Escaping the Non-Quantitative Trap

R. D. Alexander, PhD, T. P. Kelly, PhD; University of York; York, UK

Keywords: risk assessment, probability, quantitative, qualitative

Abstract

Quantitative modelling and analysis is common in safety engineering, but it is often criticised. Objections include the difficulty in acquiring probabilities (e.g. for human error), the dubious assumptions often needed to manipulate them (e.g. independence of events), and the inherent uncertainty involved in making decisions based on probabilistic predictions. Clearly, poor predictions are of little value and may be dangerous. Faced with this danger, many people respond by eliminating quantities altogether. This is a trap, as we have no guarantee that the resulting model or predictions will be better; indeed, the subtlety of expression offered by numerical probabilities has been lost. This paper discusses some alternatives to the non-quantitative trap, and explores their significance for the issue of safety case assurance.

Introduction

There is controversy over the use of quantitative claims, particularly probabilistic models, in safety engineering. In many cases, the use of quantitative approaches is at risk, and practitioners are quite happy to discard them in favour of “qualitative” approaches (or, in some cases, no explicit approach at all).

We are concerned that this situation is a dangerous one – a classic example of baby going with bathwater to the detriment of all concerned. It is an invitation to discard a valuable tool and put something worse its place. Or, worse, replace it with nothing at all.

In the next section we will go over the extant threats to quantitative modelling, and present some of the alleged alternatives. Then, we will make the case for quantitative expression – we will show how it provides tools not offered by anything else. The last two sections will discuss the genuine problems that remain with quantitative approaches, and touch upon some solutions.

Quantitative Modelling is at Risk

It is our contention that quantitative modelling and prediction is at risk in many situations. There have been many attacks on probabilistic thinking in print, for example that by Donnell (ref. 1) at ISSC last year. Similarly Dekker, in (ref. 2), criticises the assignment of probabilities to human failures. Manion, in (ref. 4), attacks the use of Fault Trees (and the resulting probabilistic risk assessments) from several different angles. In the extreme case, there are polemics like that of Angell and Straub (ref. 3), who say that probabilistic extrapolation from incident rates is “akin to numerology and astrology”. And of course Taleb’s The Black Swan (ref. 5) has been hugely successful in the popular business market.

More significant than the *printed* attacks, however, is the generally dismissive attitude to we have encountered among workers in safety-critical industries, including some safety engineers. Lines like “Oh, yes, probabilities, Well, let’s not bother with those” are often encountered, and as often as not backed with a knowing smile (“Of course, we all know they’re just an illusion”). Similarly, books such as that by Taleb (ref. 5) are often interpreted as a general condemnation of probabilistic thinking (a careful study of the text will reveal that this is not Taleb’s intent).

These attitudes are not confined to safety – they can be found in finance, climate change, and general discussions of science in policy.

One upshot of this could be a general move away from numbers and mathematics in many field and activities, replacing them instead with qualitative models. Similarly, we could replace probabilistic modelling with

possibilistic thinking, as advanced by Donnell (ref. 1) and Taleb (ref. 5). We could then consider, and plan for (and, presumably, design for) all *possibilities* – all events that are "feasible and possible" (ref. 1).

Quantitative Modelling is a Vital Tool

The threat to quantitative modelling given in the last section is a problem, because quantitative modelling and probabilistic prediction have huge power – power that we need to get a handle on the world. The following sub-sections address different aspects of this power – the power to make good predictions, the power to express certain subtle ideas, and the power to explicitly combine evidence and beliefs. A final issue, particularly relevant to possibilistic ideas, is that our resources are finite – we must have some way to prioritise what we do with them.

It's the Only Way to Make Good Predictions: To make decisions, we need predictions. After all, it makes no sense to choose an option if you don't *predict* that the outcome will be better than the alternatives. In safety, the prediction is often that a given design option (or policy, or decision to certify a product) will lead to fewer injuries and deaths.

A corollary of this is that every decision has implicit predictions – if we observe you make a decision, we can infer your decision criteria based on what you chose – "You took that decision because you predicted X…". Any such inference is clearly tenuous. However, we can also express that inference in another form: "If you follow your current (implicit) policy, the effect on variable X is likely to be…" (where 'X' might be "human lives lost"). Such an expression is compelling to someone who follows a utilitarian ethic; predictions let us consider the consequences of our actions.

