PRACTICAL APPLICATION OF BAYES' THEOREM IN ESTIMATING SAFETY SYSTEM PERFORMANCE

By Mirek Generowicz FS Expert (TÜV Rheinland #183/12)

I&E Systems Pty Ltd

Summary

Functional safety is a term used to describe the application of automated safety systems to reduce the risk associated with hazardous equipment. Functional safety relies on the effective performance of protective devices.

In this context protective devices include sensors (for instance sensing pressure, temperature, level or position), logic solvers, and final elements (such as valves, actuators, motor contactors and circuit breakers). Device failure may prevent a safe state from being achieved or maintained for the equipment under control. Device failure rates must be measured and monitored to ensure that functional safety is achieved as planned.

Some safety system practitioners promote Bayesian techniques for analysis of protective device failure rates or safety incident event rates [1-5]. For example, some suggest that Bayesian analysis can be used instead of confidence intervals to quantify the range of uncertainty in failure rates.

Bayes' theorem itself is reasonably simple: The estimated probability of an event happening can be updated if prior knowledge is available about conditions that affect the probability. More complicated techniques are derived by nesting probability estimates in hierarchical network models or by applying Bayes' theorem with probability distributions rather than with discrete values of probability.

The papers and presentations published on this topic suggest that we should make more use of Bayesian techniques in functional safety, but it is not easy to understand how that should be done and what benefits might result. At first glance it seems that complicated calculations might be necessary.

This paper reviews how Bayesian techniques might be applied in practice. It describes simple Bayesian techniques that are useful in understanding and quantifying the impact of factors that influence safety system performance.

For safety applications prior knowledge should be objective rather than subjective. Factual evidence is needed. Simple and practical rules of thumb can then be derived using Bayesian principles: If there is no evidence that the system engineering follows good practice, then failure rates and failure probabilities could be expected to be at least 3 times or even 10 times worse than normal.

Bayesian techniques can be applied together with failure modes and effect analysis and with root cause analysis. These techniques enable better understanding and control of the factors that cause variability in safety system performance. When the causes of failure are understood it becomes clear that safety system failures cannot be characterised by fixed failure rates.

Complicated and detailed calculations are neither necessary nor useful in determining failure rates and uncertainty intervals, because most failure rates do not have true values that can be measured.

The IEC 61511 (Safety instrumented systems for the process industry sector) method of selecting devices and reliability data based on prior use is an example of how Bayesian principles are applied to safety systems [6]. .

The search for dependable failure performance data

Functional safety practitioners apply mathematical analysis in the design and operation of safety-related systems. They estimate the risk of hazardous events so that targets can be set for risk reduction. They estimate the probability of failure for safety functions to demonstrate that the targets for risk reduction are achieved.

Some safety system practitioners assume that equipment failures are random. That means that the equipment failure rates remain reasonably constant, and that the true values of the rates can be determined through measurement and/or analysis. They carry out complex calculations with 2 or 3 significant figures of precision based on their estimates of the true failure rate.

The conventional approach to evaluate uncertainty in the estimated failure rates is to use confidence levels that are based on the chi-squared probability distribution. Confidence levels can only be applied to purely random failures, meaning those with constant failure rate.

That fundamental assumption of constant failure rate is valid only for purely electronic equipment in controlled environments. It is not generally valid for devices in industrial environments.

Bayesian inference has been suggested as way of evaluating uncertainty in failure rates as an alternative to the use of confidence levels [1-5]. In Bayesian inference Bayes' theorem is applied to deduce properties about a population or probability distribution from information about recorded events.

The performance of any safety-related system is variable as distinct from uncertain. It is more appropriate to discuss variability rather than uncertainty, and variability cannot be predicted from probability distributions. Variability in performance depends primarily on how much effort is put into preventing preventable failures. The level of effort varies widely between different users and different applications. Failure rates and event rates always vary over at least an order of magnitude (i.e. a factor of 10). The wide range of variation is caused by systematic factors and is inevitable [7].

Understanding randomness

To understand the difference between uncertainty and variability in safety system performance we need to review the nature of random behaviour.

Failure rates are constant if the failures result from a purely random process.

