breaks down in practice whenever there are separate causal pathways leading to the separate causes. The solution involving adding a dependency between the cause nodes breaks down for similar reasons. Finally, the proposed solution to introduce a Boolean constraint node breaks down when the priors for the causes are not 'ignorant' because the solution loses the prior information.

We have presented a practical solution to the problem that involves adding a constraint node that has states corresponding to the causes plus an NA state. The solution requires setting soft evidence weights that a) makes the NA state impossible and b) preserves the priors for the causes. We have presented (with a proof) the formulas required to calculate the soft evidence weights. The solution can be easily implemented in practice in BN tools.

Although we have described the problem and its solution in generic form this is certainly not a purely theoretical exercise. Indeed, we were motivated to find the solution as a result of our practical work in legal reasoning where, as expert witnesses, we have been using BNs to help lawyers better understand the impact of different types of evidence. It was in this context that we came across examples where we needed to model mutually exclusive causes and discovered that the standard BN solutions did not work.

ACKNOWLEDGMENT

This work has benefited from the input of a number of colleagues and experts, and we especially note the contribution from William Marsh.

REFERENCES

[1] Agena Ltd (2010). AgenaRisk, http://www.agenarisk.com.

[2] Aitken, C.G.G. and Taroni, F., Statistics and the evaluation of evidence (2nd Edition). 2004 John Wiley & Sons, Ltd.

[3] Anderson, T., Schum, D. A. and Twining, W. (2005). Analysis of Evidence, Cambridge University Press.

[4] Chase, W. G., & Simon, H. A. (1973). Perception in chess. Cognitive Psychology, 4, 55-81

[5] Cosmides, L. and Tooby, J. (1996). "Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty." Cognition 58 1-73

[6] Cowan, N. (2001) The magical number 4 in short-term memory: a reconsideration of mental storage capacity. Behavorial & Brain Sciences. 24, 87–185.

[7] Crown Court Benchbook, Judicial Studies Board, 2010.

[8] Dawid, A. P. (2002). Bayes's theorem and weighing evidence by juries. In Bayes's Theorem: Proceedings of the British Academy. R. Swinburne. Oxford, Oxford University Press. 113: 71-90.

[9] Dawid, P., Statistics on Trial. Significance, 2005. 2(1): p. 6-8. 10.1111/j.1740-9713.2005.00075.x

[10] Edwards, W. (1991). "Influence Diagrams, Bayesian Imperialism, and the Collins case: an appeal to reason." Cardozo Law Review 13: 1025-1079.

[11] Ericsson, K. & Kintsch, W. (1995). Long-term working memory. Psychological Review 102: 211–245.

[12] Ericsson, K., Charness, N., Feltovich, P., & Hoffman, R. (2006). Cambridge Handbook of Expertise and Expert Performance. Cambridge, UK: Cambridge University Press.

[13] Fenton, N. E. and M. Neil (2010). "Avoiding Legal Fallacies in Practice Using Bayesian Networks." Australian Journal of Legal Philosophy submitted.

[14] Flores, M. J., J. A. Gamez and S. Moral (2005). Abductive Inference in Bayesian Networks: Finding a Partition of the Explanation Space. ECSQARU 2005, LNAI 3571. L. Godo. 3571: 63-75

[15] Gilovich, T., Griffin, D., and Kahneman, D., Heuristics and Biases: The Psychology of Intuitive Judgment. . 2002, Cambridge: Cambridge University Press.

[16] Gobet, F., Lane, P., Croker, S., Cheng, P., Jones, G., Oliver, I. & Pine, J. (2001). Chunking mechanisms in human learning. Trends in Cognitive Science, 5, 236-243.

[17] Halford, G. S., Cowan, N. & Andrews, G. (2007). Separating cognitive capacity from knowledge: a new hypothesis. Trends in Cognitive Science, 11, 236-242.

[18] Hepler, A. B., A. P. Dawid and V. Leucari (2007). "Object-oriented graphical representations of complex patterns of evidence." Law, Probability & Risk 6(1-4): 275-293.

