Information Security Risk Revisited

Michael Barwise

"*The significant problems we face cannot be solved at the same level of thinking we were at when we created them.*"[1]

Introduction

When conversing with risk practitioners, I am often reminded of the song from Shakespeare's Twelfth Night "*what is love? 'tis not hereafter... What's to come is still unsure...*".[2] For indeed risk as we currently deal with it seems to me rather like love: we all believe we are innately equipped to recognise it; there is little agreement on what it really is, and embarrassing mistakes are remarkably common. In order to minimise failures, I believe we need to re-appraise the ways we define risk, the ways we quantify it and the quality of our judgement processes. Currently, most of us discuss risk without making reference to any fundamental definition. It is taken for granted that "everyone knows" what we are talking about. We use natural language classifications and rely on subjective experiential judgement when quantifying risks. Yet we believe our risk judgements can be trusted. I submit that this is not the case, and that the information security profession suffers from the adverse effects of a linguistic determinism that has over-simplified our concent of risk to a point where we cannot properly discuss it, let alone assess and manage real-world risks in an objective and repeatable manner. The fundamental problems we face are lack of clarity about what risk really is, lack of understanding of the statistical properties of the components of risk, and failure to appreciate the influence of uncertainties on the judgement process. To improve our risk judgement, we need not only to take on board new information, but most importantly to change the manner in which we consider it. To quote Owen Barfield "*The way we think is at least as important as what we think*"[3]

Definitions of Risk

There is a profound lack of agreement about what risk is. Having examined several hundred published definitions of risk, including an interesting collection by Hall,[4] I find that about twenty percent are purely self-referential (e.g. "*Legal Risk: the risk of loss to an institution*"[5]) and so inherently fail to advance our understanding. Sadly, the overwhelming majority of the remaining eighty percent, including "*Risk is a combination of likelihood and consequences*",[6] "*exposure to gain and loss*",[7] "*a risk is a potential problem*"[8] and even "*a 'security risk' [is] any driver who has been convicted of… crimes*"[9] serve little better to clarify the fundamentals. A vague general view of a risk emerges as an unpleasant event that might happen. Indeed the weather forecasters often warn us of "a risk of rain", but never mention "a risk of sunshine". There is of course always a risk of rain: the important question is what the probability of rain is today.

There is a similar lack of clarity about the parameters of risk. These are usually described exclusively in natural language, rendering them ambiguous. As a result it is unclear how to derive from them finite solutions to real-world risk problems. Although numerous quasi-mathematical "risk models" have been invented, this lack of clarity concerning the definition and dimensionality of the variables used undermines their practical applicability. For example, Mohindra and Larocque[10] offer the equation "*Risk = Frequency x Consequence*" without defining either of these terms in any concrete manner that would allow real-world values to be plugged into the equation. Similarly, Chicken & Posner[11] offer "*Risk = Hazard x Exposure*", defining hazard as "*.. the way in which a thing or situation can cause harm,"* and exposure as "*.. the extent to which the likely recipient of the harm can be influenced by the hazard*", again failing to define their parameters in mathematical terms. If such a definition is to be of

[1] Albert Einstein 1879-1955

[2] W. Shakespeare. Twelfth Night, Act II, scene III.

[3] Owen Barfield 1898-1997

[4] Hall DC. 2003. What is "Risk"?. Vol II. http://www.techriskmgt.com/RMWGdownloads/ What%20Is%20Risk%20Volume%20II.doc

[5] International Bar Association.

[6] Mohindra S *and* Larocque GR. 2002. Application of Risk Analysis to Security: An Introduction. TIAX LCC, Cambridge, Mass.

[7] Valuesim (http://www.valuesim.no) *in* Hall DC *op.cit.*

[8] Arizona State University in Hall DC *op.cit.*

[9] Texas Motor Transport Association *in* Hall DC *op.cit.*

[10] Mohindra S *and* Larocque GR. *op.cit.*

[11] Chicken JC and Posner T. 1998. The Philosophy of Risk.. Thomas Telford. ISBN 0-7277-2666-8.

