Safety and Reliability

ISSN: 0961-7353 (Print) 2469-4126 (Online) Journal homepage: www.tandfonline.com/journals/tsar20

Risk analysis and big data

Ben Ale

To cite this article: Ben Ale (2016) Risk analysis and big data, Safety and Reliability, 36:3, 153-165, DOI: 10.1080/09617353.2016.1252080

To link to this article: https://doi.org/10.1080/09617353.2016.1252080

SAFETY AND RELIABILITY, 2016
VOL. 36, NO. 3, 153–165
http://dx.doi.org/10.1080/09617353.2016.1252080

Taylor & Francis
Taylor & Francis Group

ARTICLE

OPEN ACCESS

Risk analysis and big data

Ben Ale

Safety Science, Delft University of Technology, Delft, the Netherlands

ABSTRACT

Big Data can help overcome various problems that exist in present risk analysis practices. By analysing systems as a whole, it is no longer necessary to define in advance what a failure is and what a success is. It is also possible to evaluate how factors that are considered to promote success can combine into catastrophic failures. Big in Big Data is relative. What was called big data 25 years ago is now called small. The continuous development of analysis techniques over the years have resulted in several operational models that use the concept of Big Data. They will become better as the technology and the accessibility of data further improves. With this new generation of systems models, accidents and incidents do not have to wait for analysis to after the fact. They can be studied beforehand in a model. Replacing hindsight by foresight can help to make the world safer, if we desire to do so.

KEYWORDS Risk analysis; Big Data; hindsight; foresight

Introduction

Risk analysis has been hampered by the lack of sufficient and reliable data since its beginning. This is especially the case for low probability events, for which any data set has to be large to have a reasonable probability that these events show up in the data set at all (Mayer-Schönberger & Cukier, 2013). This situation was even more aggravated by the impossibility to process complete data sets because of the limitations of information technology. In risk analysis, therefore, most of the time only accidents are investigated and the evaluation of the probability of events that where physically possible, but had not happened yet, at least as far as the analysts know, are mostly judgements based on similarities and extrapolations. As an example, the 2005 Buncefield explosion was considered to be impossible, despite the fact that similar explosions had taken place already at least 5 times, the first recorded one dating from 1970 (Burgess & Zabetakis, 1973). Even in hindsight, lessons are difficult to learn. The availability of

CONTACT Ben Ale ben.ale@xs4all.nl Jaffalaan 5, 2628BX, Delft, the Netherlands

larger scale data and the power to process them has improved such that analysis of all available information has become possible. The name Big Data given to it can only be interpreted as relative to the past, when data were less and therefore small, although even 25 years ago, large data sets were used to generate information, be it slowly. The development technology to use them has been developing ever since, in all sorts of projects. Big Data help us to overcome the hindsight problem. By analysing the behaviour of systems in relation to the variability of parameters, variables, circumstances and human behaviour over the full range of outcomes, from great success to utter failure, problems and accidents that might emerge in the future can be predicted before they happen in the real world. Hindsight still exists, but after the calculation, not after the accident. This will allow more founded decision making on what may be sensible, effective and economically viable interventions to make systems more safe, if we wish do to so.

Risk analysis

With risk analysis we try to describe and where possible quantify what can go wrong and what harm that can do, with the ultimate purpose of finding measures to enhance safety. Basic causes of harm have been described as deviations (Kjellén & Hovden, 1993). Deviations in the context of deviation modelling of accidents are defined in relation to specified requirements and coincide with the definition of nonconformities in the International Organization for Standardization's ISO 9000 series of standards on quality management (ISO 1994). The value of a systems variable is classified as a deviation when it falls outside a norm. Systems variables are measurable characteristics of a system, which can assume different values as found in Saari, 2016). In deviation models, it therefore is recognized that system variables can assume a range of values without leading to an accident. This range is called the norm. When a variable is outside the norm, it is assumed to be harmful and leading to an accident.

