The use of Bayes and causal modelling in decision making, uncertainty and risk

Norman Fenton[1] and Martin Neil[2]

2 June 2011

Abstract

The most sophisticated commonly used methods of risk assessment (used especially in the financial sector) involve building statistical models from historical data. Yet such approaches are inadequate when risks are rare or novel because there is insufficient relevant data. Less sophisticated commonly used methods of risk assessment, such as risk registers, make better use of expert judgement but fail to provide adequate quantification of risk. Neither the data-driven nor the risk register approaches are able to model dependencies between different risk factors. Causal probabilistic models (called Bayesian networks) that are based on Bayesian inference provide a potential solution to all of these problems. Such models can capture the complex interdependencies between risk factors and can effectively combine data with expert judgement. The resulting models provide rigorous risk quantification as well as genuine decision support for risk management.

1. Introduction

The 2008-10 credit crisis brought misery to millions around the world, but it at least raised awareness of the need for improved methods of risk assessment. The armies of analysts and statisticians employed by banks and government agencies had failed to predict either the event or its scale until far too late. Yet the methods that could have worked – and which are the subject of this paper – were largely ignored. Moreover, the same methods have the potential to transform risk analysis and decision making in all walks of life. For example:

- **Medical**: Imagine you are responsible for diagnosing a condition and for prescribing one of a number of possible treatments. You have some background information about the patient (some of which is objective like age and number of previous operations, but some is subjective, like 'overweight' and 'prone to stress'); you also have some prior information about the prevalence of different possible conditions (for example, bronchitis may be ten times more likely than cancer). You run some diagnostic tests about which you have some information of the accuracy (such as the chances of false negative and false positive outcomes). You also have various bits of information about the success rates of the different possible treatments and their side effects. On the basis of all this information how do you arrive at a decision of which

[1] Professor and Director of Risk and Information Management Research Group (Queen Mary University of London) and CEO Agena Ltd. Email: norman@eecs.qmul.ac.uk

[2] Computer Science and Statistics (Queen Mary University of London) and CTO Agena Ltd. Email: martin@eecs.qmul.ac.uk

treatment pathway to take? And how would you justify that decision if something went wrong?

- **Legal**: Anybody involved in a legal case (before or during a trial) will see many pieces of evidence. Some of the evidence favours the prosecution hypothesis of guilty and some of the evidence favours the defence hypothesis of innocence. Some of the evidence is statistical (such as the match probability of a DNA sample) and some is purely subjective, such as a character witness statement. It is your duty to combine the value of all of this evidence either to determine if the case should proceed to trial or to arrive at a probability ('beyond reasonable doubt') of innocence. How would you arrive at a decision?
- **Safety**: A transport service (such as a rail network or an air traffic control centre) is continually striving to improve safety, but must nevertheless ensure that any proposed improvements are cost effective and do not degrade efficiency. There are a range of alternative competing proposals for safety improvement, which depend on many different aspects of the current infrastructure (for example, in the case of an air traffic control centre alternatives may include new radar, new collision avoidance detection devices, or improved air traffic management staff training). How do you determine the 'best' alternative taking into account not just cost but also impact on safety and efficiency of the overall system? How would you justify any such decision to a team of government auditors?
- **Financial**: A bank needs sufficient liquid capital readily available in the event of exceptionally poor performance, either from credit or market risk events, or from catastrophic operational failures of the type that brought down Barings in 1995 and almost brought down Société Générale in 2007. It therefore has to calculate and justify a capital allocation that properly reflects its 'value at risk'. Ideally this calculation needs to take account of a multitude of current financial indicators, but given the scarcity of previous catastrophic failures, it is also necessary to consider a range of subjective factors such as the quality of controls in place within the bank. How can all of this information be combined to determine the real value at risk in a way that is acceptable to the regulatory authorities and shareholders?
- **Reliability**: The success or failure of major new products and systems often depends on their reliability, as experienced by end users. Whether it is a high end digital TV, a software operating system, or a complex military vehicle, like an armoured vehicle, too many faults in the delivered product can lead to financial disaster for the producing company or even a failed military mission including loss of life. Hence, pre-release testing of such systems is critical. But no system is ever perfect and a perfect system delivered after a competitor gets to the market first may be worthless. So how do you determine when a system is 'good enough' for release, or how much more testing is needed? You may have hard data in the form of a sequence of test results, but this has to be considered along with subjective data about the quality of testing and the realism of the test environment.

What is common about all of the above problems is that a 'gut-feel' decision based on doing all the reasoning 'in your head' or on the back of an envelope is fundamentally inadequate and increasingly unacceptable. Nor can we base our decision on purely statistical data of 'previous' instances, since in each case the 'risk' we are trying to calculate is essentially unique in many aspects. To deal with these kinds of problems consistently and effectively we need a rigorous method of quantifying uncertainty that enables us to combine data with expert judgement. Bayesian probability, which we introduce in Section 2, is such an approach. We also explain how Bayesian probability combined with causal models (Bayesian networks) enables us to factor in causal relationships and dependencies. In Section 3 we review standard statistical and other approaches to risk assessment, and argue that a proper causal approach based on Bayesian networks is needed in critical cases.

2. Bayes theorem and Bayesian networks

At their heart, all of the problems identified in Section 1 incorporate the basic causal structure shown in Figure 1.

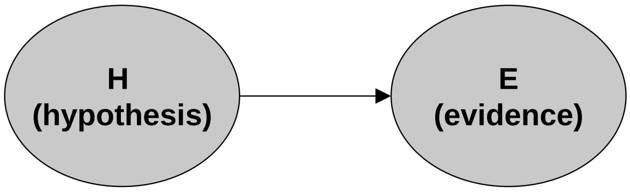

Figure 1 Causal view of evidence

There is some unknown hypothesis *H* about which we wish to assess the uncertainty and make some decision. Does the patient have the particular disease? Is the defendant guilty of the crime? Will the system fail within a given period of time? Is a capital allocation of 5% going to be sufficient to cover operational losses in the next financial year?

Consciously or unconsciously we start with some (unconditional) prior belief about *H* (for example, 'there is a 1 in a 1000 chance this person has the disease'). Then we update our prior belief about *H* once we observe evidence *E* (for example, depending on the outcome of a test our belief about *H* being true might increase or decrease). This updating takes account of the *likelihood* of the evidence, which is the chance of seeing the evidence *E* if *H* is true.