Partly, predictions are useful because they can be wrong, and we need to be wrong in order to learn. Of course, we can *be* wrong (make bad decisions) without explicit predictions, but if our predictions are implicit we may not relate them to the bad decision. We cannot learn from our mistakes (and we rarely learn from success because most success fulfils our prior expectations; it doesn't challenge them).

To learn well, then, we need predictions that are definite. 'Definite' means either certainty ("This accident will definitely not occur") or a quantified uncertainty ("There is a 1×10^{-5} per flight hour chance of the accident occurring") that can later be compared to the actual results.

It's the Only Way to Express Certain Ideas: Quantitative models and expressions are the only way to capture certain aspects of thought.

Crucially, we need numbers for precision. Numbers give us a precision that words cannot. The key use in safety is to express how much, to what degree, you believe some claim. For example, how else can you distinguish clearly between "This event will happen once a day" and "This event will happen once every thousand flight-years"? To achieve this we need numbers. At the least, we need codified labels like "remote" that code for some numerical probability, as in MIL-STD-882D (ref. 6).

We can also use numbers for subtlety and scoping. Numbers allow us to have degrees, proportions and rates – we can subdivide qualitative groups, and reason about how much (or what proportion) of them have certain qualitative properties. They also allow comparison of proportion.

This power is not about *arithmetic*. It's not about manipulation of numbers to achieve results with certainty. It's about mathematical language as an extension of your *vocabulary*, as an extension of the range of constructs you can use for explicit *modelling*. A wider range of terms means a wider range of components in your models.

There are qualitative equivalents for all the quantitative expressions discussed here (examples include "regularly", "hardly ever", "some", "almost all"), but they are very blunt tools. When we use them instead of numbers, we pay a great cost in precision and subtlety.

It's the Only Way to Combine Certain Things: Probabilities let us combine certain types of belief, confidence or uncertainty. First, they let us combine the individual probabilities of multiple events into a single probability for a related event. In safety, the most common use of this is in Fault Trees. Secondly, they let us take several expressions of confidence or uncertainty, based on diverse evidence, and combine them. For example, see Ord et al in (ref. 7) – they combine the 'aleatoric' probability of certain events (which they believe to be random[1]) with the epistemic probability that their model of the situation is wrong.

Certainly, we can combine beliefs and confidence levels in our heads – we do it all the time. However, this intuitive combination is forever condemned to be implicit, subjective and personal. Using probabilities (and their associated mathematics) lets you render this explicitly and externally for others to challenge and consider.

It's Necessary Because Resources Are Finite: Finally, in answer to the possibilistic criticism, we are stuck with numbers because we must prioritize potential futures and allocate resources accordingly. Many things that we need are finite, and not just tangible things like money and time – we also have finite capacity for degraded capability in essential infrastructure (e.g. when safety measures impair the performance of fire or ambulance services), finite tolerance for military casualties (when safety features on a military system impair its performance). A purely possibilistic approach is inadequate – we must distinguish between the possible and the probable.

Quantitative Modelling *Does* Have Problems

We won't deny that there are problems with numbers and probabilities – even on the quantitative side of the trap there are many smaller traps lurking for the unwary. For all the power we get from numbers, there remain long-standing problems with quantitative modelling. This is especially true of probabilistic models.

The Realist's Problem – "You Can't Get the Numbers": The first objection is that although probabilities could theoretically be used for reasoning about system safety, we cannot get the probabilities that we need because we cannot observe them. Often, we don't want probabilities for simple events in systems that exist; we want them for complex events in systems that we haven't created yet. An extreme example is to contrast the probability of getting '16' on one spin of a roulette wheel with the probability of a latent systemic fault being present in a complex software system. This is perhaps the most common objection.

The Mathematician's Objection: A common solution to the problem above is to declare that the probabilities you are working with do not need to be derived from physical reality – they, instead, represent subjective beliefs. This, however, can lead to the objection that the *combination* you are performing on probabilities is not valid – the probabilities you have acquired are 'not real ones'. In other words, they cannot be explained in terms of the simple physical systems on which intuitions of probability are based. The argument follows that any derived probability (e.g. by combining them in a fault tree) is equally invalid. This is the primary objection to the use of probability in safety engineering identified by Watson in (ref. 8) – the authority of probabilistic models has traditionally been assigned on a basis of their claim to be objective representations of the probabilistic behaviour of the world.