practical use, we must be told how to measure "*.. the way in which a thing ... can cause harm*" and "*.. the extent to which [something] can be influenced...*" in a consistent manner to derive values that can be multiplied together. What are the units of hazard, exposure, frequency and consequence? In reality, these are potentially complex issues requiring not only considerable conceptual understanding of statistics (which most information risk practitioners in the front line do not possess), but also of the events and processes under discussion (about which we frequently know much less than we would like). The impetus to simplify is therefore strong, and in the absence of a sound understanding of statistical theory, it is difficult to know when to stop. For example, to simplify practical risk decision-making, ordinal scales or rankings ("high", "medium", "low") are widely used in risk equations to substitute for values of parameters. But there is a general failure to recognise that, from the mathematical standpoint, quantisation should be the last, not the first, step: that, for the results to be useful in any given context, allocation of measurements to an ordinal scale is an outcome, not a starting point. Furthermore, if the definition of parameters remains vague or ambiguous, allocation to rankings on such a scale cannot be rigorous. But if it relies to any significant extent on the semantics of natural language, this lack of rigour may go unnoticed until subsequent events prove that risks were wrongly evaluated.

Such over-simplification may conceal fundamental defects in risk evaluation, but it also, and more importantly, can conceal inadequacies in the definition itself. However, semantic simplification of complex issues, as seen in the prevalence of self-referential "definitions" of risk (which actually elucidate nothing), is an innate human tendency that is well recognised as feeding back on, and significantly biasing, patterns of thought. As pointed out by Whorf, "*We cut nature up, organize it into concepts, and ascribe significances as we do, largely because we are parties to an agreement to organize it in this way--an agreement that holds throughout our speech community and is codified in the patterns of our language.*"[12] I suggest that, due to a lack of understanding of statistics and probability theory, coupled with the ubiquitous imprecise use of the word "risk" in common parlance, we are preconditioned to think we can recognise and understand risks, and not to question our assessment methods. Because we are experientially equipped to make reasonably reliable judgements about when it is safe to cross a busy road, we believe we can handle much more complicated risk decisions using the same methods. Indeed, it may not even be apparent that the additional complexity exists. This then is the problem we have to address.

Risk Judgement

There is generally held assumption that frequency, consequence, hazard, exposure *et al* are finitely quantifiable absolutes. This is not the case. If it were, we would not be discussing risk at all. There will always be an element of uncertainty about future events over which we exercise imperfect control. Some emerging risk definitions do indeed suggest this has been recognised. For example the Institute of Internal Auditors[13] has defined risk as "*a measure of uncertainty*", and the UK OGC (in its risk practitioners' guide)[14] as "*uncertainty of outcome*". But even these definitions assume our judgement is perfect. Sadly, it cannot be perfect. How good it is depends largely on how we approach the process of decision-making, but clearly any risk decision we make will be influenced by both factors. As a basic example, consider a horse race. It is impossible to absolutely certain of the winner in advance regardless how much research you do, as there are uncontrollable factors continuing to operate after the moment of decision-making (placing your bet). These continue to influence the outcome right up to the end of the race, so there is an element of irreducible uncertainty as to which horse will win. This is *systemic uncertainty*: the uncertainty intrinsic to the system. However, if you have been unable to the form, your judgement will additionally suffer from *epistemic uncertainty*: uncertainty as to the quality of your judgement. Note that this does not mean "how convinced you are that you are right" but specifically the probability of your judgement being objectively wrong.

Particularly in cases where information is sparse or where the statistics of the situation are poorly understood or unstable, epistemic uncertainty can be a significant parameter, and can even dominate the equation. This has been conclusively proved by numerous authoritative studies in spheres other

[12] Whorf BL. 1940. *in* Carroll, J.B. (ed). 1956. Language, Thought and Reality: Selected Writings of Benjamin Lee Whorf. MIT press: Cambridge MA. pp213-214.

[13] http://www.theiia.org/iia/index.cfm?doc_id=1605

[14] Office of Government Ccommerce. 2002. Management of Risk: Guidance for Practitioners. The Stationery Office. ISBN 0-1133-0909-0.

than information technologies,[15] but is still generally ignored at the front line of information risk assessment. As a result, many “risk assessments” are in reality based on little more than untested guesswork: convincing enough for “compliance” audits but insufficiently accurate to support the development of cost-effective countermeasures. I believe the situation could be significantly improved by bringing to bear established methods of reducing uncertainties of judgement that have been applied to good effect in other spheres. The first step is, however, to recognise the need. In the words of Oliver Cromwell “*I beseech you ... think it possible you may be mistaken*”[16]

Circulated as a white paper, July 2005

[15] for an excellent overview, see Morgan MG and Henrion M. 1998. Uncertainty. Cambridge University Press.

[16] Oliver Cromwell 1650, in a letter to the general assembly of the Church of Scotland.