EXPLORING SUBJECTIVITY IN HAZARD ANALYSIS[1]

Felix Redmill
Redmill Consultancy
Email: Felix.Redmill@ncl.ac.uk

INTRODUCTION

In an earlier paper (Redmill 2002) risk analysis was defined as comprising four stages - scope definition, hazard identification, hazard analysis, and risk assessment. In the first stage, the scope and terms of reference of the analysis are defined. In the second, the hazards that could lead to breaches of safety are identified. In the third, the risks associated with the hazards are determined, and in the fourth the tolerability of each risk is assessed against predetermined criteria.

In considering the subjectivity in risk analysis, the previous paper particularly addressed the first and second stages and briefly discussed the fourth. This paper examines the subjectivity inherent in the third, hazard analysis, stage.

Once the scope and terms of reference of a risk analysis have been defined, the second, hazard identification, stage lays the technical foundation of the analysis. It cannot be guaranteed that all possible hazards are uncovered, or that all hazards that will be discovered are identified, as it would be extraordinary if further hazards were not revealed later. But the coverage of later analysis is constrained by the thoroughness of the identification process: hazards not identified are neither analysed nor mitigated.

Given that risk is taken to be a function of probability and consequence, hazard analysis, the third stage of risk analysis, involves determining these two parameters, which may be done quantitatively or qualitatively, depending on the information available and the confidence that can be placed in numeric values. One method is to carry out a bottom-up analysis, starting with each hazard and working forwards towards its system-level consequences. The initial effect of a hazard (e.g. the failure of a component or person) may be local, but the hazardous effect of interest is almost always at the system level - i.e. at the boundary between the system and 'the rest of the world'. A second method of hazard analysis is to take a top-down approach, commencing with the top hazardous events and working backwards towards ultimate causes, creating fault trees in which successive causal events are identified.

The two approaches are complementary and neither is likely to offer a complete analysis. Because their results are usually inconsistent in many respects, they should be compared and integrated. Yet, far from comparing them, some modern computer-based tools automatically derive fault trees from the results of a bottom-up technique such as FMEA (fault modes and effects analysis). Naturally it can then be claimed that the fault tree is consistent with respect to the model produced from

---

[1] Published in Engineering Management Journal (IEE), Vol. 12, No. 3, June 2002

the FMEA. But, being the result of human judgement, that model is almost certainly neither complete nor wholly correct. Moreover, the opportunity for cross-checking between bottom-up and top-down methods is lost, and there is likely to be misplaced confidence in the correctness of the fault trees.

The results of a risk analysis depend on the techniques employed, the ways in which they are used, and the consistency with which they are used with respect to each other - all of these factors being subjects of human discretion.

The more sophisticated hazard identification techniques, e.g. hazard and operability studies (HAZOP) and FMEA, also include bottom-up hazard analysis. Discussion of these, and the subjectivity implicit in their use, has already been provided (Redmill 2002) and will not be repeated here. In examining subjectivity in hazard analysis, this paper first considers the two aims of the process - to determine consequences and likelihood. It examines the ways in which subjectivity affects the numbers arrived at in both cases. Then it considers the use of the most usual top-down method, fault tree analysis (FTA), in arriving at them. Finally, there is a discussion of the findings.

CONSEQUENCES

At first glance, evaluating consequences may seem objective, but what we evaluate depends on where we look, and this is determined by a number of decisions.

First, the consequences to be evaluated depend on which event in a chain of events is considered to be 'final' or of interest in the circumstances. For example, in transport a human or component failure might lead to the loss of control of a vehicle, which could lead to an accident, which in turn could lead to a loss of life. Each of these might be of interest as a 'final event', depending on the purpose of the risk analysis, and each carries different assumptions and a different probability of occurrence. Railtrack's safety management manual (Railtrack 2000) takes loss of life to be of interest, but the motor industry's guidelines (MISRA 1994) focus on controllability (or loss of it) of the vehicle.