The Risk of False Precision: Numbers can be very precise. Numbers, where they correspond to an easily-observable physical fact (the length or mass of a solid object) can be precisely correct to within a well-defined measurement tolerance. Not all numbers have this property, and some will have very large errors, indeed errors that are difficult to quantify. Treating these highly-uncertain numbers in the same way as precise measurements is unsound.

The Risk of False Confidence: Once in numbers, it is easily to stay in numbers – once you have a quantitative model, it's easy to continue working in terms of those numbers and disregard its (perhaps tenuous) correspondence with the real world. This is similar to the situation with computer models discussed by the authors in (ref. 9) – one can be lulled into working in terms of a model, not in terms of the world.

The Stakeholder Problem: Many stakeholders in your analysis can't (or won't) understand the numbers you want to present. This especially common when your stakeholders are non-engineers: managers, elected officials, and "the general public". Sometimes, these stakeholders will reject quantitative data. Perhaps worse, they may think they understand, but in fact misinterpret and misuse your numbers. The latter is worse, because giving them numbers is now dangerous rather than merely futile.

---

[1] See Taleb in reference 5 for a criticism of the aleatoric/epistemic divide

Steps Towards a Solution – Embracing Subjective Probability

We can point towards some work that offers solutions to these problems, although as ever these solutions raise problems of their own.

Probabilities as an Argument Tool: Watson, in (ref. 8), develops idea that a probabilistic model can be used to express an argument. Indeed, he argues that this is the *only* sound way in which probabilistic models can be used. He sums his core suggestion like this: *"If some individual agrees that the input probabilities fairly measure the degree of uncertainty about the uncertain input variables which he or she has, or is prepared to accept on the advice of others, and if he or she further agrees that the relationship between these input variables is adequately represented by the models used in the analysis, then it is reasonable to adopt the output probability distributions as measures of how uncertain he or she ought to be about these variables"*. In other words, we can treat a probabilistic quantitative model as an explicit argument that we present to others in support the validity of its output.

If anyone objects to an argument in this form, Watson notes, then they are making an argument themselves – a counter-argument to the first argument. At this point, we can assess the relative merits of the two arguments using whatever judgement we have. Where we can't agree on this, we can turn the assessment over to others, which seems to be Watson's expectation of the normal course: *"Whether our argument was more acceptable than theirs would depend on other's detailed judgements of both" (ref. 8).*

We can see echoes here of Manion's view that Fault Trees give highly subjective, value-laden results, rather than an 'objective' assessment of risk (ref. 4). Similarly, Manion suggests that it is the social duty of risk assessors to make clear to the public exactly how they arrived at the risk figures they did (given the inevitable subjectivity involved in the creation of a fault tree). This position (that a quantitative analysis is an argument subject to public assessment rather than a formalised discovery of objective truth) is obviously destructive to the role that the engineer traditionally aspires to, but it is far more supportable given the reality of the quantitative analyses that we can actually do. Watson is keenly aware of the difficulties this raises, but suggest that his approach *"would allow us to use [Probabilistic Safety Analysis] as it should be used, as a vehicle for thinking clearly about the future"* (ref. 8).

In any case, it may be that the realist position is entirely unsound. Watson clearly thinks so, approvingly quoting de Finetti's "PROBABILITY DOES NOT EXIST". Watson seems hopeful that his approach resolves the traditional objections to probabilistic models in safety. Indeed, goes some way to tackle four of the five problems raised in the previous section. It does nothing, however, to resolve the stakeholder problem. It may even make it worse, as it is now necessary for stakeholders to not just understand the input and outputs of the model, but also the details of the model itself.

Combining Bayesian Belief Nets with GSN: Some steps towards an instantiation of Watson's ideas can be seen in the work by Wu and Kelly in (ref. 10). This work combines Bayesian Belief Nets (BBNs) (for modelling architectural properties of a system) with the Goal Structuring Notation (for arguing the validity of the properties derived).