Purely random processes involve a **series of completely independent events**. These processes have no memory and no dependencies. They are classed as **stochastic** processes. The term stochastic means that behaviour can be described by a random probability distribution. The probability of each event is predetermined by some inherent physical constraint, characteristic or process. The probability does not change with time and does not depend on external influences.

Many processes appear to be random, but few are purely random. Most processes are **dependent** rather than purely random.

Appendix A discusses a series of football matches as an example of a dependent process that appears to be random but cannot be characterised by constant rates.

The tossing of a coin, the rolling of dice or the spinning of a roulette wheel are examples of purely random processes. The probability of tossing a head or rolling a pair of sixes depends on the symmetry of the coin or dice. The probability remains constant over time.

The probability of tossing a head with a given coin can be estimated by dividing the long-term cumulative total of heads by the number of tosses.

Games like football matches are not purely independent events. The outcome of a match depends on many factors. A football team's average win-rate can be measured to gauge its success. Over longer terms the average win-rate will always appear to be reasonably constant. That might lead to the false impression that the behaviour is as random as the tossing of a coin. A closer examination of the variation in performance from match to match reveals that the win-rate is not fixed. There is no reason why football team win-rates would ever be constrained to some inherent constant value.

The probability of a football team winning any individual match has no relationship with its long-term average win rate.

Confidence intervals for random failure rates

Safety system designers need estimates of equipment failure rates so that the probability of system failure can be calculated.

If a failure process is purely random the long-term average rate will reveal the true value of the constant failure rate. Short-term averages can be used as an approximation. The error in the approximation reduces as the number of measured events increases.

The chi-square function allows an event rate in a random process to be estimated with any required level of confidence.

A rate estimate at a 95% confidence level means that there is only a 5% chance that the true rate will be higher than the estimate.

A confidence interval can be used to estimate the expected range of error in a measured rate. For example, a 90% confidence interval is the range between an estimate at the 5% confidence level and the estimate at the 95% confidence level. There is a 10% chance that the long-term average (i.e. the true value) of the rate will be outside the 90% confidence interval.

The width of the confidence interval depends only on the total number of events measured.

Appendix B shows how the chi-square function can be used to determine a confidence interval for a purely random process but not for a dependent process.

The chi-square function can accurately predict the long-term average of a coin toss trial but not for a series of football matches. There is a true value for the probability of tossing a head. There is no true value for the probability of a team winning a match.

Randomness of failures in safety-related systems

The probability calculation methods described in the functional safety standards IEC 61508-6 [8] and ISO/TR 12489 [9] are intended for random failures. The calculations all assume that failure rates remain reasonably constant.

The calculations are intended for use in the design of systems and devices assembled from electronic components. ISO/TR 12489 Annex B explains that only electronic components are subject to purely random failure and constant failure rate (refer to ISO/TR 12489 Figure B5). All other components have variable failure rates.

The assumption of constant failure rate is justified for electronic components in controlled environments. Reasonably constant failure rates can be measured for electronic devices in a factory test environment.

Appendix C below explains why the assumption of constant failure rate might not be valid for electronic devices and systems installed in uncontrolled environments.

In short, most failures (typically at least 95%) are systematic rather than purely random in nature. Systematic failures can be prevented, avoided or controlled to some extent. The probability of systematic failure on demand can be estimated (within half an order of magnitude) but systematic failures do not occur at predictable rates. Systematic failures are related to pre-existing faults or degradation.

The overall failure rates of electronic devices and systems are not fixed and constant because they are partially systematic. The failure rates of their component elements depend on the operating environment, on the specific application, and on the condition of the component. The failure rate of each component can vary over at least an order of magnitude [7].

Failures in non-electronic components do not occur at fixed rates because the failure mechanisms are all dependent on design and on environmental factors.

Whether electronic or non-electronic, no devices in industrial applications have fixed failure rates with true values that can be measured. It is always possible to measure failure rates, but the rates are variable.

The measured failure rates and event rates reflect the past performance of safety systems. Performance can be changed through action and inaction, whether deliberate or inadvertent.

Confidence intervals based on the chi-square function cannot be applied to devices in operation. The chi-square function is applicable to uncertainty in measurement only of parameters that have a true constant value. It may be useful for testing of devices in a controlled test environment.