This definition is helpful in quality control and safety management. Keeping variables inside the norm assures standard quality and thus safety. But, does it? This depends on what norms are taken into account. Although by one norm all instrumentation should always work according to specification, by another norm repairs can often wait until a scheduled maintenance operation. Although defective instrumentation is outside the norm, it is nevertheless considered normal not to repair it. In many cases the operational procedures do not demand that a system is halted because of a defect. In a motor car the only indicator which is imperative in this respect is the 'oil' light, because insufficient lubrication will ruin the engine. For all other warnings repairs can wait until one gets to a workshop or until scheduled maintenance, which may be 15,000 km or more away. In this and in

most cases there is considerable safety built into the norm. The edge of the norm does not coincide with the edge of sellable quality, collapse or failure. Unfortunately being inside the norm does not guarantee safety or quality. Every product has a warranty. It is accepted that eve with tight quality control, such as in the production of light bulbs, a product may be defective on sales – hence, the option to check whether the bulb works in the shop where it is bought – or not function for the specified duration. Even if all parameters are conform the norm, undisturbed functioning is not guaranteed.

What is unsafe?

This creates an obvious problem in a hazard and operability study or any study aimed at constructing accidents sequences from base failures, defined as deviation. There does not seem to be a hard criterion for what constitutes a failure. A broken part is not necessarily outside the norm.

On the other end of the path from cause to accident, there is a similar problem. How to define a system to be unsafe?

The above is even more complicated for the classification of human actions. Is ignoring a warning signal unsafe, or is it part of the accepted standard operating procedure: lights can wait or even lights need to wait. Hollnagel (2014) describes in his book Safety I and Safety II the example of signals passed at danger (SPAD) in railway operations. The operation of getting to a stand-still when a signal is at danger is performed by the train driver. This leads in a relatively large variety of the point where the train gets to a complete stand still with respect to where the signal is. In a number of cases the train actually stops after the signal post. As long as the train has not reached the potential point of collision, i.e. where the crossing or the points actually are, even then there is not necessarily any potential for a collision.

In a block system it is advantageous if a train reaches the end of a block and clears it as fast as possible. This maximizes the capacity of the track. However, the closer a train follows another train the higher the probability of encountering a signal at danger. A skilled operator will have a mental estimate of where the other train is and usually the signal changes from danger to free, before he reaches it. It is therefore almost good operatorship to risk a SPAD as long as the train can be stopped before the real danger point (Van den Top, 2010). Whereas, the safety people will consider this unsafe, train operators usually are quite comfortable with this way of operating a train and statistics support them. According to Hollnagel there were 130 SPAD events in Belgium in 2012, of which one third were serious. It was also estimated that there were about 13,000,000 cases per year of trains stopping at a red signal, which means that the probability of a SPAD

is 10^{-5}. According to Hollnagel, a value of 10^{-5} means that the system is safe, although not ultra-safe. What is missing in this example is the metric. The metric is SPADS per encountering a red signal. As described above, the block system maximizes the encounters of red lights. Therefore the system can be improved by informing the operators where the other trains are so that they can adjust their speed to that of the train in front of them rather than 'finding out' whether they are to close by looking at the signal (Van den Top, 2010).

The above is an example of the problem of defining what and when a system is safe. It is also interesting to note that the example given in the introduction of Hollnagels Book is an illustration of the difficulty in defining what success really means, which makes safety II just as poorly defined as Safety I.

Variability

Variability not only is unavoidable it often also is desirable. As Orton & Weick (1990) and Perrow (1984) explain, loosely coupled systems are more forgiving in case of delays or faults. However the coupling can be too loose, in which case the output of a system becomes unpredictable. On the other hand, if a system is too tightly coupled, there is no room for error. Systems are like ball bearings. Too loose and they rattle, too tight and they heat up. What is precisely right is approximately known and a system designer has to choose the desired level of coupling and live with the residual variability. In heavily used railway systems being on time is demanded by the customers and need tight coupling. To be safe needs some margins and requires loose coupling. These contradicting demands need to be brought together in a single system.

In any case if the behaviour of parts of the system is variable, so will the outcome. The extent of variability is often underestimated. The normal way of estimating the probability of a deviation of a certain size is to take a sample, determine the mean and the standard deviation and assume a normal distribution. However it has been shown now in several books (Taleb, 2007; Hand, 2014) that the normal distribution underestimates the probability that a value is far from the mean for many situations and that there are distributions with similar shape, for which this probability is much higher. As an example, the probability of a value that is 5 standard deviations away from the mean is $2.97{\cdot}10^{-6}$ in the bell shaped normal distribution and 1.22×10^{-2} in the equally bell shaped Cauchy distribution, a 4 orders of magnitude difference (Figure 1). In such a distribution, the occurrence of a black swan is about as probable as the occurrence of a white elephant. Accidents are no longer rare, they are normal accidents.