[19] Jensen, F. V. and T. Nielsen (2007). Bayesian Networks and Decision Graphs, Springer-Verlag New York Inc.

[20] Kahneman, D., Slovic, P., and Tversky, A., Judgement Under Uncertainty: Heuristics and Biases. 1982: New York: Cambridge University Press.

[21] Krynski, T. R., & Tenenbaum, J. B. (2007). The role of causality in judgment under uncertainty. Journal of Experimental Psychology: General, 136, 430-450.

[22] Lagnado, D. (2010). Thinking about evidence. To appear in Dawid, P, Twining, W., Vasilaki, M. Eds. Evidence, Inference and Enquiry. British Academy/OUP.

[23] Lagnado, D. A. (2010b). Adverse inferences about alibi evidence. Manuscript submitted for publication.

[24] Lagnado, D. A. and N. Harvey (2008). "The impact of discredited evidence." Psychonomic Bulletin & Review 15: 1166-73.

[25] Lagnado, D. A., Waldmann, M. A., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation. Cues to causal structure. In A. Gopnik, A., & L. E. Schultz, L. (Eds.), Causal learning: Psychology, philosophy, and computation (pp. 154-172). Oxford University Press

[26] Lam, F. C. and W. K. Yeap (1992). "Bayesian updating: On the interpretation of exhaustive and mutually exclusive assumptions." Artificial Intelligence 53(2-3): 245-254.

[27] Lauritzen, S. L. and Spiegelhalter, D. J. (1988). Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems (with discussion). J. R. Statis. Soc. B, 50, No 2, pp 157-224.

[28] Neil M, Fenton NE, Nielsen L, "Building large-scale Bayesian Networks", The Knowledge Engineering Review, 15(3), 257-284, 2000.

[29] Pearl, J. (2000). Causality: Models Reasoning and Inference, Cambridge University Press.

[30] Pearl, J. (1988). Probabilistic reasoning in intelligent systems. Palo Alto, CA, Morgan Kaufmann

[31] Pennington, N. & Hastie, R. (1986). Evidence evaluation in complex decision making. Journal of Personality and Social Psychology, 51, 242-258.

[32] Pennington, N., & Hastie, R. (1992). Explaining the evidence: test of the story model for juror decision making. Journal of Personality and Social Psychology, 62, 189-206.

[33] Perry, B. B. and T. V. Allen (2005). "Causal Reasoning Engine: An Explanation-Based Approach to Syndromic Surveillance." Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS'05) 6: 143b.

[34] Taroni, F., Aitken, C., Garbolino, P., and Biedermann, A., Bayesian Networks and Probabilistic Inference in Forensic Science 2006: John Wiley & Sons.

Norman Fenton is Professor of Computing at Queen Mary University of London and also CEO of Agena, a company that specialises in risk assessment software for critical systems. Norman has been involved in both the research and practice of software engineering for 25 years, but his recent research has focused on causal models (Bayesian Nets) for risk assessment. In addition to extensive work on causal models for software defect prediction, this work covers a wide range of application domains such as legal reasoning (he has been an expert witness in major criminal and civil cases), medical trials, vehicle reliability, embedded software, transport systems, and financial services.

Martin Neil Martin is Professor in Computer Science and Statistics at the Department at Queen Mary University of London, where he teaches decision and risk analysis and software engineering. Martin is also a joint founder and Chief Technology Officer of Agena Ltd, and a Visiting Reader

in the Faculty of Engineering and Physical Sciences, University of Surrey. Martin has over twenty years experience in academic research, teaching, consulting, systems development and project management. His interests cover Bayesian modeling and/or risk quantification in diverse areas: operational risk in finance, systems and design reliability (including software), software project risk, decision support, simulation (using dynamic discretisation as an alternative to Monte Carlo) cost benefit analysis, AI and personalization, and statistical learning.

David Lagnado is Senior Lecturer in Cognitive and Decision Sciences at the Division of Psychology & Language Sciences, University College London. His research focuses on the psychological processes that underlie human learning, reasoning and decision-making. His recent work has explored the role of causal models in evidential reasoning and juror decision making..