Uncertainty and variability

OREDA (Offshore Reliability Data) was established in 1981 to analyse failure statistics collected by a consortium of offshore hydrocarbon producers [10, 11]. A series of OREDA handbooks has been published over several decades to summarise the failures recorded in successive time periods. The results were first published in 1984. The work was extended to include onshore statistics in the 6th edition, published in 2015.

The statistics collected over the past 40 years show that failure rates of similar equipment vary widely between different users and different applications. Variation in performance spans at least one order of magnitude between different users. The variation spans over three orders of magnitude for some types of equipment. Performance also varies over time.

OREDA recognises that confidence intervals cannot be applied to datasets combined from different users because each user is measuring a different failure rate. OREDA uses uncertainty intervals instead of confidence intervals.

The definition of the term 'uncertainty' varies. ISO 14224 [12] defines uncertainty of a quantity as:

'the inability to determine accurately what is or will be the true value of a quantity'

This definition limits the application of uncertainty to parameters that have a **true value.**

ISO 14224 includes a note clarifying that uncertainty can have different meanings. It explains that uncertainty can mean stochastic uncertainty or epistemic uncertainty. Stochastic uncertainty results from random processes and is characterised by probability distributions. Epistemic uncertainty can result from a lack of knowledge or understanding of the process or can result from inaccuracy in a mathematical model.

It may be better to use the term 'variability' instead of 'uncertainty' to describe the range over which dependent parameters vary.

Estimating uncertainty intervals with Bayesian analysis

Appendix D below provides an introduction to Bayes' theorem.

Unlike confidence intervals, the application of Bayesian analysis is not limited to parameters that have fixed true values. Bayesian analysis can be applied with any probability distribution [13, 14]. It can be used with distributions that are characterised by several parameters rather than only by mean and variance.

Bayesian analysis can be applied to estimate the relative likelihood that selected values of parameters accurately represent the behaviour observed in a trial.

Bayesian analysis is based on some prior knowledge regarding the expected range of parameters.

In general applications prior knowledge can be objective or subjective. It can be based on theoretical analysis or on previously measured performance. Prior knowledge might also be based on hypothesis, inference or conjecture, though that would not be acceptable in safety applications.

The expected distribution of values is the 'prior' distribution. Recent performance of the system is evaluated using the prior parameter values to calculate how likely that performance would have been if that set of parameters were valid. This yields an updated estimate for the likelihood of possible parameter values. The updated set is called the 'posterior' value or distribution.

With no prior knowledge, a flat distribution can be used as the prior distribution. This is called a weak prior or an uninformed prior. It assumes that all values of the failure rate are equally likely. Upper and lower bounds of an uncertainty interval or a range of variability can then be inferred from the resulting posterior distribution.

Appendix E shows how Bayesian analysis can be applied to a simulated coin toss trial as an example of a random process and to football match results as an example of a dependent process. The range of variation or uncertainty observed between sets of coin toss trials and between a football team's performance in different seasons appears to be similar. There are no obvious differences between the random process and the dependent process.

The effectiveness with which Bayesian analysis reveals variability depends on the choice of the prior distribution.

This type of analysis needs to be treated with caution because a strong prior or a biased prior (rather than a flat prior) will skew the posterior results to reinforce the assumption made in the prior - even if that assumption is not valid. Appendix E includes an example of a biased prior that assumes that the probability of tossing a head is 0.6 rather than 0.5. After only 20 tosses the analysis seems to confirm that 0.6 is the true value. The posterior distribution will move towards the true value as more information is accumulated from measurements.

The prior distribution might be updated successively from previous posterior distributions to find the 'true value' of the rate. The successively updated distributions 'shrink' to reveal the long term mean of the rate.

The range of variability in industrial equipment failure rates is already well known and the reasons for variability are understood. Measured failure rates vary because performance varies. Uncertainty in measurement is not relevant in this context.

Posterior distributions derived from Bayesian analysis do not provide further insight. The posterior distributions can identify the long-term mean, but that mean value is not a true value for the failure rate. Most failure rates do not have true values that can be measured.

This type of analysis requires complicated calculations and is unlikely to be of value in safety systems.