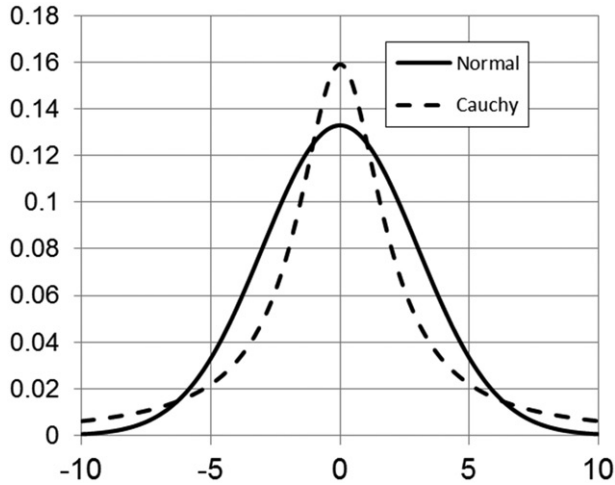

Figure 1. Two bell-shaped probability distributions.

Uncertainty by itself increases the probability of events with larger outcomes. One possibility of expressing risk in a mathematical form is the FN curve. This curve depicts the cumulative frequencies that consequences of size N are exceeded. For this curve to exist it is only necessary that the consequences can be expressed in a single metric. It is assumed that the probability of occurrence can be expressed as a frequency. This implies that the event can repeat itself. The inverse of the frequency is the return period. It can be argued that events never really repeat themselves, because the circumstances will have changed, or the technological structure has been rebuilt after the previous event and therefore are not the same. However, the assumption that near enough or similar is the same as equal is a necessary assumption to be able to make any estimate of probability. The next die taken from a box is not the same die as was taken earlier, LPG tanks are not the same either. But, they are sufficiently alike that they can be treated as being the same for the purposes of say the calculation that 6 will come up or that an LPG tank will explode. Uncertainties in the estimates of extents and frequencies can be integrated into the FN curve. When there are uncertainties and there is a desire to explicitly deal with them, each F, N pair expands into a little FN curve. FN curves can be summed from large N to small N just as the FN pairs. What results is a single line diagram depicting the aggregated result of all uncertainties: aleatory and epistemic. With respect to the FN curve with non-uncertainties, the curve will be rotated counter clockwise. When the total probability of an accident remains the same, the intersection with the vertical axis will not change (Figure 2).

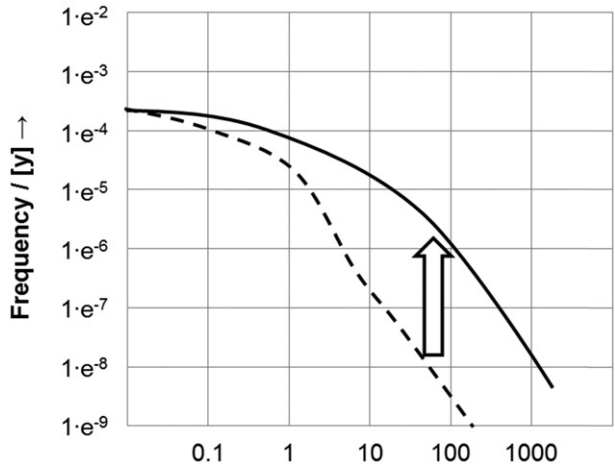

Figure 2. Variability and uncertainty rotate the FN curve counter-clockwise.

Evaluating accidents

In many cases the causes of an accident or even a type of accident is derived from the analysis of a single or a few accidents. After an inquiry into an accident recommendations are made to prevent reoccurrence. In such a case, it is rarely known whether such an intervention is indeed effective. The accident might be a unique coincidence of factors, or the solution might make things worse. A much better approach would be to investigate all the accidents in a certain industry or with a certain activity (Bellamy et al. 1999). That would allow drawing conclusions on dominant common accident causes. Today this would be called the use of big data. But, several systems have been developed in the past that used all available data, be it that they did not have to be extracted from the cloud. They were already in an accident database. An early example is the analysis of all reported pipeline accidents and incidents in the EU (EGIG. 1982).