It must also be decided, at the scope-definition stage of the analysis, whether estimations should be based on the worst possible consequence, the worst credible, or the most likely, and the risk values are influenced by the choice. Further, each scenario is not clearly defined and waiting to be 'measured', but is a potential outcome whose parameters must be subjectively defined - perhaps in line with the goals or mind-set of a particular industry sector.

Then, the values of the possible losses need to be identified and estimated, and they may be distorted in a number of ways. For example, some costs, such as those of damage limitation and of 'cleaning up' after an incident, are frequently omitted. Further, when there is no previous experience of the hazardous event, it is easy to over- or under-estimate the consequences. For example, in the UK the losses attributed to a 'hundred-year' flood are now found to be a great deal higher than previously assumed. Then there is the potential to induce distortions by deliberate adjustments of costs or benefits.

Thus, there is always subjectivity in the estimation of consequences, and more so

when there is little or no experience of the hazardous event. Not only is there a degree of uncertainty about a potential future event, there is also error, inaccuracy, and the use of discretion and judgement in the description and valuation of what might occur.

LIKELIHOOD

Likelihood may be determined quantitatively (as a probability) or qualitatively. Qualitative analysis is by definition approximate, but quantitative analysis is often assumed to be wholly objective. Yet there is considerable subjectivity in the analysis process. In spite of the appearance of accuracy, quantitative analysis is subject to assumptions that are not always made explicit. One is that of randomness. Yet, in many modern systems - particularly those in which software and humans are involved - the assumption is invalid. Unlike mechanical and electromechanical components, software does not wear out. Its failures are not random but systematic and would be repeated if the triggering circumstances recurred, so even a history of use and failure is not a firm basis for prediction. Forecasts of the rate of software failure should only be made with great care.

Similarly, the behaviour of humans cannot be assumed to be random. Yet many models for human reliability assessment (HRA) are probabilistic. However, although 'probabilities' are derived, the approach taken in most cases is based principally on human judgement. The results are at best reasonable approximations and at worst wild guesses, but always they include considerable subjectivity.

Even for hardware, whose failure may indeed be random, it is questionable to what extent the derived probabilities are truly representative. Some are of incredible magnitude and, though stated to several decimal places and based on elaborate calculations, may contain huge errors because of the omission of failure modes from the model on which the calculations are based. Results then are representative of the model but not of reality, as was highlighted by the failure of the British nuclear submarine *Tireless* in 2000.

The problem was a crack in a cooling-system pipe. Such a crack had not been considered in the probabilistic analysis of submarine failure, and (perhaps consequently) the pipe was never checked during maintenance. By implication, the occurrence of such a crack was considered to be incredible. Yet, when the crack occurred on *Tireless*, and checks of the other submarines in the fleet were made, seven of the twelve were found to have indications of similar cracking and the other five were not wholly exonerated (BBC 2000). The calculated probability of failure was out by several orders of magnitude - with huge detrimental implications for national defence. Subjective assumptions and omissions through human judgement or negligence can have enormous implications on the accuracy and relevance of probabilistic calculations. Failure or accident may occur for reasons not considered in the analysis. Precision is not the same as accuracy and should not be assumed to imply it.

Another issue is statistical inference - the reliance on historic data for the prediction of future events. Some analyses rely on inadequate or inappropriate data and are therefore wrong. But even when the data are statistically valid, it is important for the conditions in which the data are to be applied to be the same as (or adequately

similar to) those in which they were collected. Further, historical results often rest on crucial conditions that are unnoticed, unrecorded, or unrepeatable, so using them as the basis of prediction means applying subjective judgement. Suppose, for example, that a component of risk in system operation is operator error. Suppose that a company plan to reduce its operating budget involves the recruitment of less qualified staff and a cut-back on training. The error frequency prior to the cut-back would be expected to be lower than after it. It would not be a valid predictor of future performance because of the change in conditions. Past history is an unreliable guide to the future if consistency of conditions is not guaranteed.

Even when the historic data is extensive, and the conditions remain constant, the past frequency only tends to become an accurate predictor (in theory) as the time of observation tends towards infinity. This is the law of large numbers. Now, what may be taken as a reasonable approximation to infinity depends on the application, and in some cases may be quite small, but care is required in making assumptions about the predictive value of historic data. Many risk-takers have lost their shirts in casinos because they have implicitly assumed that the law of large numbers applies to small numbers.