Long-term means are stable even with variable performance

With either frequentist and Bayesian methods the calculated confidence interval or range of variability becomes narrower as more measurements are taken, even if the rate is not truly constant.

Both methods will produce a similar estimate of the long-term average rate (i.e. the arithmetic mean). It does not matter whether the rate is fixed and constant or varying widely between trials.

A mean rate can always be measured, but in a dependent process that is only a measure of past performance. We cannot assume that the future performance will be consistent with past performance unless the performance is constrained to a fixed rate by some fixed physical, chemical or biological characteristics.

Long-term means measured in constrained processes will be close to the true values of the parameters that determine performance.

All mean values will become progressively more stable as the number of samples increases. Failure performance measured over many failures appears to be random because long-term mean event rates appear to be constant.

The search for patterns

Humans are naturally inclined to search for patterns. We ascribe meaning and significance to the patterns that we find even when the patterns have no meaning [15].

Fenton and Neil explain some of these concerns in '*Debunking Bad Statistics*', chapter 2 of their book '*Risk Assessment and Decision Analysis with Bayesian Networks*' [13].

Many safety system practitioners assume that failure rates are reasonably constant because the long-term means are reasonably constant. They assume that the true values of device failure rates can be determined through measurement and/or analysis.

The fact is that there are no reasons that would ever cause or constrain failure rates in safety-related systems to be constant. Failure rates always vary over at least an order of magnitude because of systematic factors [7]. Belief in constant failure rates is an example of 'anchoring bias'. People assume and believe that failure rates are constant because that is what they were initially taught. We tend to cling tenaciously to what we learnt, despite finding overwhelming conflicting evidence.

Complex mathematical modelling and analysis are attractive because they give the impression of precision and accuracy. In practice our models are usually imperfect. The statistician George Box summarised it well: *'Essentially, all models are wrong, but some are useful'* [16].

We need to be careful not to be drawn into unnecessarily complex analysis and not to read too much into the results of any analysis. The purpose of the analysis is to inform our decision making rather than to predict the future precisely and accurately.

Predicting future behaviour

Techniques such as Monte Carlo analysis or Bayesian analysis might be applied to predict the overall behaviour of large populations if the probability distributions remain reasonably stable over time.

Analysis cannot provide precise and accurate predictions of the future behaviour of individuals, but it facilitates a better understanding of causes and effects. The analysis enables decision makers to identify influencing factors and to understand their potential effects.

These same techniques cannot be used reliably to predict future performance of individual units within large populations or over limited time periods if the processes are dependent (systematic) rather than purely independent (random).

The example of a football team in Appendix E below shows that complex Bayesian analysis does not necessarily provide insights into the factors that influence the overall win rate of an individual football team. The analysis shows the range of variation in win rate, but the variation can also be observed by simply comparing short term averages with the long-term average.

Similarly, complex Bayesian analysis based on probability distributions is not useful in predicting the performance of safety systems.

It might be possible to formulate theories linking performance to specific factors. Bayesian networks might then be used to examine how well the theories predict performance [13]. Bayesian networks join individual Bayesian estimates together. The prior probability for one Bayesian estimate is taken from the posterior probability of a previous estimate.

For example, a football team's win rate might depend on whether they play at their home ground or away. The composition of the team could be considered. Skill, experience and health all affect performance. If Liverpool were to play Southend, we would usually expect Liverpool to win. A probability of winning could be estimated given knowledge of factors such as the players selected for each team. The estimate might be reliable, but we would not expect the probabilities to be estimated with precision.

Bayesian techniques in epidemiology

Epidemiological studies provide an interesting example of how Bayesian techniques can be applied in processes that are not random. [17-20].

Bayesian techniques can be applied to analyse health and disease in large populations. The behaviour of a pandemic disease will vary widely between different populations and over time depending on the effectiveness of preventive and remedial actions.

Analysis and comparison of different populations allows researchers to gain understanding of the relative effectiveness of different strategies for disease prevention and treatment.

Probability distributions can be used to model the relationships between various risk factors and morbidity or mortality.

Bayesian networks may be useful in assessing which mathematical models are more likely to be accurate. The models can be useful in demonstrating how morbidity and mortality can reduced by deliberately changing behaviours to minimise risk factors.