The occupational risk model was developed in 1993 (Hurst, Bellamy, Geyer, & Astley, 1991; Ale, 2006). It was based on the analysis of all reported occupational accidents in the Netherlands in over 5 years (Papazoglou & Ale, 2007). This allowed some ground-breaking conclusions. The root cause of accidents in the construction industry, which in the Netherlands is the sector with the largest number of casualties, is motivational. Although safety measures and protection equipment usually is supplied and present, they are not used because people do not care. This led to a nationwide campaign aimed at that industry to improve motivation rather than trying to invent yet another technical measure. It also indicated that reduction of the number of accidents in this sector would be difficult, as it depended on

the motivation of the individual workers, for who the threat of being killed apparently was insufficient motivation to work more safely.

Another example is this realm is the measure to replace ladders by scaffolds for a whole series of activities at height. In fact, this proved to be counterproductive. The analysis of all ladder and scaffold accidents combined with a nationwide inventory of the number of missions on ladders and scaffolds and the duration of these missions showed that in fact the use of scaffolds for many applications was more risky than the use of ladders (Bellamy et al. 2007). The use of the data of all exposures to a risky activity including the circumstances is another example of the use of big data, although here the data were not even available at the time and needed to be generated in a nationwide survey among the Dutch labour force.

Along the same lines the Causal Model for Air Transport Safety was developed and quantified on the basis of an analysis of some 10000 accidents and incidents in the ADREP data base of air traffic incidents and accidents (Ale et al. 2009). This model later was used among others to search for the most likely cause of an airline crash in the Netherlands in 2009 (Ale et al., 2010). Also, this development would now be designated as a use of big data.

Fortunately the Netherlands is small, only 40,000 km^2 of land mass. This allowed nationwide calculations of risk exposure on a detailed level already in the early 1990s. At that time the first database with buildings, including houses, and the number of occupants was created from what previously was an essentially a paper record of 'where everything was'. From this, among other, population density numbers were derived in a resolution of 100*100m squares, all 4 million of them. This in turn allowed the development of a zoning policy for hazardous chemical installations (Ale, Lahey, & Uijt de Haag, 1996) and the evaluation of the geographical distribution of the accumulation of various environmental and health risks in the population (Pruppers et al. 1997). These number of crunching exercises took place on hardware with capacities which today would be considered insufficient even for a moderately powerful mobile phone, but nevertheless were early examples of attempts to 'analyse them all'.

Although after the accident many if not all experts will say that you could have seen it coming, apparently this is not the case, or the accident would have been prevented (Dekker, 2002). Nevertheless there are many situations in which it is assumed that the probability of an accident is higher than it should be, a situation that we call unsafe. Although in many instances this may be true, there are also situations in which we are misled. One of the reasons is that it is known what factors were present in situations where there was an accident, but what is not known is in how many situations the same factors were present when no accident

happened. It is not even known whether certain factors that are designated as being causal to the accident are more dominant in situations where accidents happened than in situations where it did not. Only because of the mission survey, which investigated all missions with ladders and not only those where an accident happened, it appeared that it made no difference at all whether the ladders used had an EU certificate or not. There were as many uncertified ladders in the missions that did not result in an accident as in the missions that did (Bellamy et al., 2007).

Modelling the system rather than the failures

Therefore, it would be preferable if a model could be built in which it is not necessary to know what causes an accident, or to know in advance what is right and what is wrong. The results of evaluations with such a model could be interpreted just as situations are interpreted in reality. If the outcome contains events that are judged to be negative or an accident, the model construct is unsafe when all negative outcomes are judged unacceptable. If the outcome is a probability or the frequency of occurrence of an accident, the construct can be called unsafe if the probability is judged to be too high, it can be called safe enough if the frequency is judged to be acceptable. If then the observed frequency of accidents falls within the reliability bounds of the model one can argue that the accident, how unfortunate it may be, is normal. One can even decide not to take measures to decrease the probability further.