When done formally this type of reasoning is called Bayesian inference, named after Thomas Bayes who determined the necessary calculations for it in 1763. Formally, we start with a prior probability P(*H*) for the hypothesis *H*. The likelihood, for which we also have prior knowledge, is formally the conditional probability of *E* given *H*, which we write as P(*E*|*H*).

Bayes's theorem provides the correct formula for updating our prior belief about *H* in the light of observing *E*. In other words Bayes calculates P(*H*|*E*) in terms of P(*H*) and P(*E*|*H*). Specifically:

$$P(H \mid E) = \frac{P(E \mid H)P(H)}{P(E)} = \frac{P(E \mid H)P(H)}{P(E \mid H)P(H) + (E \mid notH)P(notH)}$$

Example 1: Assume one in a thousand people has a particular disease *H*. Then:

$P(H) = 0.001$, so $P(not\ H) = 0.999$

Also assume a test to detect the disease has 100% sensitivity (i.e. no false negatives) and 95% specificity (meaning 5% false positives). Then if E represents the Boolean variable "Test positive for the disease", we have:

$P(E \mid not\ H) = 0.05$
$P(E \mid H) = 1$

Now suppose a randomly selected person tests positive. What is the probability that the person actually has the disease? By Bayes Theorem this is:

$$P(H \mid E) = \frac{P(E \mid H)P(H)}{P(E \mid H)P(H) + (E \mid notH)P(notH)} = \frac{1 \times 0.001}{1 \times 0.001 + 0.05 \times 0.999} = 0.01963$$

So there is a less than 2% chance that a person testing positive actually has the disease.

Bayes theorem has been used for many years in numerous applications ranging from insurance premium calculations [10], through to web-based personalisation (such as with Google and Amazon). Many of the applications pre-date modern computers (see, e.g. [12] for an account of the crucial role of Bayes theorem in code breaking during World War 2).

However, while Bayes theorem is the only rational way of revising beliefs in the light of observing new evidence, it is not easily understood by people without a statistical/mathematical background. Moreover, the results of Bayesian calculations can appear, at first sight, as counter-intuitive. Indeed, in a classic study [2] when Harvard Medical School staff and students were asked to calculate the probability of the patient having the disease (using the exact assumptions stated in Example 1) most gave the wildly incorrect answer of 95% instead of the correct answer of less than 2%. The potential implications of such incorrect 'probabilistic risk assessment' are frightening. In many cases, lay people only accept Bayes theorem as being 'correct' and are able to reason correctly, when the information is presented in alternative graphical ways, such as using event trees and frequencies (see [3] and [6] for a comprehensive investigation of these issues). But these alternative presentation techniques do not scale up to more complex problems.

If Bayes theorem is difficult for lay people to compute and understand in the case of a single hypothesis and piece of evidence (as in Figure 1), the difficulties are obviously compounded when there are multiple related hypotheses and evidence as in the example of Figure 2.

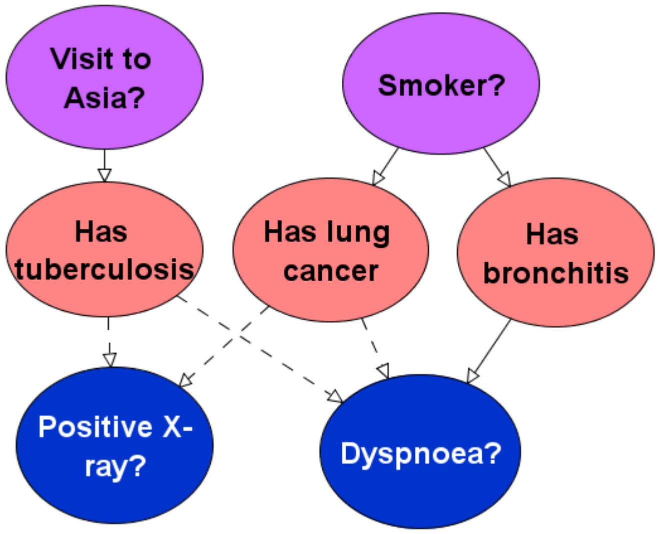

Figure 2 Bayesian network for diagnosing disease

As in Figure 1 the nodes in Figure 2 represent variables (which may be known or unknown) and the arcs represent causal (or influential) relationships. Once we have relevant prior and conditional probabilities associated with each variable (such as the examples shown in Figure 3) the model is called a ***Bayesian network (BN)***.

Yes	0.01
No	0.99

Probability Table for "Visit to Asia?"

Smoker?	Yes	No
Yes	0.6	0.3
No	0.4	0.7

Probability Table for "Bronchitis?"

Figure 3 Node Probability Table (NPT) examples

The BN in Figure 2 is intended to model the problem of diagnosing diseases (TB, Cancer, Bronchitis) in patients attending a chest clinic. Patients may have symptoms (like dyspnoea – shortness of breath) and can be sent for diagnostic tests (X-ray); there may be also underlying causal factors that influence certain diseases more than others (such as smoking, visit to Asia).

To use Bayesian inference properly in this type of network necessarily involves multiple applications of Bayes Theorem in which evidence is 'propagated' throughout. This process is complex and quickly becomes infeasible when there are many nodes and/or nodes with multiple states. This complexity is the reason why, despite its known benefits, there was for many years little appetite to use Bayesian inference to solve real-world decision and risk problems. Fortunately, due to breakthroughs in the late 1980s that produced efficient calculations algorithms [9], [11] there are now widely available tools such as [1] that enable anybody to do the Bayesian calculations without ever having to understand, or even look at, a mathematical formula. These developments were the catalyst for an explosion of interest in BNs. Using such a tool we can do the kind of powerful reasoning shown in Figure 4.