The work is based on failure logic modelling, where the properties of interest are whether particular failures will occur. Wu and Kelly uses a hierarchical modelling approach, whereby the overall system architecture is represented by a BBN, and individual components can be represented by BBNs in turn, or by conditional probability tables for the failures of concern. By composing the hierarchy according to accepted means of BBN analysis, failure probabilities for the overall system can be derived. If the resulting failure probability is acceptable, then the model can be used in support of a safety case for the system.

Wu and Kelly acknowledge that the conditional failure probabilities for the low-level architectural elements will be highly subjective, and that even the causal relationships within the wider architecture may be controversial: *"given the same deviation, different architects may have different BBN models in terms of the DAG models (i.e. a set of causal factors) and conditional probabilities (i.e. the relative strengths of these casual factors)"* (ref. 10). However, Wu and Kelly's model is explicit and is therefore open to external criticism and challenge. It does not purport to be a closed model that generates a 'true' result. To go further, indeed to embrace this subjectivity, Wu and Kelly show how a structured argument (expressed in GSN) can be created to justify that the model supports the claim "the system is acceptably safe".

The arguments created by Wu and Kelly explain *why* the BBN is an adequate representation of the system, and both identify the assumptions made and justify why those assumptions are acceptable. The use of an explicit argument as

part and parcel of their approach is a tacit admission that all models are flawed, but that models of some kind are essential and that we will have to assess their validity in individual cases. The admission, and the position of relative humility that goes with it, is a step towards the alternate relationship envisaged by Manion (ref. 4) and discussed above.

With respect to the stakeholder problem, the work has some potential. The idea of certifying a system safe based on a structured safety argument has made great inroads into UK industry and regulatory bodies. Indeed, many standards (such as the military equipment standard Def Stan 00-56 (ref. 11)) now make it mandatory. GSN has been at the forefront of this movement. At the very least, the work represents a way forward for getting subjective probability into general use. In particular, there is potential for BBNs to form an explicit stage in a continuum going from "Hard Quantitative" to "Pure Qualitative". This continuum may be valuable in getting people to understand how the quantitative and the qualitative interrelate, and that there *is* something of worth in the space in between.

The Common-Sense Problem: The solutions discussed above are problematic because they unravel a number of common-sense assumptions. In particular, Watson challenges the assumption that there is an objective, 'correct' value to which all analyses aspire. This is well outside our scope here, but we can note that this may create a further problem for dealing with stakeholders in that they expect the aim of analysis to be objective truth. Paradoxically, the subjectivity of safety assessment that makes their involvement ethically necessary may be an obstacle to their successful engagement.

Can We Afford the Stakeholder Problem?: One solution to the stakeholder problem is to discard the quantitative and probabilistic representations that they cannot understand. This is an understandable response, but it is a dangerous one.

It is possible that many "scientific" objections to the use of numbers ("you can't get the numbers" etc) are fronts for social ones (numbers, no matter how valid, are practically useless because key stakeholders *can't* correctly interpret them). To override these stakeholders is politically impossible (e.g. they are the employers and paymasters of the analyst) or ethically unacceptable (e.g. there is an expectation that system developers will consult with the general public, or at least with their elected representatives). But if this *is* the case, we should be honest about it. At the very least, it will make this problem apparent.

Meanwhile, when you remove quantitative models you are depriving the skilled and intelligent of tools they could use for better communication. Many young people, and many of the marginally educated, are stuck in world of absolutes – without quantities, probabilities, or measures of degree, they're stuck with crude categories. If you doubt this, try reading the letters page of a tabloid newspaper, or visiting any internet discussion forum. When you refuse to use and combine quantities you're reducing yourself to that level.

On a related note, the objection "But it could be misused!" is a poor one when you're referring to your peers, to your fellow scientists and engineers. If you refuse to use numbers because your peers may misuse them, then you're transmitting a pretty poor opinion of those peers. It is possible that this opinion is justified, but in that case we have far worse problems, and abandoning numbers will not solve them.

Conclusions

Numbers are threatened in safety engineering, particularly those numbers that are of uncertain validity. In any situation, there's a risk they'll be swept away in favour of a "more qualitative" approach. It is our position that, far from being an easy choice, this is a sacrifice that is rarely warranted.