Cats

A first attempt to build such a model was the Causal Model for Air Transport Safety (Ale et al. 2005). In the model the air-traffic system is considered from the point of view of air travel. This means that a successful flight from an airport of origin to an airport of destination is the activity of central concern. Building an aircraft, maintaining it, loading it, fuelling it, getting cargo and passengers aboard and guiding it through airspace all are considered as activities that are associated with this journey. In previous attempts at modelling the risks associated with air traffic the focus was, as it usually is in accident models, on catastrophic events and in particular the crash of an airliner into inhabited areas. In CATS the air traffic system and its safety functions are modelled in such a way that the relationship between the various components of the system and the management system in the model are sufficiently realistic as to make it possible to model the population of flights, determine what combinations of factors result in successful arrival and what flights result in an incident or accident. For the latter flights risk reduction alternatives within a given (sub) system – such

as an airfield – can be investigated and also differences between different systems. The model has sufficient capabilities to allow quantification of these differences to support cost benefit comparisons (Ale et al. 2008). The model uses Bayesian Belief Nets as the means to capture these relationships. BBNs represent a concise way of representation of joint probability distribution of a set of variables. By definition, a BBN is a directed acyclic graph in which nodes represent random variables and arcs represent probabilistic or functional influences The introduction of the distribution free BBNs (Kurowicka & Cooke, 2004) allows the variables to be either discrete or continuous, and the influences to be represented by rank correlations or functional relations. The rank correlation between two probabilistic nodes represent the degree of association of high values of one variable with high (for positively correlated) or low (for negatively correlated) values of the other variable. On the other hand, functional nodes can be any analytical function having as arguments its parent (or influencing) nodes. This implies that also the logical deterministic relations, such as when a law of nature is involved, can be transformed into BBNs. The model was populated with data on the basis of the analysis of accident and incident data and by expert judgement exercises using the method described by Cooke and Goosens (2000). In the model all variables are distributed. This means that the model can be run with all variables set to the average values derived from the data analyses as described in the earlier papers referred to above; it can be run in a Monte Carlo type mode in which selections are based on the probability of occurrence of certain values given the mean and the distribution of the variable in the dataset or a value for a parameter can be selected that corresponds to a known condition. The model BBN consists of approximately 1400 nodes and 5000 arcs. A method was developed to incorporate human behaviour and factors that shape performance in the BBN as well (Lin et al., 2008).

Platypus

With the technology developed in the CATS project, a further development was made in the Platypus project. Platypus models the left hand or fault-tree side of a bowtie model for safety process by using a Bayesian Belief Net (BBN). That is to say, it calculates the frequency of occurrence of leaks of arbitrary size in process equipment. A leak or loss of containment (LoC) event does not necessarily lead to a catastrophe but in some cases leads a loss of containment can have severe consequences ranging from accidents or catastrophes. Platypus contains engines to read equipment lists, to change parameters in subsets of process equipment in a plant and copying subsets of the plants. This combination of features has led to a piece of software that enables BBN-based risk analyses of entire chemical or process

plants, on the basis of a plant model that is generated from the complete inventory list, where in earlier models a comprehensive shortlist of failure outcomes needed to be defined à priori (Uijt de Haag, Ale, & Post, 2000). The use of a BBN as an advanced method for risk calculations yields new possibilities for risk analysis. Platypus generates risk analyses that yield risk distributions rather than point estimations. These risk distributions automatically address aleatory uncertainties in risk analysis (whose origin lies in randomness) and identifies instances of high risk that are overlooked in setting risk standards. The distribution also yields new leading risk indicators and new methods and instruments for risk management. When estimates of epistemic uncertainty ranges are available, these can be incorporated in the overall distributions as well. Platypus provides insight into the factors that influence the frequency of LoC events before accidents have taken place, whereas traditionally that information could only be derived from historical data. In Platypus, technical, organizational and human factors elements are integrated which enables an in-depth analysis of causes of LoC and the effectiveness of measures. Due to its efficient calculation engine, initial LoC estimates for entire sites can be calculated in a matter of minutes which makes it very easy to minimize the LoC frequencies in process plants in the future. Therefore Platypus is a leap forward in estimating leak frequency rates for process plants. From the point of view of data handling, this is an example of the use of all available information, rather than a subset, which is one of the characteristics of Big Data (Mayer-Schönberger & Cukier, 2013).

Hindsight before the fact

With the use of Big Data technology, it is no longer necessary to define failures and failure paths beforehand. The full behaviour of a system can be modelled including the variability of parameters. In a BBN events and parameters can have a distribution and the influences between events can be probabilistic just as in the real world. The model will show what Hollnagel (2006) calls emergent behaviour just as reality does. The model results can be observed and unwanted behaviour can be detected. This would be hindsight after the calculation, but before the unwanted behaviour emerges in reality. The causes can then be explored which can be the combination of factors that are considered a bit extreme but accepted as normal. Having airplanes diverted to small airfields was considered normal, just as having pilots and co-pilots with large differences in experience, air traffic controllers with limited command of the English language and taking off with limited visibility. Nevertheless they combined in 1977 at Tenerife in the largest airline accident to date.