When historic data are not available, as in the case of new technologies and products, mathematical models are often devised for assessing probabilities, and these carry assumptions. For example, if a new drug is found to (or not to) induce cancer in mice but cannot be tested in large doses on humans for ethical reasons, how valid is a projection of its effect on mice to its effect on humans? Assumptions must be made; and it is usual for the judgements of experts to differ, for they depend on the assumptions made about the relationship between the observed effect and the administered dose.

At a recent conference on risk, one paper presented a mathematical model, based on probabilistic equations, for determining the likelihood of certain hazardous events. At the end of the talk, a delegate referred to the presenter's statement that in some cases there were very sparse data for input into the equations, and he asked what was done in such cases. The presenter dismissed the question. 'That is not a problem,' he replied. 'When there is sparse data we just employ an expert to provide an opinion.' The mathematician neglected to wonder what the expert based his opinion on if historic accident data did not exist. A computing acronym, GIGO (garbage in, garbage out) is also appropriate to mathematical risk models, but the results of risk analyses are often taken to be accurate and the assumptions and inaccuracies in their derivation unrecorded and forgotten.

The foregoing discussion has concerned quantitative risk analysis. In qualitative analysis techniques, subjectivity is an integral and obvious part of the process.

THE USE OF FAULT TREES

The previous two sections have shown that there is considerable subjectivity in the estimation of the two components of risk, consequence and likelihood. But what about the techniques used in combining them? In the previous paper (Redmill 2002), the subjectivity in bottom-up techniques, such as hazard and operability studies (HAZOP) was highlighted. In this section, fault tree analysis (FTA), a top-down technique is examined.

A fault tree is a cause and effect network. It starts with a final hazardous event and works backward logically to ultimate causes. The ways in which events combine to cause higher-level events are represented in the tree by AND and OR logical units, or gates (see Figure 1). If the occurrence of any of a number of events would cause a higher-level event, the causal events are combined by an OR gate; if a given result would require the occurrence of two or more events, these are combined by an AND gate. When probabilities are attached to the tree, their combinations can be calculated all the way up to the top level. Those combined by an OR gate are added and those combined by an AND gate multiplied.

