What is Risk Anyway?

Michael Barwise

In addition to Donald Rumsfeld's much quoted 'known unknowns' and 'unknown unknowns' there's another class of misapprehension that he failed to recognise - things you're convinced of that happen to be wrong (in his case, WMD). For many of us in corporate information assurance, the real nature of risk is in this category. We all perform what we believe to be 'risk assessments' but are they really any good?

Fortunately (or for me as a methodologist - unfortunately), information breaches are quite rare so the majority of our risk assessments never get tested. If they did, I suspect they would exhibit a long-term success rate approaching 50 per cent. That makes them about as useful as tossing a coin - not because we're stupid, but because we're simultaneously largely uninformed, and widely misinformed by 'experts', about the true nature of risk. So what is Risk? Many supposedly authoritative sources refer to events or outcomes as 'risks', whereas risk is actually an *attribute* of an event: a measure of its probable consequence. Nevertheless 95 per cent of risk professionals responding to a survey in 2001[1] agreed that 'a risk' is an event.

Having defined risk, let's consider how to quantify it. We might hope to be guided here by standards, but the ISO Guide 73[2] definition that has influenced almost all other risk-related standards is '*effect of uncertainty on objectives*': a definition at once irrefutable and effectively useless, as it's entirely abstract. Many attempts to create operationally functional definitions have been made, ranging from the elementary '*risk=likelihood x consequence*' to quite complex combinations of '*vulnerability*', '*threat*', '*opportunity*', '*impact*' among other terms, multiplied and summed in various ways. However I question whether many of these ostensible mathematical relationships are valid, and whether their parameters are specified in ways that allow mathematical operators to be used at all. How do you multiply (or add) '*wooden shed*', '*small boy with box of matches*', '*pyromaniac tendency*', and '*value of contents*' to arrive at '*loss of tools*'?

But this is not my only concern. Finding evidence-based or 'quantitative' risk decision-making rather hard work, risk practitioners have mostly resorted to 'qualitative' methods (a.k.a. guesswork), resulting in a drastic loss of both accuracy and repeatability. Such sloppy thinking encourages the use of crude risk rankings - '*high*', '*medium*', '*low*' - that make it impossible to distinguish with confidence anything but extreme differences in risk. Furthermore, cross referencing 'medium impact' and 'medium likelihood' may yield 'medium risk' or 'high risk', depending solely on my personal preconceptions when I designed the corporate Risk Matrix. There's no demonstrably valid (or even accepted arbitrary) axiom to guide us, so someone else in the same organisation (and even the same role) might create a risk matrix quite different from mine, delivering different answers. These failings combine to cause cultural dynamics ('office politics' and personal attitudes to taking gambles) to swamp objectivity. Unwillingness both to bring bad news and to stick one's neck out frequently leads to most risks being ranked 'medium', so we don't make much progress in prioritising our risk treatment, even supposing our criteria were trustworthy in the first place. Considering the large number of entries in a realistic corporate risk register, granularity is essential - there's no real hope of prioritising the treatment of 649 'medium risks'. And ultimately, that's what corporate risk management is about: not arriving at absolute values of risk as an intellectual exercise, but working out the optimum priorities when allocating a limited protection budget.

These failings (ill-defined formulae, low resolution poorly quantified 'risk scales' and uncontrolled or biased guesswork) contribute significantly to what are often essentially meaningless risk decisions. However their malign influences are usually dwarfed by an overriding conceptual error. Corporate risk assessments commonly assume a single cause leads to a single outcome with a single (if vaguely expressed) likelihood and consequence. Unfortunately, the real world ain't quite like that. This has frequently been ignored, even in life-critical arenas. For example, a coincidence of one 'medium risk' and two 'low risk' independent events is unlikely to be considered a high risk, even supposing we know what low, medium and high mean in the first place. But those were the assumed risks of the three most significant causal factors of the NASA Challenger accident[3] (reduced rocket segment 'O' ring resilience at low temperatures, distorted re-usable rocket segments, leaks in segment joint insulating putty).

Most adverse incidents (and indeed many business opportunities) result from the coincidence of

multiple independent events. Some of these events, individually or in concert, may trigger dependent intermediate events, forming chains of causality. The ultimate result may sometimes be a single outcome, or there may be multiple alternative or simultaneous outcomes. Each event has a likelihood of occurring at the required point in the mesh of causality, depending not only on its intrinsic properties but also on the properties of any other events that contribute to it. Furthermore, some events are binary (they happen or they don't) and some have multiple discrete effects, but the effect of many events varies over a range, and not necessarily in an intuitively determinable way. The probability of each possible outcome is a function of the aggregate of probabilities of all the events in all the contributory chains of causality, so clearly there is usually a range of possible outcomes and consequences. Such ranges of possibility are '*probability distributions*' . Although they can be highly informative, they are almost universally ignored in the sphere of information risk management, partly because they are hard to define for some of the events we deal with, but mainly because most practitioners don't even know they exist.

That said, events driven by human decisions - including most cyber attacks - tend not to have constant probability distributions over time. This is why trying to deduce future exposure from past information breaches is so uncertain. Although Fault Tree analysis is a widely adopted technique for finding causes after the fact, because we often don't know the probability distributions of the contributory factors at the time of the incident it's often impossible to recreate its causal matrix with confidence. However, before the fact it's possible to use the reverse of fault tree analysis - *consequence analysis* - to map the possible outcomes of coincidences of events. This is not strictly 'risk assessment' unless it's possible to assign probability distributions to the events, but it's nevertheless a powerful tool for identifying possible adverse outcomes that might otherwise escape identification due to the apparent insignificance of their causal factors when considered individually in isolation. Such outcomes can then be pre-empted by preventing as many of the causal factors as practicable from acting.

In summary, to make reliable risk judgements we must: understand the basic principles of probability, including recognising that statistics don't describe individual events and knowing when probability theory can't be usefully applied; completely understand the process or system risk being assessed; have no vested interest in the outcome of the assessment; and apply a repeatable standard process that demonstrably yields results that consistently accord with reality. And sometimes likelihood's contribution to business exposure must be ignored, particularly when a potential outcome could be catastrophic.

[1] Hillson, D. What is 'Risk'? Results from a Survey Exploring Definitions. http://www.risk-doctor.com/pdf-files/def0202.pdf

[2] ISO Risk management — Vocabulary, 2009

[3] Report of the Presidential Commission on the Space Shuttle Challenger Accident, Chapter IV: The Cause of the Accident http://history.nasa.gov/rogersrep/v1ch4.htm (section 70: Findings)

Originally appeared in ITNow, Summer 2014