Average chemical plants do not explode on average, nor do average airplanes fall from the sky or average cars collide. But then, the world is not

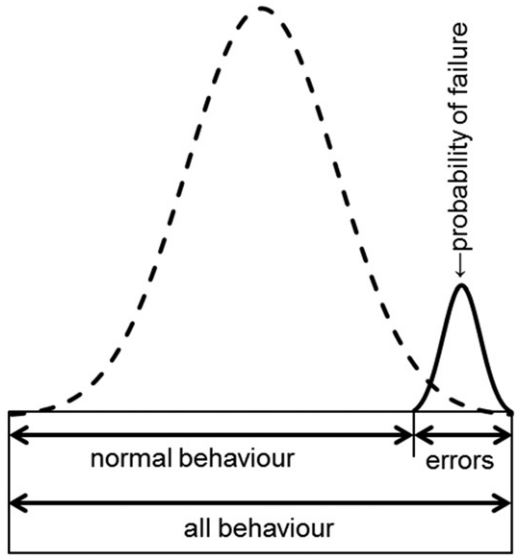

Figure 3. Model simulations show potential failure conditions.

an average. The world on the average of the universe should not exist, but it does. Now that in many areas of technology the probability of failure is low but the consequences of failure could be catastrophic, it is no longer sufficient to look at predefined abnormalities. Unfortunately looking at success, or why a system keeps working is not sufficient either. The factors listed above that combined into the Tenerife disaster can all be designated as contributing to speedy and successful operation of an airline. And they were, before 1977. It was the combination of success factors that created the disaster. The result of modelling all behaviour is conceptually depicted in Figure 3. What is unsafe will reveal itself.

Using BBN technology the correlation between increased probabilities of success and values of parameters in system can be explored. From that it can be decided whether to curtail variability by selection or exclusion of certain actions.

Speed is the main factor in traffic accidents. The first roundabout was constructed in Bath (UK) in 1754. The modern roundabouts invented in the UK in the 1960s. They reduce the variability in speed and take the high speeds away, which in turn led to a significant decrease in accidents.

Conclusion

Several models have been developed in the past decennia that can exploit the availability of large data sets. By using Bayesian Belief Nets the variability of the world can be modelled as it is without the need to pre-set what is

right or what is wrong. The more data become available the more detailed these models can be and the better the description of complicated systems can answer the question what interventions contribute to success without also laying the foundations for a catastrophe. When used wisely Big Data can help to make the world a safer place.

Disclosure statement

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of this article.

Notes on contributors

B.J.M. Ale, is emeritus professor at TU-Delft in Safety Science and Disaster Management. He was full professor from 2002 to 2012. He also was professor at the University of Ghent and of EPFL in Lausanne. B.J.M. Ale is visiting professor in Risk management at the University of Antwerp.

References

Ale, B.J.M., Lahey, G.M.H., & Uijt de Haag, P.A.M. (1996). *Zoning instruments for major accident prevention*. Paper presented at International Conference on Probabilistic Safety Assessment and Management, Crete. Proceedings (C. Guedes Soares ed.), Springer 1997, p 2191–2196.

Ale, B.J.M., Bellamy, L.J., Roelen, A.L.C., Cooke, R.M., Goossens, L.H.J., Hale, A.R., … Smith, E. (2005). *Development of a causal model for air transport safety*. Paper presented at Proceedings of IMECE, Orlando, Florida, Nov 5–11.

Ale, B.J.M. (2006). *The occupational risk model* TU-Delft/TBM report: 2006073, ISBN 90-5638-157-1, Delft, 2006.

Ale, B.J.M., Bellamy, L.J., Van der Boom, R., Cooke, R.M., Goossens, L.H.J., Hale, A.R., … Spouge, J. (2008). *Further development of a causal model for air transport safety (CATS): The complete model*. Paper presented at Ninth International Probabilistic Safety Assessment and Management Conference, Hong Kong, 18–23 May.

Ale, B.J.M., Bellamy, L.J., Van der Boom, R., Cooper, J., Cooke, R.M., Goossens, L.H.J., … Spouge, J. (2009). Further development of a Causal model for Air Transport Safety (CATS): Building the mathematical heart. *Reliability Engineering & System Safety*, *94*, 1433–1441.

