Norman Fenton and Martin Neil

Measuring your Risks:

Numbers that would make sense to Bruce Willis and his crew

By destroying the meteor in the film Armageddon Bruce Willis saved the world. The probability of the meteor strike was so large, and the consequences so great, that nothing much else mattered except to try to prevent the strike. Combining the 'probability' and 'impact' of a risk in order to define its 'size' is standard practice. But in most cases it's irrational, and it certainly would not have explained to Bruce Willis and his crew why their mission made sense. To get rational measures of risk you need a causal model ('risk map') that links triggers, controls, events, mitigants and consequences. Once you do this measuring risk starts to make sense – and it's much easier.

1. The utility measure of risk is usually meaningless

Think about the 'risks' that might cause your next project to fail to complete properly. Some of your key staff might leave, you might run out of funds, or a piece of technology you were depending on might fail.

Whether deliberately or not, you will have *measured* such risks. The very act of listing and then prioritising risks, means that mentally at least you are making a decision about which risks are the *biggest*. A meteor strike that destroys your office (with all the staff in it) would probably not appear on your list. Although its impact is more devastating than any of the others, unlike in Armageddon the chances of it actually happening are so low that you discounted it.

What you probably did, at least informally, is what most standard texts on risk propose [4]. You decompose risks into two components:

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

You then define risk as the measure:

Risk = Probability x Impact

This 'utility' type measure of risk is quite useful for prioritising risks (the bigger the number the 'greater' the risk) but it's normally meaningless. More importantly, you normally cannot get the numbers you need to calculate it.

To see why let's go back to the Armageddon scenario as an example. According to the above approach to risk, we have a picture like the one in Figure 1.

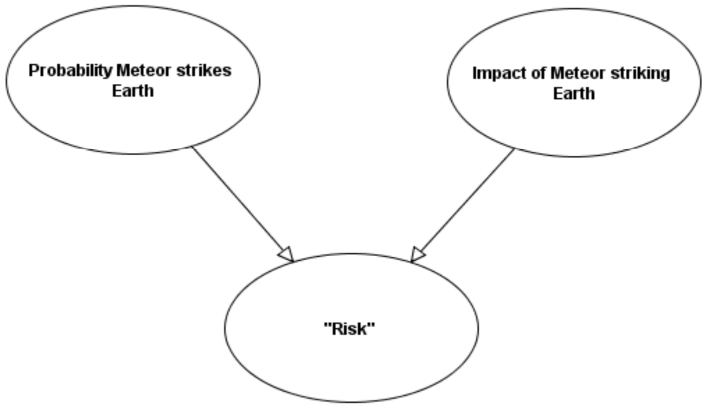

Figure 1 "Risk" of meteor strike

The problems with this are:

- *We cannot get the Probability number*. According to the NASA scientists in the film the meteor was on a direct collision course. Does that make the probability of it striking Earth equal to one? Clearly not, because if it was one 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 other 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, triggers and mitigants**. This is especially significant when there are, for example, controls that have never been used before (like destroying the meteor with a nuclear explosion).

- *We cannot get the Impact number*. 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 [5].

- *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.

2. Getting sensible risk measures: think of causal models (risk maps)

In [3] we explained how the rational way to think of risks was in terms of causal models (risk maps) with trigger events, control events, risk events, mitigant events and consequence events. For the meteor risk a relevant risk map is shown in Figure 2.