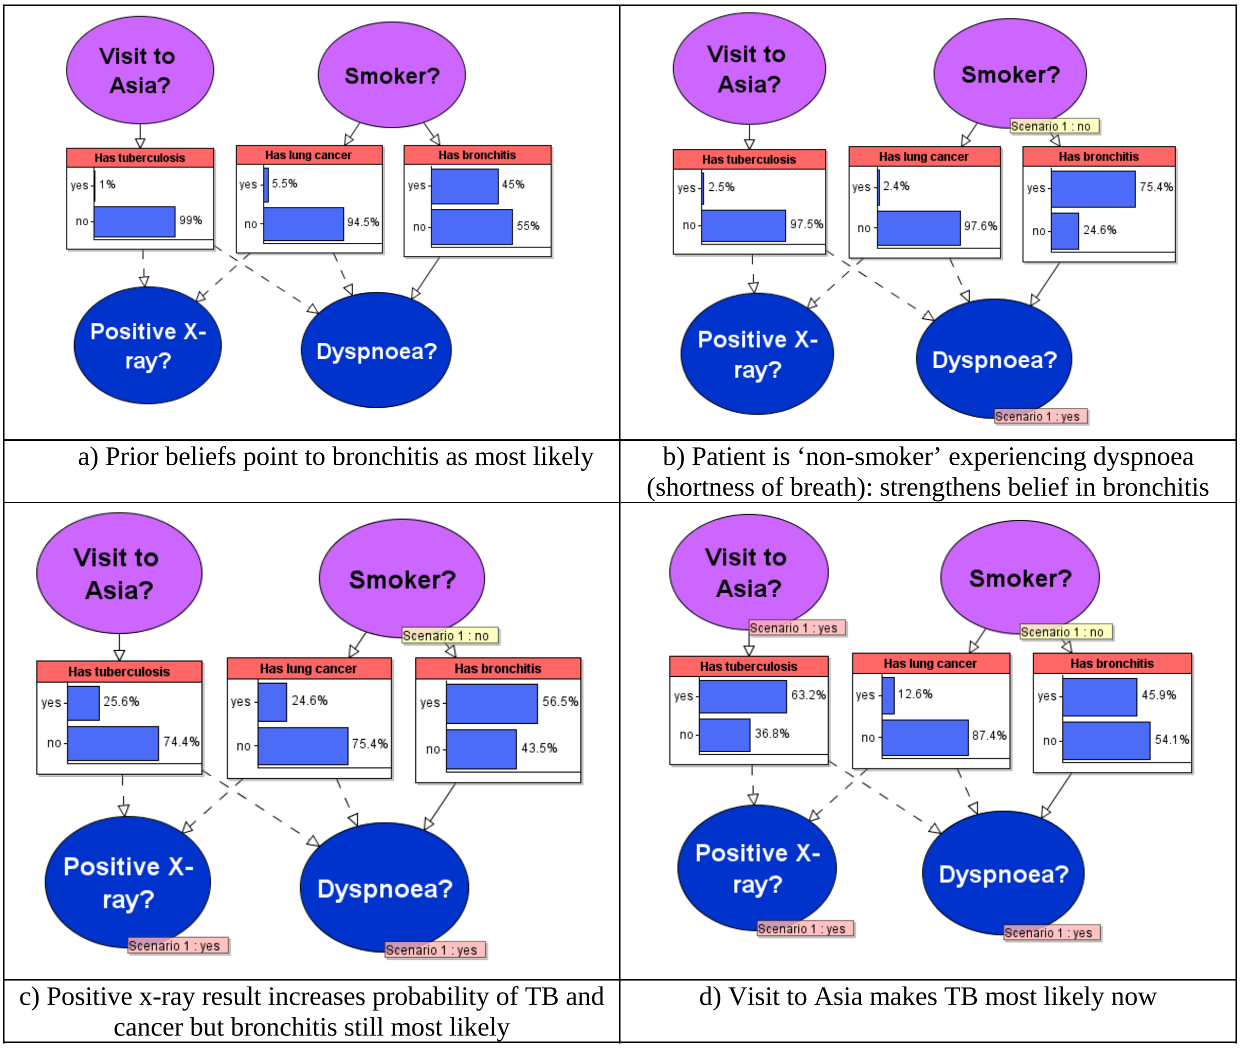

a) Prior beliefs point to bronchitis as most likely

b) Patient is 'non-smoker' experiencing dyspnoea (shortness of breath): strengthens belief in bronchitis

c) Positive x-ray result increases probability of TB and cancer but bronchitis still most likely

d) Visit to Asia makes TB most likely now

Figure 4 Reasoning within the Bayesian network

Specifically:

- With the prior assumptions alone (Figure 4a) Bayes theorem computes what are called the prior marginal probabilities for the different disease nodes (note that we did not 'specify' these probabilities – they are computed automatically; what we specified were the conditional probabilities of these diseases given the various states of their parent nodes). So, before any evidence is entered the most likely disease is bronchitis (45%).
- When we enter evidence about a particular patient the probabilities for all of the unknown variables get updated by the Bayesian inference. So, (in Figure 4b) once we enter the evidence that the patient has dyspnoea and is a non-smoker, our belief in bronchitis being the most likely disease increases (75%).
- If a subsequent X-ray test is positive (Figure 4b) our belief in both TB (26%) and cancer (25%) are raised but bronchitis is still the most likely (57%).
- However, if we now discover that the patient visited Asia (Figure 4d) we overturn our belief in bronchitis in favour of TB (63%).

Note that we can enter any number of observations anywhere in the BN and update the marginal probabilities of all the unobserved variables. As the above example demonstrates, this can yield some exceptionally powerful analyses that are simply not possible using other types of reasoning and classical statistical analysis methods.

In particular, BNs offer the following benefits:

- Explicitly model causal factors:

- Reason from effect to cause and vice versa

- Overturn previous beliefs in the light of new evidence (also called 'explaining away')

- Make predictions with incomplete data

- Combine diverse types of evidence including both subjective beliefs and objective data.

- Arrive at decisions based on visible auditable reasoning (Unlike black-box modelling techniques there are no "hidden" variables and the inference mechanism is based on a long-established theorem).

With the advent of the BN algorithms and associated tools, it is therefore no surprise that BNs have been used in a range of applications that were not previously possible with Bayes Theorem alone. A comprehensive (and easily accessible) overview of BN applications, with special emphasis on their use in risk assessment, can be found in [5].

It is important to recognise that the core intellectual overhead in using the BN approach is in defining the model structure and the NPTs – the actual Bayesian calculations can and must always be carried out using a tool. However, while these tools enable large-scale BNs to be executed efficiently, most provide little or no support for users to actually build large-scale BNs, nor to interact with them easily. Beyond a graphical interface for building the structure, BN-builders are left to struggle with the following kinds of practical problems that combine to create a barrier to the more widespread use of BNs:

- Eliciting and completing the probabilities in very large NPTs manually (e.g. for a node with 5 states having three parents each with 5 states, the NPT requires 625 entries);
- Dealing with very large graphs that contain similar, but slightly different "patterns" of structure ;
- Handling continuous, as well as discrete variables.