Bayesian networks can be useful when the relationships between causes and effects are complex and dependant on many variables if the relationships can be represented in a network model [20].

Variability in safety-related system failures

It is important to understand that the performance of safety systems is never determined by a failure rate or by any other set of probability distribution parameters.

Probability distribution parameters such as failure rates are simply indicators of historical safety system performance.

Performance can always be changed through action or inaction, whether deliberate or inadvertent.

Performance depends on the strategies and resources applied in designing, manufacturing, installing, operating and maintaining the systems.

Estimating variation in failure rate using Bayes' theorem

A Bayesian approach can be relatively simple and straightforward. It can be combined with conventional frequentist methods (i.e. statistical inference based on the measured frequency of recorded events).

For example, ISO 14224 C.3.3 [12] suggests a simple Bayesian approach to estimate the failure rate of a device when no failures have been measured.

A limiting value for the expected failure rate can be estimated from the total number of device-hours in fault-free service: $\lambda \approx 0.7/t$.

Though a failure rate cannot yet be measured, the information about the length of fault-free service enables a prediction to be made of the likely range for failure rate. The estimated value will continue to reduce as the number of fault-free hours increases.

Bayes' theorem estimates the probability of an event based on prior knowledge of conditions that might be related to the event.

Appendix F below has two examples of how Bayes' theorem can be applied to estimate the expected increase in failure rates if equipment is not maintained promptly and effectively. The examples show that probability of failure can typically be reduced by a factor of 3 or more if sub-standard equipment is found by inspection and corrected before it fails.

The conclusions from these examples are consistent with the findings reported by Bukowksi and Stewart in their paper *Quantifying the Impacts of Human Factors on Functional Safety* [21]. They showed that failure rates may increase by at least a factor of 3 or 4 if maintenance is not done effectively and promptly.

Failure rates depend not only on maintenance effectiveness but also on the suitability of the system design and system components for the environment.

The concept of prior use in IEC 61511 (ANSI/ISA S84.00.01) applies a Bayesian approach [6]. Device failure rates assumed in the design of safety functions are based on failure performance measured in prior use of that same type of device in a similar operating environment.

Bayes' theorem applied to LOPA

Layers of Protection Analysis (LOPA) is a method used to evaluate the levels of risk reduction needed for high-consequence hazardous scenarios. The risk is evaluated as the combination of probability of

occurrence and severity of consequences. The risk might be reduced through independent, dependable and reliable protection methods such as pressure relief valves. Further risk reduction is required if the evaluated risk exceeds a company's target for tolerable risk.

Subjective or objective prior knowledge can be applied in hazard and risk studies to guide users in estimating future event frequencies. Guide phrases might be used, for instance:

- *'known to have happened once or twice at our plant in the last 10 or 20 years'*
- *'typically occurs once or twice per year per one hundred devices in similar service'.*

Prior knowledge can be used in estimating the probability of failure of protective equipment such as pressure relief valves. There needs to be some evidence to support the assumptions made.

A good starting point is the Center for Chemical Process Safety (CCPS) handbook *Guidelines for Initiating Events and Independent Protection Layers in Layer of Protection Analysis*' [22]. It provides guidelines on the risk reduction credits (to the nearest order of magnitude) that can be taken for independent protection layers.

This is essentially a Bayesian approach because the expectation of future performance is informed by experience, even though the event frequencies vary widely and can never be known with certainty.

Grattan and Brumbaugh [1] suggest that '*simple Excel™ based Bayesian calculations* [can be used] *to address issues such as uncertainty, establishing confidence intervals, properly evaluating LOPA gaps, and incorporating site specific data, all related to IPLs and barriers used to meet LOPA targets.'*

Calculations are useful if data are available but the precision in the calculations can never be better than +/- half an order of magnitude. We could perhaps make use of prior knowledge in a heuristic way, without necessarily requiring calculations. Given the subjectivity and variability in the prior knowledge we might as well simply apply multiplying factors such as times 3 or times 10.

The assumptions made in taking risk reduction credit for protection layers always need to be justified. The CCPS guidance cannot be taken for granted. Multiplying factors cannot be based on opinion without factual evidence. **We need some evidence to show that those failure rates and probabilities are feasible in the target environment**.