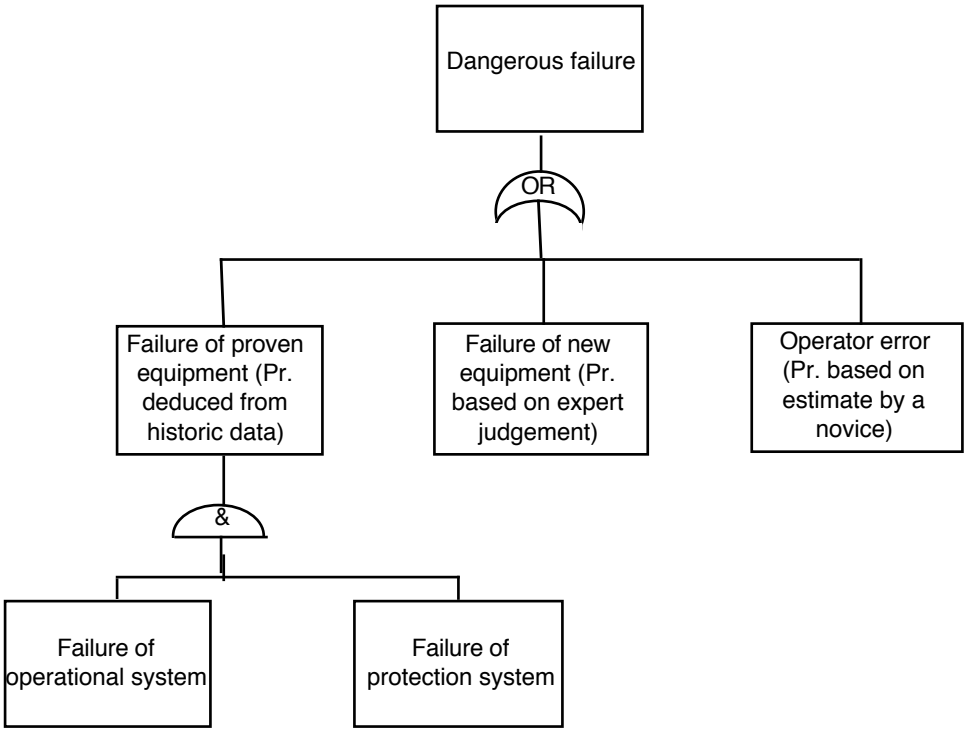

Figure 1: A simple fault tree

Fault tree analysis was developed in the field of reliability theory, where, typically, systems comprise known and identifiable subsystems and components. In such cases, a fault tree may be represent a system directly. Traditionally, too, components were predominantly mechanical or electromechanical, and their fault histories allowed probabilistic determination of the likelihood of the top event. FTA is therefore often thought of as an objective technique.

Attaching probabilities to fault trees assumes not only randomness (already discussed above) of individual events (e.g. component failures) but also independence of events from common causes. But independence may not apply, often for unrecognised reasons such as the derivation of power from a single supply or the control of several componentss by a single operator.

The extension of fault tree analysis to situations of uncertainty, such as policy decisions and safety scenarios, means that the trees are subjectively constructed rather than modelled to reflect directly the combinations of components in a system.

This introduces judgement as to what should be included in the tree and the possibility of omissions because of ignorance or error.

Fischhoff, Slovic and Lichtenstein (1978) have shown that the construction and use of fault trees are subject to variability. They point out that during construction such questions arise as, Which faults should be identified separately and which should be lumped under 'others'? Which items should be grouped together? What sort of graphic display should be used? What level of detail is most appropriate? Answers depend on a number of factors, including the purpose of the analysis and how much of an effect each branch of the tree or each component is thought to have. Thus, the construction of a fault tree is subjective, and a tree prepared for a given purpose by one person is unlikely to be reproduced by another under the same circumstances.

Omissions of relevant pathways in a fault tree are possible because of ignorance, poor memory, and lack of imagination, among other causes, and such omissions can lead to understatement of the relevant probabilities. Fischhoff, Slovic and Lichtenstein concluded that in the creation of a fault tree humans are likely to be biased in favour of information readily available to them (the availability bias), and that when omissions do occur people are, in the main, insensitive to them. In their study, this was found to be the case not only in tests of groups of college students but also of groups of experts (car mechanics). Such insensitivity to omission occurred in the case of the submarine *Tireless*, discussed above.

Fischhoff, Slovic and Lichtenstein also found that the perceived importance of a particular branch of a fault tree was increased if it was represented in pieces (i.e. as two separate component branches). Thus, the probabilities estimated by experts were dependent on the construction of the tree, which in turn was subject to human frailties and biases.

Then, in the absence of reliable historic data, at least some probabilities are likely to be derived from sources of 'low pedigree'. Funtowicz and Ravetz (1990) show that a system's probability of failure may be dependent on the combination of items of information provided by various persons. One source, they say, may be historic failure data, from which a reliability estimate may be derived (see the second layer of Figure 1). They refer to this source as being of high pedigree. But for the failure probability of a new piece of equipment, the information may be an estimate by acknowledged expert, and they refer to this as being of medium pedigree. Another part of the tree might require an assessment of the expected reliability of staff during maintenance, and this might be made by (say) a recent graduate who, Funtowicz and Ravetz suggest, might be a source of low pedigree.

Even high-pedigree data sources may lead to false probabilities. As shown above, crucial conditions during data collection and the assumptions made in deriving results are often not recorded, and the conditions under which the derived probabilities are used predictively may be very different from those under which the data were collected. But confidence levels are not commonly assigned to probabilities, so there is no recognition by fault-tree users of when they are low. In any case, how can confidence be derived in very low probabilities? In most instances adequate statistical data are not available for the assessment of rare events, and it could take years to discover if the assumptions on which estimates are based are valid, or even reasonable.