Fortunately, recent algorithm and tool developments (also described in [5]) have gone a long way to addressing these problems and may lead to a 'second wave' of widespread BN applications. But before BNs are used more widely in critical risk assessment and decision making, there needs to be a fundamental cultural shift away from the current standard approaches to risk assessment, which we address next.

3. From statistical models and risk registers to causal models

3.1 Prediction based on correlation is not risk assessment

Standard statistical approaches to risk assessment seek to establish hypotheses from relationships discovered in data. Suppose we are interested, for example, in the risk of fatal automobile crashes. Table 1 gives the number of crashes resulting in fatalities in the USA in 2008 broken down by month (source: US National Highways Traffic Safety Administration). It also gives the average monthly temperature.

Table 1 Fatal automobile crashes per month

Month	Total fatal crashes	Average monthly temperature (F)
January	297	17.0
February	280	18.0
March	267	29.0
April	350	43.0
May	328	55.0
June	386	65.0
July	419	70.0
August	410	68.0
September	331	59.0
October	356	48.0
November	326	37.0
December	311	22.0

We plot the fatalities and temperature data in a scatterplot graph as shown in Figure 5.

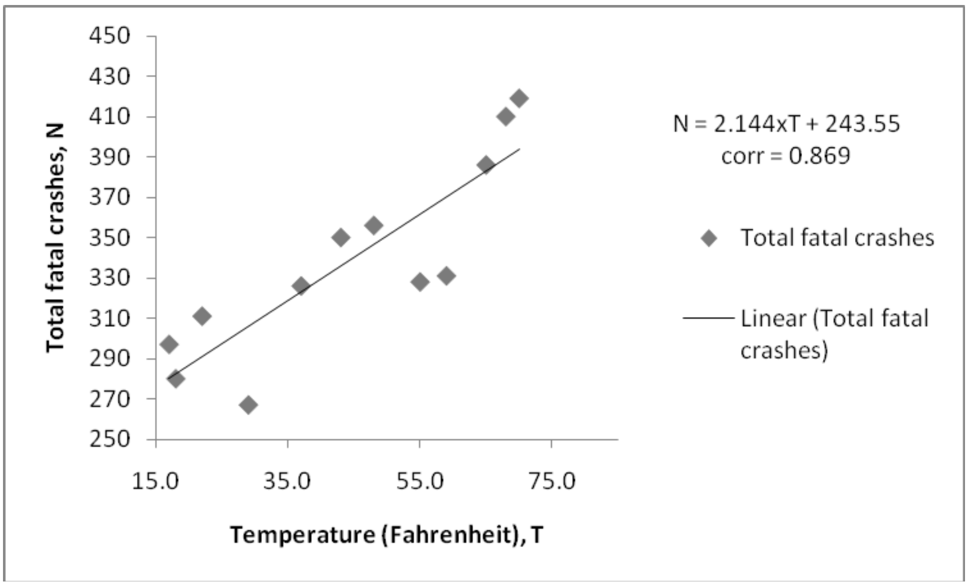

Figure 5 Scatterplot of temperature against road fatalities (each dot represents a month)

There seems to be a clear relationship between temperature and fatalities – fatalities increase as the temperature increases. Indeed, using the standard statistical tools of correlation and *p*-values, statisticians would accept the hypothesis of a relationship as 'highly significant' (the correlation coefficient here is approximately 0.869 and it comfortably passes the criteria for a *p*-value of 0.01).

However, in addition to serious concerns about the use of *p*-values generally (as described comprehensively in [13]), there is an inevitable temptation arising from such results to infer causal links such as, in this case, higher temperatures cause more fatalities. Even though any introductory statistics course teaches that correlation is not causation, the regression equation is typically used for prediction (e.g. in this case the equation relating N to T is used to predict that at 80F we might expect to see 415 fatal crashes per month).

But there is a grave danger of confusing prediction with risk assessment. For risk assessment and management the regression model is useless, because it provides no explanatory power at all. In fact, from a risk perspective this model would provide irrational, and potentially dangerous, information: it would suggest that if you want to minimise your chances of dying in an automobile crash you should do your driving when the highways are at their most dangerous, in winter.

One obvious improvement to the model, if the data is available, is to factor in the number of miles travelled (i.e. journeys made). But there are other underlying causal and influential factors that might do much to explain the apparently strange statistical observations and provide better insights into risk. With some common sense and careful reflection we can recognise the following:

- Temperature influences the highway conditions (which will be worse as temperature decreases).
- Temperature also influences the number of journeys made; people generally make more journeys in spring and summer and will generally drive less when weather conditions are bad.
- When the highway conditions are bad people tend to reduce their speed and drive more slowly. So highway conditions influence speed.
- The actual number of crashes is influenced not just by the number of journeys, but also the speed. If relatively few people are driving, and taking more care, we might expect fewer fatal crashes than we would otherwise experience.

The influence of these factors is shown in Figure 6:

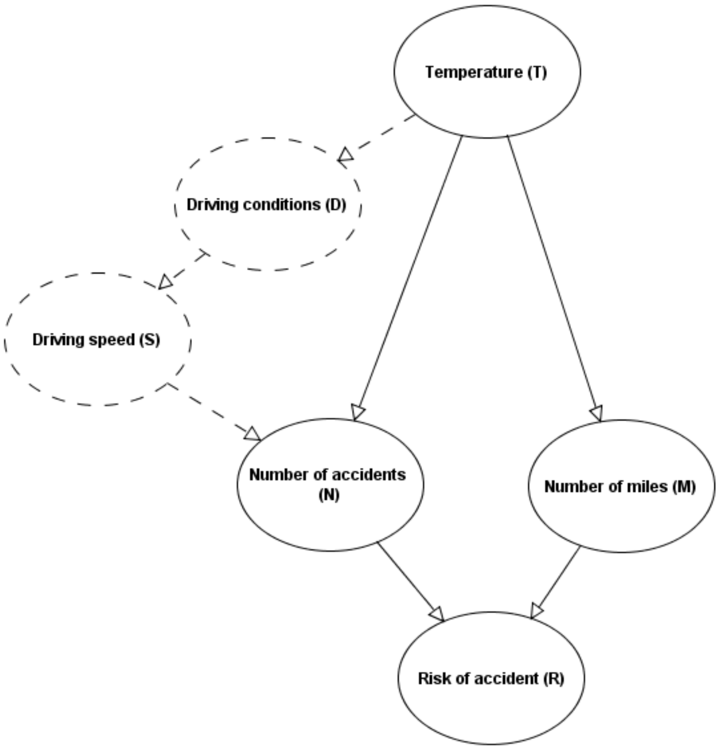

Figure 6 Causal model for fatal crashes

The crucial message here is that the model no longer involves a simple single causal explanation; instead it combines the statistical information available in a database (the ‘objective’ factors) with other causal ‘subjective’ factors derived from careful reflection. These factors now interact in a non-linear way that helps us to arrive at an explanation for the observed results. Behaviour, such as our natural caution to drive slower when faced with poor road conditions, leads to lower accident rates (people are known to adapt to the perception of risk by tuning the risk to tolerable levels. - this is formally referred to as risk homeostasis). Conversely, if we insist on driving fast in poor road conditions then, irrespective of the temperature, the risk of an accident increases and so the model is able to capture our intuitive beliefs that were contradicted by the counterintuitive results from the simple regression model.

The role played in the causal model by driving speed reflects human behaviour. The fact that the data on the average speed of automobile drivers was not available in a database explains why this variable, despite its apparent obviousness, did not appear in the statistical regression model. The situation whereby a statistical model is based only on available data, rather than on reality, is called “conditioning on the data”. This enhances convenience but at the cost of accuracy.

By accepting the statistical model we are asked to defy our senses and experience and actively ignore the role unobserved factors play. In fact, we cannot even explain the results without recourse to factors that do not appear in the database. This is a key point: with causal models we seek to dig deeper behind and underneath the data to explore richer relationships missing from over-simplistic statistical models. In doing so we gain insights into how best to control risk and uncertainty. The regression model, based on the idea that we can predict automobile crash fatalities based on temperature, fails to answer the substantial question: how can we control or influence behaviour to reduce fatalities. This at least is achievable; control of weather is not.

3.2 Risk Registers do not help quantify risk

While statistical models based on historical data represent one end of a spectrum of sophistication for risk assessment, at the other end is the commonly used idea of a 'risk register'. In this approach, there is no need for past data; in considering the risks of a new project risk managers typically prepare a list of 'risks' that could be things like:

- Some key people you were depending on become unavailable
- A piece of technology you were depending on fails.
- You run out of funds or time

The very act of listing and then prioritising risks, means that mentally at least risk managers are making a decision about which risks are the *biggest*. Most standard texts on risk propose decomposing each risk into two components:

- 'Probability' (or likelihood) of the risk
- 'Impact' (or loss) the risk can cause

The most common way to measure each risk is to multiply the probability of the risk (however you happen to measure that) with the impact of the risk (however you happen to measure that) as in Figure 7.

$$\textbf{Risk} = \boxed{\textbf{Probability}} \times \boxed{\textbf{Impact}}$$

Figure 7 Standard impact-based risk measure

The resulting number is the 'size' of the risk - it is based on analogous 'utility' measures. This type of risk measure is quite useful for prioritising risks (the bigger the number the 'greater' the risk) but it is normally impractical and can be irrational when applied blindly. We are not claiming that this formulation is wrong. Rather, we argue that it is normally not sufficient for decision-making.

One immediate problem with the risk measure of Figure 7 is that, normally, you cannot directly get the numbers you need to calculate the risk without recourse to a much more detailed analysis of the variables involved in the situation at hand.

> **Example**: By destroying the meteor in the film Armageddon Bruce Willis saved the world. Both the chance of the meteor strike and the consequences of such a strike were so high, that nothing much else mattered except to try to prevent the strike. In popular terminology what the world was confronting was a truly massive 'risk'. But if the NASA scientists in the film had measured the size of the risk using the formula in Figure 7 they would have discovered such a measure was

irrational, and it certainly would not have explained to Bruce Willis and his crew why their mission made sense. Specifically:

- *Cannot get the Probability number* (for meteor strikes earth). According to the NASA scientists in the film the meteor was on a direct collision course with earth. Does that make it a certainty (i.e. a 100% chance) of it striking Earth? Clearly not, because if it was then there would have been no point in sending Bruce Willis and his crew up in the space shuttle. The probability of the meteor striking Earth is *conditional* on a number of ***control*** events (like intervening to destroy the meteor) and ***trigger*** events (like being on a collision course with Earth). It makes no sense to assign a direct probability without considering the events it is conditional on. **In general it makes no sense (and would in any case be too difficult) for a risk manager to give the unconditional probability of every 'risk' irrespective of relevant controls and triggers**. This is especially significant when there are, for example, controls that have never been used before (like destroying the meteor with a nuclear explosion).

- *Cannot get the Impact number (*for meteor striking earth). Just as it makes little sense to attempt to assign an (unconditional) probability to the event "Meteor strikes Earth', so it makes little sense to assign an (unconditional) number to the *impact* of the meteor striking. Apart from the obvious question "impact on what?", we cannot say what the impact is without considering the possible ***mitigating*** events such as getting people underground and as far away as possible from the impact zone.

- *Risk score is meaningless.* Even if we could get round the two problems above what exactly does the resulting number mean? Suppose the (conditional) probability of the strike is 0.95 and, on a scale of 1 to 10, the impact of the strike is 10 (even accounting for mitigants). The meteor 'risk' is 9.5, which is a number close to the highest possible 10. But it does not measure anything in a meaningful sense

- *It does not tell us what we really need to know.* What we really need to know is the probability, given our current state of knowledge, that there will be massive loss of life.

In addition to the problem of measuring the size of each individual risk in isolation, risk registers suffer from the following problems:

- However the individual risk size is calculated, the cumulative risk score measures the total project risk. Hence, there is a paradox involved in such an approach: the more carefully you think about risk (and hence the more individual risks you record in the risk register) the higher the overall risk score becomes. Since higher risk scores are assumed to indicate greater risk of failure it seems to follow that your best chance of a new project succeeding is to simply ignore, or under-report, any risks.
- Different projects or business divisions will assess risk differently and tend to take a localised view of their own risks and ignore that of others. This "externalisation" of risk to others is especially easy to ignore if their interests are not represented when constructing the register. For example the IT department may be forced to accept the deadlines imposed by the marketing department.
- A risk register does not record "opportunities" or "serendipity" and so does not deal with upside uncertainty, only downside.
- Risks are not independent. For example, in most circumstances cost, time and quality will be inextricably linked; you might be able to deliver faster but only by sacrificing quality. Yet "poor quality" and "missed delivery" will appear as separate risks on the

register giving the illusion that we can control or mitigate one independently of the other. In the subprime loan crisis of 2008 there were three risks: 1) extensive defaults on subprime loans, 2) growth in novelty and complexity of financial products and 3) failure of AIG (American International Group Inc.) to provide insurance to banks when customers default. Individually these risks were assessed as 'small'. However, when they occurred together the total risk was much larger than the individual risks. In fact, it never made sense to consider the risks individually at all.