Appendix G describes two hypothetical cases of how a heuristic Bayesian approach might be adopted in LOPA. Examples are given of the evidence that is needed to support assumptions made in LOPA.

Bayesian techniques can be combined with failure analysis

Safety system devices fail for many different reasons. Each component of each device may be susceptible to several different failure modes.

Purely random failures are not preventable. The rate of failure can be expected to be reasonably constant for each of the purely random failure modes of each component.

Confidence intervals or uncertainty intervals may be useful in determining the most likely range for the value of a constant failure rate if enough data can be collected.

A large volume of operational experience is necessary. Individual electronic components typically have failure rates in the range between 5 to 50 failures per 100,000 device-years (or approximately 10^9 device-hours). This estimate includes failures of all types and all modes. The failure rate for each purely random mode of failure can be expected to be in the order of 1 failure in 100,000 device-years [7].

At least 3 or 4 failures of each type need to be measured to provide a reasonably accurate estimate for the failure rate of an individual component. A volume of experience including around 300,000 device-years might be sufficient if the failure rate were in the order of 1 failure in 100,000 device-years.

The overall total number of failures for any type of composite device will include failures of its different components, each with perhaps several different modes and different causes.

It might not be practicable to collect enough data for analysis. A Bayesian approach can still be useful because it can be used to estimate failure rates for composite devices without the need to measure any failures of that device.

A Bayesian approach might use knowledge of the expected failure modes and failure effects of the component parts. For example, failure mode effects and diagnostic analysis (FMEDA) uses knowledge of a device's component parts to predict failure modes and failure rates for the overall device.

Most safety system devices include many non-electronic components. Typical safety system devices have overall failure rates in the range of between 0.3 to 3 failures per 100 device-years (or approximately 10^6 device-hours) [7-11, 23-25]. These failure rates are frequent enough to be measured with relatively small populations (e.g. 100 devices over 3 years), but the failure rates vary over at least an order of magnitude depending on the effectiveness of maintenance [21] as well as on the suitability of the devices for the operating environment and on quality in design and manufacturing [7].

Statistical techniques can be used to analyse failure performance of safety-related devices over large populations of devices or over extended periods of time. The overall performance measured provides an indication of the performance that can be achieved in practice. The variance in the measurements indicates the degree of variability in performance between different applications and different operating environments.

Information gathered from consolidated datasets of failure statistics (such as OREDA [10, 11] and FARADIP [7]) shows the performance that can typically be achieved and the range of variation that can be expected. The variability in performance usually spans several orders of magnitude.

FMEDA results published by *exida* in the Safety Equipment Reliability Handbooks (SERH) [23 and 24] are consistent with the OREDA statistics. The *exida* SERH handbooks present failure rates that 70% of users could be expected to achieve when applying recognised and generally accepted good engineering practices for design, installation and maintenance.

Failure data collected from prior use of devices in a specific operating environment gives the most useful indication of the performance that can be expected in that environment. IEC 61511 prefers data based on field feedback from prior use in a similar operating environment [6]. Sub-clause 11.5.2.2. of the standard notes that data available from manufacturers may not be valid in all applications. Prior use data can be combined with industry databases to inform decisions made in the design, operation and maintenance of safety systems. The evidence for reliability must be credible, traceable, documented and justified.

FMEDA and prior-use data are both examples of how Bayesian thinking can be applied in the selection of appropriate failure rates in the design of safety functions.

Some failures will inevitably occur during operation of a safety system. Most failures will have specific causes that can be found and explained. The priority is always to determine and eliminate those causes rather than to produce precise estimates of failure rate.

Not all failures can be easily eliminated if time and resources for maintenance are constrained.

Root cause analysis [12, 26, 27] is always appropriate for failures of safety system devices because it is useful even with only one failure. Root cause analysis should be conducted on every individual failure to determine if the failure has causes that can be corrected to reduce the likelihood of future failures. Remedial action can be taken without waiting to count further failures.

Root cause analysis of each failure can distinguish purely random failures from those of a deterministic nature with specific causes that can be remedied.

The value of remedial action can be assessed by considering the cost of the remedial action in comparison with the consequences of failure.