But if some probability estimates are inaccurate, what is their significance? In a simple example, Freudenburg (1992) shows that in some cases omissions can make a huge difference to the results. Suppose, Freudenburg suggests, that in drawing up a fault tree, analysts arrived at a failure probability of 10^{-6}, having identified all but two risk factors - one of which made the system more safe and the other less safe. And suppose that the system would operate at the 10^{-6} risk level for 80% of the time, but the 'real' risk would in fact include operation at 10^{-3} for 10% of the time and 10^{-9} for the other 10%. It might at first be assumed that the two errors would cancel each other out in the risk calculation, but this is not the case. The actual calculation produces a result of: $(0.1 \times 10^{-9}) + (0.8 \times 10^{-6}) + (0.1 \times 10^{-3}) = 0.0001008001$, which is slightly greater than 10^{-4}. Thus, the omitted higher risk is not cancelled by the omitted lower risk but dominates the result, even though it exists for only 10% of the time. Its omission leads to a distorted belief in the safety of the system.

Thus, not only are subjective omissions and inaccuracies almost inevitable, but they can also be of great significance to the result of a risk analysis. In the case of the submarine *Tireless*, not only were the probabilistic calculations inaccurate by many orders of magnitude but also the consequences of this (on the defence of the nation) were potentially catastrophic.

DISCUSSION

All aspects of risk-analysis, from planning to the interpretation of results, rely on subjectivity. This paper shows that in the third stage of the process, hazard analysis, the derivation of values for both consequence and probability is based on numerous judgmental decisions. Moreover, the method most used for determining cause and effect and for combining probabilities, fault tree analysis, was seen to be dependent on human factors for both its construction and its use.

While the language of the paper suggested that quantitative analysis was being addressed, it should be added that qualitative analysis, extensively used on many of today's systems, is by definition based on human judgement.

It was seen that inputs to FTA may arrive from disparate sources, of varying pedigree and trustworthiness, so the error in the estimation of probabilities can be great, without there being any human sensitivity to it. It would therefore be useful to introduce an activity in risk analysis in which information and its source and pedigree were reviewed. There should also be a formal requirement to place a confidence level on information and to make a decision on whether there is a need to seek further information from a source of higher pedigree.

It would also be useful to revise the risk-analysis syllabus to cover the ways in which subjectivity is introduced and the effects that it has, and to make the process's assumptions more explicit to analysts. Thus, analysts would be taught to understand not only the mathematical assumptions but also their own human biases. There would then be an increased chance that subjectivity would be considered, and partly neutralised, during the analysis and management of risk, and that evidence would exist for placing confidence figures on results.

While this paper points to the subjectivity in hazard analysis and invites analysts to understand it better, it should also be emphasised that human creativity and

decision-making are major strengths of the process. Transparency and a methodical approach are others. In spite of the criticism that it receives, risk analysis and its techniques serve industry well and could beneficially be transferred from science and engineering to other domains where there is a recognised need for a better understanding of risks.

REFERENCES

BBC (2000). *File On Four*. BBC Radio 4, 12 December

Fischhoff B, Slovic P and Lichtenstein S (1978). Fault Trees: Sensitivity of Estimated Failure Probabilities to Problem Representation. *Journal of Experimental Psychology: Human Perception and Performance,* Vol 4, No. 2, 330-344

Freudenburg W R (1992). Heuristics, Biases, and the Not-So-General Publics: Expertise and Error in the Assessment of Risks. In Krimsky S and Golding D (eds): *Social Theories of Risk*. Westport, Praeger

Funtowicz S O and Ravetz J R (1990). *Uncertainty and Quality in Science for Policy*. Kluwer, Dordrech

MISRA (1994). *Development Guidelines for Vehicle Based Software*. The Motor Industry Research Association, UK

Railtrack (2000). *Engineering Safety Management Guidance (The Yellow Book)*. Issue 3, Railtrack, on behalf of the UK Rail Industry

Redmill F (2002). Risk Analysis - A Subjective Process. *Engineering Management Journal*, 12, 2, April 2002