Hence, risk analysis needs to be coupled with an assessment of the impact of the underlying events, one on another, and in terms of their effect on the ultimate outcomes being considered. The accuracy of the risk assessment is crucially dependent on the fidelity of the underlying model; the simple formulation of Figure 7 is insufficient. Instead of *going through the motions* to assign numbers without actually doing much thinking, we need to consider *what lies under the bonnet.*

Risk is a function of how closely connected events, systems and actors in those systems might be. Proper risk assessment requires a holistic outlook that embraces a causal view of interconnected events. Specifically to get rational measures of risk you need a causal model, as we describe next. Once you do this measuring risk starts to make sense, but it requires an investment in time and thought.

3.2.1 Thinking about risk using causal analysis

It is possible to avoid all the above problems and ambiguities surrounding the term risk by considering the causal context in which risks happen (in fact everything we present here applies equally to opportunities but we will try to keep it as simple as possible). The key thing is that a risk is an ***event*** that can be characterised by a causal chain involving (at least):

- the event itself
- at least one consequence event that characterises the impact
- one or more *trigger* (i.e. initiating) events
- one or more *control* events which may stop the trigger event from causing the risk event
- one or more *mitigating* events which help avoid the consequence event

This is shown in the example of Figure 8.

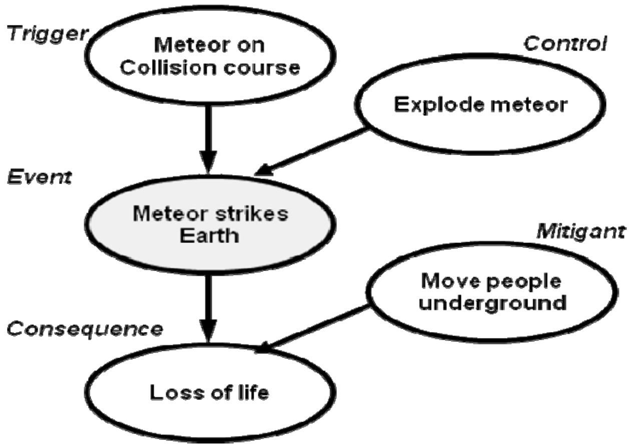

Figure 8 Meteor strike risk

With this causal perspective, a risk is therefore actually characterised not by a single event, but by a set of events. These events each have a number of possible outcomes (to keep things as simple as possible in the example here we will assume each has just two outcomes true and false so we can assume "Loss of life" here means something like 'loss of at least 80% of the world population').

The 'uncertainty' associated with a risk is not a separate notion (as assumed in the classic approach). Every event (and hence every object associated with risk) has uncertainty that is characterised by the event's probability distribution. Triggers, controls, and mitigants are all inherently uncertain. The sensible risk measures that we are proposing are simply the probabilities you get from running the BN model. Of course, before you can run it you still have to provide the prior probability values. But, in contrast to the classic approach, the probability values you need to supply are relatively simple and they make sense. And you never have to define vague numbers for 'impact'.

Example. To give you a feel of what you would need to do, in the Armageddon BN example of Figure 8 for the uncertain event "Meteor strikes Earth" we still have to assign some probabilities. But instead of second guessing what this event actually means in terms of other conditional events, the model now makes it explicit and it becomes much easier to define the necessary *conditional* probability. What we need to do is define the probability of the meteor strike given each combination of parent states as shown in Figure 9.

Meteor on collision course	False		True	
Explode meteor	False	True	False	True
False	1.0	1.0	0.0	0.8
True	0.0	0.0	1.0	0.2

Figure 9 Conditional probability table for "Meteor strike Earth"

For example, if the meteor is on a collision course then the probability of it striking the earth is 1, if it is not destroyed, and 0.2, if it is. In completing such a table we no longer have to try to 'factor in' any implicit conditioning events like the meteor trajectory.

There are some events in the BN for which we *do* need to assign unconditional probability values. These are represented by the nodes in the BN that have no parents; it makes sense to get unconditional probabilities for these because, by definition, they are not conditioned on anything (this is obviously a choice we make during our analysis). Such nodes can generally be only triggers, controls or mitigants. An example, based on dialogue from the film, is shown in Figure 10.

False	0.0010
True	0.999

Figure 10 Probability table for "Meteor on collision course with Earth"

Once we have supplied the priors probability values a BN tool will run the model and generate all the measures of risk that you need. For example, when you run the model using only the initial probabilities the model (as shown in Figure 11) computes the probability of the meteor striking Earth as 99.1% and the probability of loss of life (meaning at least 80% of the world population) is about 94%.