A Bayesian approach can be used to apply knowledge gained from root cause analysis of failures. Estimates of failure rate or failure probability can be modified accordingly if the causes of the failure cannot be remedied at a proportionate cost.

A simple Bayesian approach can be adopted to account for factors that are known to affect failure rate or failure probability (such as quality management and maintenance effectiveness).

For example, the work by Bukowksi and Stewart [21] can be applied in a Bayesian approach without the need for detailed analysis. Bukowksi and Stewart proposed a simple index to categorise the effectiveness of maintenance practices at individual process plants. Bukowksi and Stewart found that failure rates can be higher than normal by up to a factor of 4 if maintenance is not done effectively and promptly.

Equipment failure commonly results when deterioration is not detected or corrected promptly. Ineffective maintenance is known to cause increased failure rates.

Safety applications require objective evidence

Designers and operators of hazardous equipment have a duty of care in managing hazards. Duty of care requires demonstration of due diligence: keeping evidence that reasonable efforts are made to apply appropriate standards and practices, and monitoring the effectiveness of those efforts.

Safety systems are applied as part of that duty of care to reduce risk associated with hazardous equipment.

In non-safety applications Bayes' theorem can be applied with prior knowledge that is either objective or subjective, but in safety applications prior knowledge needs to be as objective as far as practicable, so that due diligence can be demonstrated.

Some degree of subjectivity is inevitable but subjective opinions need to be supported by evidence, such as evidence of:

- Safety and quality management systems in compliance with the appropriate standards
- Suitability of devices and systems for the application
- Maintenance effectiveness
- Failure analysis and appropriate corrective actions.

Conclusions

Bayesian techniques are essential in the design and operation of automated safety systems, but complex calculations are not necessary or useful.

The assumption of constant failure rate is not valid. Failure rates and event rates in safety-related systems are never fixed and constant. The rates always vary over at least one if not two or three orders of magnitude.

Confidence intervals and uncertainty intervals have limited application in safety systems. In practice failure rates are variable rather than uncertain. Failure rates are not constrained to true values.

Lack of precision in failure rate data limits the validity of complex calculations. The precision in predictions of failure rate and probability is necessarily limited to +/- half an order of magnitude at best.

A simplified Bayesian approach is more appropriate. Probability of failure depends on the condition of the equipment and can be changed through deliberate action.

A simple Bayesian approach can be taken by using failure rates that are based on OREDA, FARADIP and SERH. Those failure rates are achievable provided that the engineering of the system follows recognized and generally accepted good engineering practices.

If there is no evidence to prove that the system engineering follows good practice, then the failure rates and failure probabilities could be expected to be at least 3 times or even 10 times worse than are achieved with good practice.

Failure rates can be reduced deliberately through consistent application of reliability-centred maintenance [28]. OREDA statistics show that some users report failure rates that are 3 times or even 10 times lower (better) than normal good practice.

According to IEC 61511, such low failure rates could only be assumed in the design if objective documentary evidence were available from prior use in the intended operating environment.

IEC 61511 requires that the reliability data '***shall be credible, traceable, documented, justified and shall; be based on field feedback from similar devices used in a similar operating environment'***. The evidence must include demonstration of the quality management and configuration management of the devices.

The application of reliability data from prior use in IEC 61511 is an example of Bayesian thinking, but it requires objective prior knowledge.