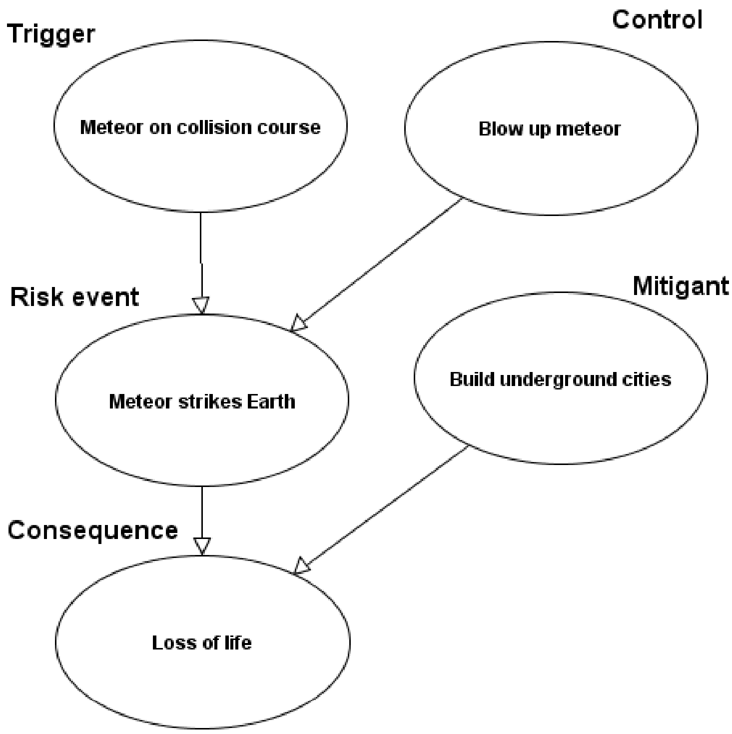

Figure 2 Meteor strike risk

A risk is therefore characterised by a set of uncertain events. Each of these events has a set of outcomes. For simplicity we assume that these events have two outcomes — *true* and *false* (in practice we can extend the outcomes to incorporate more states). So, for example, "Loss of life" here means 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*.

The sensible risk measures that we are proposing are simply the probabilities you get from running a risk map. Of course, before you can run a risk map you still have to provide some 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'.

To give you a feel of what you would need to do, in the risk map of Figure 2 the uncertain event "Meteor strikes Earth" still requires us to assign a so-called *conditional probability distribution*. 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 3.

Blow up meteor	False		True	
Meteor on collision course with Earth	False	True	False	True
False	1.0	0.0	1.0	0.8
True	0.0	1.0	0.0	0.2

Figure 3 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 blown up, and 0.2, if it is blown up. In filling in 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 risk map for which we *do* need to assign unconditional probability values. These are the nodes that have no parents; it makes sense to get unconditional probabilities for these because, by definition, they are not conditioned on anything. Such nodes can generally be only triggers, controls or mitigants. An example, based on dialogue from the film, is shown in Figure 4.

False	0.05
True	0.95

Figure 4 Probability table for “Meteor on collision course with Earth”

We are not suggesting that assigning the probability tables in a risk map is always easy. You will generally require expert judgement or data to do it properly (AgenaRisk provides a wealth of tools to make the task as easy as possible). What is important is that it is *easier* than the classic alternative. At worse, when you have no data, purely subjective values can be supplied.

The wonderful thing about risk maps is that once you have supplied the prior probability values a Bayesian inference engine (such as the one in AgenaRisk) will run the model and generate all the measures of risk that you need. For example, when you run the model using only the prior probabilities the model (as shown in Figure 5) computes the probability of the meteor striking Earth as just under 0.8 and the probability of loss of life (meaning at least 80% of the world population) is about 0.77.

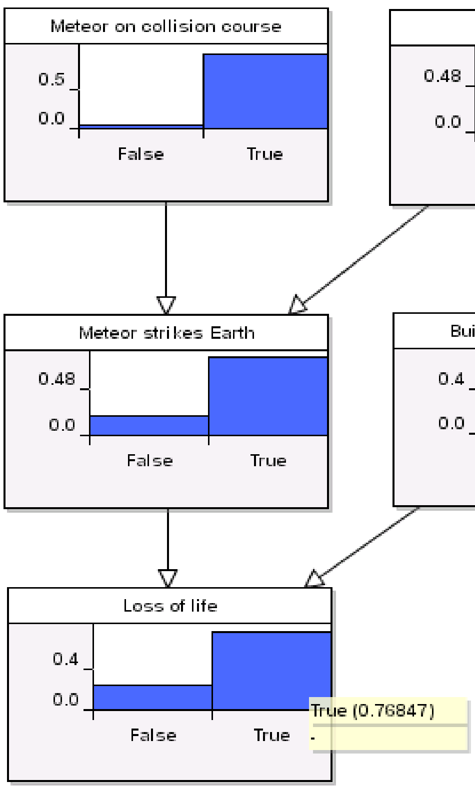

Figure 5 Initial risk of meteor strike

In terms of the difference that Bruce Willis and his crew could make we run two scenarios: One where the meteor is blown up and one where it is not. The results of both scenarios are shown in Figure 6.