Quantitative models, especially probabilistic ones, greatly extend our vocabulary. They provide means of expression that we can't get elsewhere. Numbers provide a tremendous range of subtlety and expression, whereas purely qualitative arguments can be terribly blunt. It follows that we should not limit ourselves to those cases where hard numbers are available. Certainty in numbers is a luxury we cannot afford.

Despite their strengths, quantitative approaches have many of problems. We can resolve the worst of these by embracing quantitative models as an argument tool. However, it remains true that many stakeholders cannot (or will not) understand quantitative models, or even their results, and this is likely to be a problem for the foreseeable future.

References

1. A. P. Donnell, The Black Swan and Nuclear Weapon Safety, in Proceedings of The 26th International System Safety Conference (ISSC '08), 2008.
2. S. Dekker, Don't Errors Exist?, in Ten Questions About Human Error: A New View of Human Factors and System Safety: CRC Press, 2004.
3. I. O. Angell and B. Straub, Rain-Dancing with Pseudo-Science, Cognition, Technology and Work, vol. 1, pp. 170-196, 1999.
4. M. Manion, The Epistemology of Fault Tree Analysis: an Ethical Critique, International Journal of Risk Assessment, vol. 7, pp. 382-430, 2007.
5. N. Taleb, The Black Swan: The Impact of the Highly Improbable: Allen Lane, 2007.
6. MIL-STD-882D - System Safety Program Requirements, US Department of Defence, 2000.
7. T. Ord, R. Hillerbrand, and A. Sandberg, Probing the Improbable: Methodological Challenges for Risks with Low Probabilities and High Stakes. arXiv:0810.5515v1 [physics.soc-ph], 2008.
8. S. R. Watson, The Meaning of Probability in Probabilistic Safety Analysis, Reliability Engineering and System Safety, vol. 45, pp. 261-269, 1994.
9. R. Alexander and T. Kelly, Simulation and Prediction in Safety Case Evidence, in Proceedings of the 26th International System Safety Conference (ISSC '08), 2008.
10. W. Wu and T. Kelly, Combining Bayesian Belief Networks and the Goal Structuring Notation to Support Architectural Reasoning About Safety, in Proceedings of SAFECOMP 2007, 2007.
11. MoD Interim Defence Standard 00-56 Issue 4 - Safety Management Requirements for Defence Systems, Ministry of Defence, 2007.

Biography

Dr Robert Alexander, Ph.D., Department of Computer Science, University of York, Heslington, York, YO10 5DD, UK, telephone – +44 1904 432792, facsimile – +44 1904 432767, e-mail – robert.alexander@cs.york.ac.uk

Dr Rob Alexander is a Research Associate in the High Integrity Systems Engineering (HISE) group in the Department of Computer Science at the University of York. His research focus is on safety engineering for unmanned and autonomous systems, although he is active in a range of other areas including System of Systems engineering, software safety, and real-time risk awareness. He is also involved in research on the evaluation and accreditation of complex simulation models. Rob has published papers in international conferences on a range of systems safety issues. He graduated from Keele University in 2001 with first class honours in Computer Science, and was awarded his doctorate in 2008 by the University of York.

Dr Tim Kelly, PhD., Department of Computer Science, University of York, Heslington, York, YO10 5DD, UK, telephone – +44 1904 432764, facsimile – +44 1904 432708, e-mail – tim.kelly@cs.york.ac.uk

Dr Tim Kelly is a Senior Lecturer in software and safety engineering within the Department of Computer Science at the University of York. He is also Academic Theme Leader of the UK MOD Software Systems Engineering Initiative Dependability Theme. His expertise lies predominantly in the areas of safety case development and management. His doctoral research focused upon safety argument presentation, maintenance, and reuse using the Goal Structuring Notation (GSN). Tim has provided extensive consultative and facilitative support in the production of acceptable safety cases for companies from the medical, aerospace, railways and power generation sectors. Before commencing his work in the field of safety engineering, Tim graduated with first class honours in Computer Science from the University of Cambridge. He has published a number of papers on safety case development in international journals and conferences and has been an invited panel speaker on software safety issues.