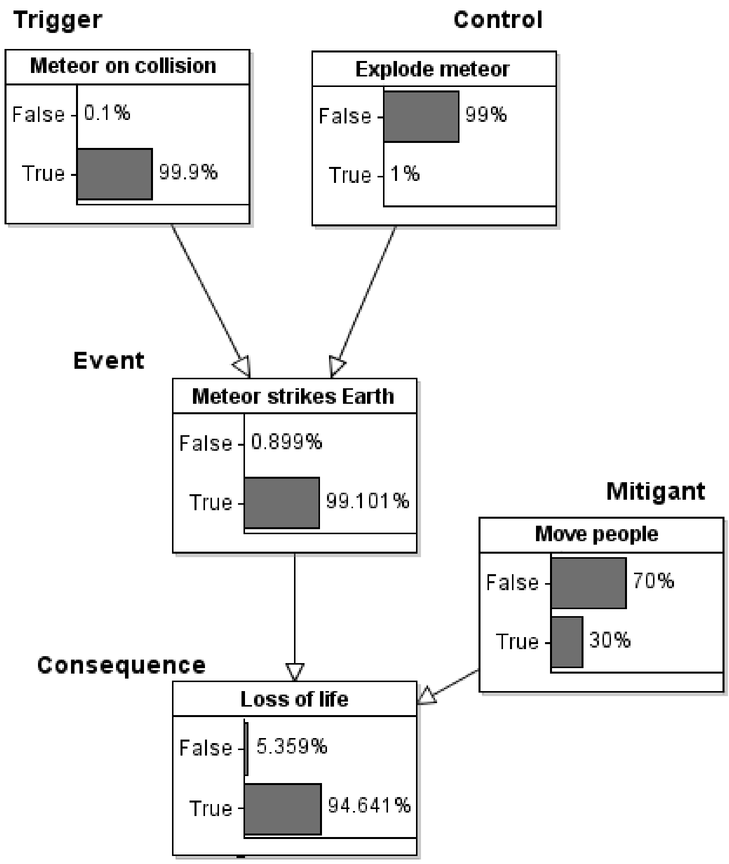

Figure 11 Initial risk of meteor strike

In terms of the difference that Bruce Willis and his crew could make we run two scenarios: One where the meteor is exploded and one where it is not. The results of both scenarios are shown together in Figure 12.

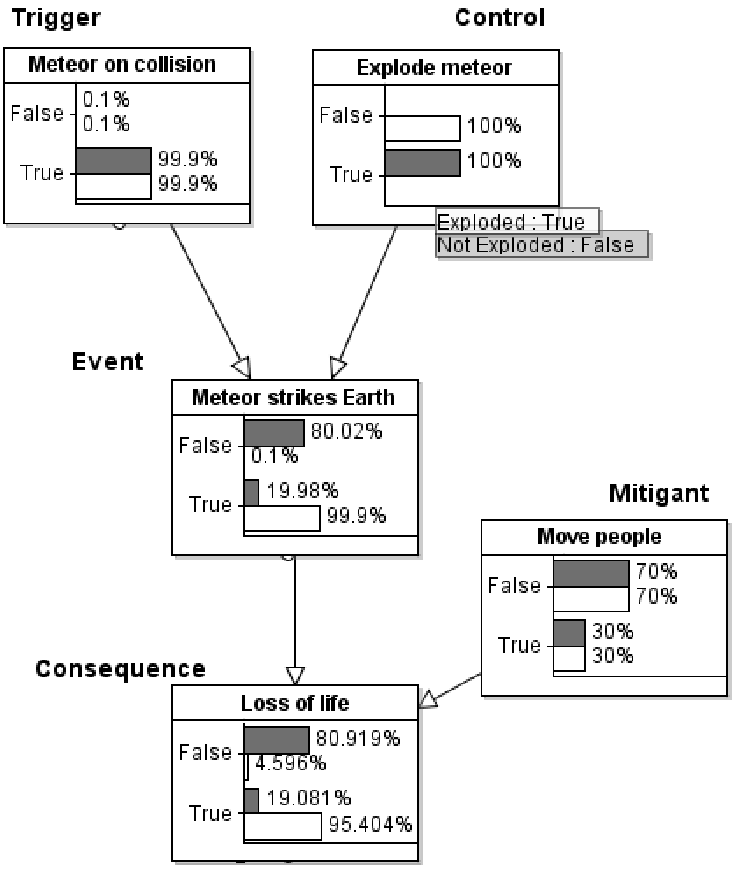

Figure 12 The potential difference made by Bruce Willis and crew

> Reading off the values for the probability of "loss of life" being false we find that we jump from just over 4% (when the meteor is not exploded) to 81% (when the meteor is exploded). This massive increase in the chance of saving the world clearly explains why it merited an attempt.

Clearly risks in this sense depend on stakeholders and perspectives, but these perspectives can be easily combined as shown in Figure 13 for 'flood risk' in some town.