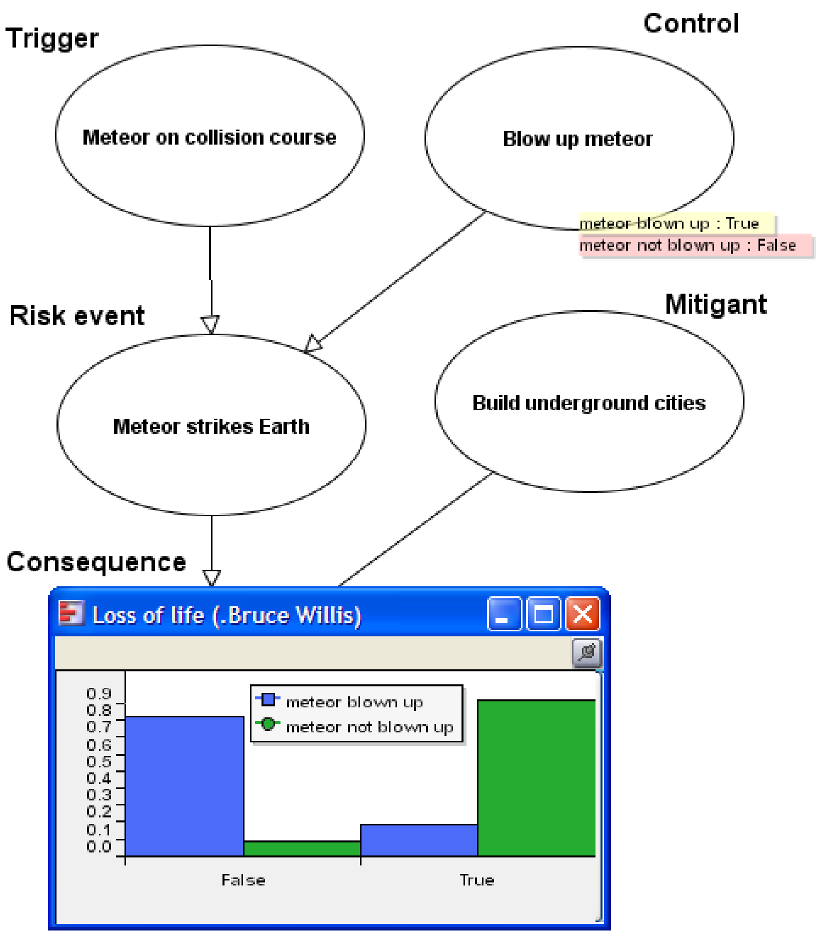

Figure 6 The potential difference made by Bruce Willis and crew

Reading off the values for the probability of "loss of life" being false we find that we jump from 0.085 (when the meteor is not blown up) to 0.82 (meteor blown up). This near tenfold increase in the probability of saving the world clearly explains why it merited an attempt.

3. Benefits

The causal, risk map approach satisfies minimalist requirements described by Chapman and Ward in [4] where they recommend that any approach to risk quantification:

> "should be so easy to use that the usual resistance to appropriate quantification based on lack of data and lack of comfort with subjective probabilities is overcome".

Moreover, the approach ensures that:

- Every aspect of risk measurement is meaningful in the context – the risk map tells a story that makes sense. This is in stark contrast with the "risk = probability x impact" approach where not one of the concepts has a clear unambiguous interpretation.

- Every aspect of uncertainty is fully quantified since at any stage we can simply read off the current probability values associated with any event.

- It provides a visual and formal mechanism for recording and testing subjective probabilities. This is especially important for a risky event where you do not have much or any relevant data about (this is after all Humankind's first mission to land on a meteorite).

Although the approach does NOT explicitly provide an overall risk "score" and prioritisation these *can* be grafted on in ways that are much more meaningful and rigorous. For example, we could:

- Simply read off the marginal probability values for each risk event given your current state of knowledge. This will rank the risks in order of probability of occurrence. (this tells you which are most likely to happen given your state of knowledge of controls and triggers);

- Set the value of each risk event in turn to be true and read off the resulting probability values of appropriate consequence nodes. This will provide the probability of the consequence given that each individual risk definitively occurs. The risk prioritisation can then be based on the consequence probability values.

- Introduce 'dashboard' type features by identifying key consequences, setting thresholds below with they should not fall and then building traffic light components based around that.

Above all else the approach explains why Bruce Willis's mission really did make sense.

4. References

[1] Agena Ltd, www.agenarisk.com, 2006

[2] Chapman C and Ward S, 'Estimation and evaluation of uncertainty: a minimalist first pass approach', International Journal of Project Management, 18, 369-383, 2000

[3] Fenton NE and Neil M, 'Visualising Risk', www.agenarisk.com, 2006

[4] Lowrence W, The Nature of Risk, in Societal risk assessment: how safe is safe enough?, R Schwing & W Albers (eds), 5-14, New York: Plenum, 1980

[5] Roberts FS, 'Measurement Theory with Applications to Decision Making, Utility and the Social Sciences', Addison Wesley, 1979