Ale, B.J.M., Bellamy, L.J., Cooper, J., Ababei, D., Kurowicka, D., Morales, O., & Spouge, J. (2010). Analysis of the crash of TK 1951 using CATS. *Reliability Engineering & System Safety*, *95*, 469–477. doi: 10.1016/j.ress.2009.11.014.

Bellamy, L.J., Papazoglou, I.A., Hale, A.R., Aneziris, O.N., Ale, B.J.M., Morris, M.I., & Oh, J.I.H. (1999). I-*Risk: Development of an integrated technical and management risk control and monitoring methodology for managing and quantifying on-site and off-site risks* (Contract ENVA-CT96-0243). Report to European Union.

Bellamy, L.J., Ale, B.J.M., Geyer, T.A.W., Goossens, L.H.J., Hale, A.R., Oh, J., … Whiston, J.Y. (2007). Storybuilder—A tool for the analysis of accident reports. *Reliability Engineering & System Safety*, *92*, 735–744.

Burgess, D.S., & Zabetakis, M.G. (1973). *Detonation of a flammable cloud following a propane pipeline break* (Report of Investigation 7752). Bureau of Mines.

Cooke, R.M., & Goosens, L.J.H. (2000). *Procedures guide for structured expert judgement* (EUR 18820 EN, 2000). European Commission.

Dekker, S. (2002). *The field guide to understanding human error*. New York: Ashgate.

EGIG. (1982). *European Gas pipeline Incident data Group*. Retreived from https://www.egig.eu/about-egig.

Hand, D. (2014) *The improbability principle*. London: Corgi.

Hurst, N.W., Bellamy, L.J., Geyer, T.A.W., & Astley, J.A.A. (1991). A classification scheme for pipework failures to include human and sociotechnical errors and their contribution to pipework failure frequencies. *Journal of Hazardous Materials*, *26*, 159–186. doi: 10.1016/0304-3894(91)80003-7.

Hollnagel, E. (2006) *Resilience engineering*. New York: Ashgate.

Hollnagel, E. (2014) *Safety I and Safety II*. London: Ashgate.

Kjellén, U., & Hovden, J. (1993). Reducing risks by deviation control—A retrospection into a research strategy. *Safety Science*, *16*, 417–438.

Kurowicka, D., & Cooke, R.M. (2004) Non-parametric continuous Bayesian belief nets with expert judgment *(pp*. 2785–2790). Paper presented at *Proceedings of the 4th International Conference on Probabilistic Safety Assessment and Management*. New York: Springer.

Lin, P.H., Hale, A.R., Van Gulijk, C., Ale, B.J.M., Roelen, A.L.C., & Bellamy, L.J. (2008) *Testing a safety management system in aviation.* Paper presented at Proceedings of the Ninth International Probabilistic Safety Assessment and Management Conference, Hong Kong, 18–23 May.

Mayer-Schönberger, V.M., & Cukier, K. (2013) *Big Data*. London: John Murray.

Orton, J.D., & Weick, K. (1990). Loosely coupled systems: A reconceptualization. *The Academy of Management Review*, *15*, 203–223. doi: 10.2307/258154.

Perrow, C. (1984) *Normal accidents*. New York: Basic Books.

Pruppers, M.J.M., Janssen, M.P.M., Ale, B.J.M., Pennders, R.M.J., Van den Hout, K.D., & Miedema, H.M. (1997) *Accumulation of environmental risks to human health: Geographical differences in the Netherlands*. Paper presented at RISK 97, Amsterdam, October 21–24.

Saari, J. (2016) Accident deviation models. *Encyclopedia of occupational health and safety (Chapter 56)*. Retrieved from http://www.ilocis.org/documents/chpt56e.htm.

Taleb, N.N. (2007) *The black swan: The impact of the highly improbable*. London: Penguin.

Uijt de Haag, P.A.M., Ale, B.J.M., & Post, J.G. (2000) *Guideline for quantitative risk assessment: Instructions for a quantitative risk analysis in the Netherlands*. G.I. Schuller and P. Kafka (eds) Safety and Reliability. ISBN 905809 109 0, Balkema, Rotterdam, The Netherlands.

Van den Top, J. (2010) *Modelling risk control measures in railways* (PhD thesis). Delft University of Technology.