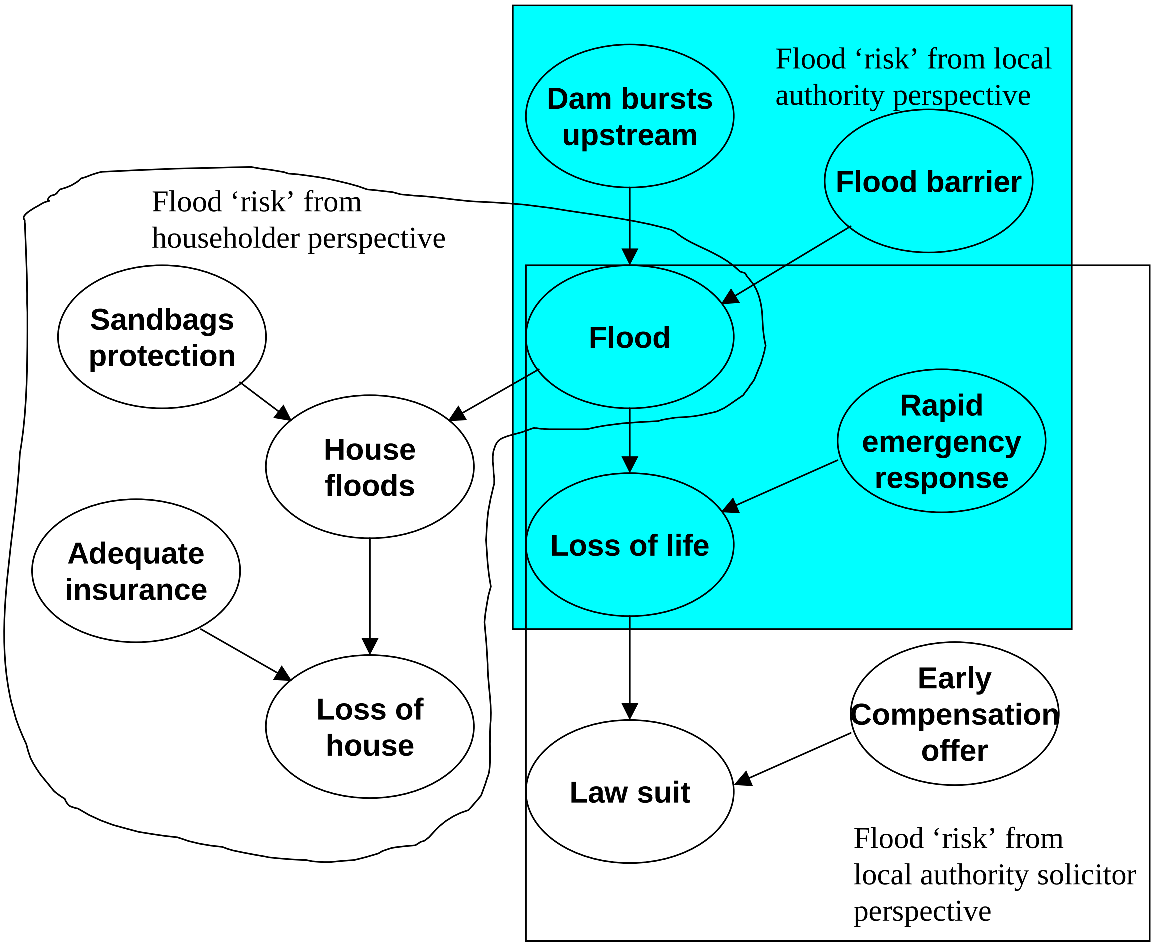

Figure 13 Incorporating different risk perspectives

The types of events are all completely interchangeable depending on the perspective. From the perspective of the local authority the risk event is 'Flood' whose trigger is 'dam bursts upstream' and which has 'flood barrier' as a control. Its consequences include 'loss of life' and also 'house floods'. But the latter is a trigger for flood risk from a Householder perspective. From the perspective of the Local Authority Solicitor the main risk event is 'Loss of life' for which 'Flood' is the trigger and 'Rapid emergency response' becomes a control rather than a mitigant.

This ability to decompose a risk problem into chains of interrelated events and variables should make risk analysis more meaningful, practical and coherent. The BN tells a story that makes sense. This is in stark contrast with the "risk equals probability times impact" approach where not one of the concepts has a clear unambiguous interpretation. Uncertainty is quantified and at any stage we can simply read off the current probability values associated with any event.

The causal approach can accommodate decision-making as well as measures of utility. It provides a visual and formal mechanism for recording and testing subjective probabilities. This is especially important for a risky event for which you do not have much or any relevant data.

4. Conclusions

We have addressed some of the core limitations of both a) the data-driven statistical approaches and b) risk registers, for effective risk management and assessment. We have demonstrated how these limitations are addressed by using BNs. The BN approach helps to identify, understand and quantify the complex interrelationships (underlying even seemingly simple situations) and can help us make sense of how risks emerge, are connected and how we might represent our control and mitigation of them. By thinking about the hypothetical causal relations between events we can investigate alternative explanations, weigh up the consequences of our actions and identify unintended or (un)desirable side effects.

Of course it takes effort to produce a sensible BN model:

- Special care has to be taken to identify cause and effect: in general, a significant correlation between two factors A and B (where, for example A is 'yellow teeth' and B is 'cancer') could be due to pure coincidence or a causal mechanism, such that:

    - A causes B

    - B causes A

    - Both A and B are caused by C (where in our example C might be 'smoking') or some other set of factors

  The difference between these possible mechanisms is crucial in interpreting the data, assessing the risks to the individual and society, and setting policy based on the analysis of these risks. In practice causal interpretation may collide with our personal view of the world and the prevailing ideology of the organisation and social group, of which we will be a part. Explanations consistent with the ideological viewpoint of the group may be deemed more worthy and valid than others irrespective of the evidence. Hence simplistic causal explanations (e.g. 'poverty' causes 'violence') are sometimes favoured by the media and reported unchallenged. This is especially so when the explanation fits the established ideology helping to reinforce ingrained beliefs. Picking apart over-simplistic causal claims and reconstructing them into a richer, more realistic causal model helps separate ideology from reality and determine whether the model explains reality. The richer model may then also help identify more realistic possible policy interventions.

- The states of variables need to be carefully defined and probabilities need to be assigned that reflect our best knowledge.

- It requires an analytical mindset to decompose the problem into "classes" of event and relationships that are granular enough to be meaningful, but not too detailed that they are overwhelming.

If we were omniscient we would have no need of probabilities; the fact that we are not gives rise to our need to model uncertainty at a level of detail that we can grasp, that is useful and which is accurate enough for the purpose required. This is why causal modelling is as much an art (but an art based on insight and analysis) as a science.

The time spent analysing risks must be balanced by the short term need to take action and the magnitude of the risks involved. Therefore, we must make judgements about how deeply we model some risks and how quickly we use this analysis to inform our actions.

5. References

[1] Agena 2010, http://www.agenarisk.com

[2] Casscells, W., Schoenberger, A. and Graboys, T. B. (1978). "Interpretation by physicians of clinical laboratory results." New England Journal of Medicine 299 999-1001.

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

[4] Druzdzel MK and van der Gaag LC, Building probabilistic networks: where do the numbers come from? IEEE Transactions on Knowledge and Data Engineering 12(4), 481-486 (2000).

[5] Fenton, N. E. and M. Neil (2007). Managing Risk in the Modern World: Bayesian Networks and the Applications, London Mathematical Society, Knowledge Transfer Report. 1. http://www.lms.ac.uk/activities/comp_sci_com/KTR/apps_bayesian_networks.pdf

[6] Fenton, N. and M. Neil (2010). "Comparing risks of alternative medical diagnosis using Bayesian arguments." Journal of Biomedical Informatics 43: 485-495.

[7] Fenton NE, Neil M, and Caballero JG, Using ranked nodes to model qualitative judgements in Bayesian Networks. IEEE Transactions on Knowledge and Data Engineering 19(10), 1420-1432 (2007).

[8] Gelman A, Carlin JB, Stern HS, and Rubin DB. Bayesian Data Analysis (2nd Edition), Chapman and Hall, 2004, pp. 209-302.

[9] Lauritzen SL and Spiegelhalter DJ, Local computations with probabilities on graphical structures and their application to expert systems (with discussion). Journal of the Royal Statistical Society Series 50(2), 157-224 (1988).

[10] Hossack, I. B., J. H. Pollard and B. Zehnwirth (1999). Introductory statistics with applications in general insurance, Cambridge University Press.

[11] Pearl, J. (1986). "Fusion, propagation, and structuring in belief networks." Artificial Intelligence 29(3): 241-288.

[12] Simpson, E. (2010). "Bayes at Bletchley Park." Significance 7(2): 76-80.

[13] Ziliak, S. T. and D. N. McCloskey (2008). The Cult of Statistical Significance, The University of Michigan Press.