References

TABLE 1. REFERENCE DOCUMENTS

Ref	Title
1	Grattan, D. and Brumbaugh, K. *'Reverend Bayes, meet Process Safety - Use Bayes' Theorem to establish site specific confidence in your LOPA calculation'.* aeSolutions
2	Thomas, S. '*A Hierarchical Bayesian Approach to IEC 61511 Prior Use*'. Presented at the American Institute of Chemical Engineers 2018 Spring Meeting and 14th Global Congress on Process Safety. <https://sisengineer.com/hierarchical-bayesian-paper/>
3	Thomas, S. *'Uncertainty and Basic Bayesian Inference for SIS'* Published by SISEngineer.com, 2019 <https://youtu.be/AcATkRuOe94>
4	Lennox, K. '*All About that Bayes: Probability, Statistics, and the Quest to Quantify Uncertainty*'. Presented at Lawrence Livermore National Laboratory, 2016 <https://youtu.be/eDMGDhyDxuY> .
5	Recchia, C. *'Bayesian Methods in Reliability Engineering'* Presented as a webinar for the American Society for Quality, 2012 <https://vimeo.com/158000548>
6	IEC 61511:2016 *'Functional Safety – Safety instrumented systems for the process industry sector'*
7	Smith, D. J. *'Reliability, Maintainability and Risk'*, 9th Ed. Butterworth Heinemann. 2017
8	IEC 61508-6: 2010 *'Functional safety of electrical/electronic/programmable electronic safety-related systems Part 6 Guidelines on the application of IEC 61508-2 and IEC 61508-3'*
9	ISO/TR 12489: 2013 '*Petroleum, petrochemical and natural gas industries — Reliability modelling and calculation of safety systems'*
10	OREDA '*Offshore Reliability Data Handbook'* Volume 1, 5th Ed. SINTEF. 2009
11	OREDA '*Offshore and Onshore Reliability Data Handbook*' Volume 1, 6th Ed. SINTEF. 2015
12	ISO 14224: 2016 '*Petroleum, petrochemical and natural gas industries — Collection and exchange of reliability and maintenance data for equipment'*
13	Fenton, N. and Neil, M., '*Risk Assessment and Decision Analysis with Bayesian Networks'*, 2nd Ed. CRC Press, Taylor & Francis Group, 2019
14	Atwood C.L., LaChance JL., Martz H.F., Anderson, D.L Englehardte, M. Whitehead, D. and Wheeler, T. NUREG/CR-6823 (SAND2003-3348P) '*Handbook of Parameter Estimation for*

Ref	Title
	Probabilistic Risk Assessment' Sandia National Laboratories / U.S. Nuclear Regulatory Commission Office of Nuclear Regulatory Research Washington, DC, 2003
15	Shermer, M. 'Patternicity: The tendency to find meaningful patterns in meaningless noise'. Scientific American. 299. 48. 10.1038/scientificamerican1208-48. <https://www.scientificamerican.com/article/patternicity-finding-meaningful-patterns/>
16	Box, G. E. P. '*Science and statistics*', Journal of the American Statistical Association 1976, 71 (356): 791–799
17	MacLehose, R.F. and Hamra, G.B. '*Applications of Bayesian Methods to Epidemiologic Research'.* Springer International Publishing AG 2014
18	Greenland, S. '*Bayesian perspectives for epidemiological research: I. Foundations and basic methods'* International Journal of Epidemiology 2006 Vol. 35 pp 765-775
19	Dunson, D.B. *'Commentary: Practical Advantages of Bayesian Analysis of Epidemiologic Data'* American Journal of Epidemiology 2001 Vol. 153 No 12
20	Fenton, N., '*Diagnostic Testing: impact of confirmatory testing on false positives using Bayesian networks*' Queen Mary University of London, 2021 <https://youtu.be/GLnTC4LLLLA>
21	Bukowski, J.V. and Stewart, L. '*Quantifying the Impacts of Human Factors on Functional Safety' exida.* Presented at the American Institute of Chemical Engineers' 12th Global Congress on Process Safety, Houston, Texas. 2016
22	*'Guidelines for Initiating Events and Independent Protection Layers in Layer of Protection Analysis'*, Center for Chemical Process Safety, Wiley, 2015
23	*exida* '*Safety Equipment Reliability Handbook*' (SERH), 3rd Ed. 2007
24	*exida* '*Safety Equipment Reliability Handbook*' (SERH), 4th Ed. 2015
25	*'Failure Rates for Process Industry Applications'* <silsafedata.com>
26	*'FactSheet: The Importance of Root Cause Analysis During Incident Investigation'.* OSHA; EPA.<https://www.osha.gov/Publications/OSHA3895.pdf>
27	*'Guidelines for Investigating Chemical Process Incidents',* Center for Chemical Process Safety, 2nd Ed., Wiley, 2003
28	Moubray, J., *'Reliability-Centered Maintenance RCM 2.1'*, 2nd Ed. Butterworth-Heinemann, 1999
29	*exida 'Site Safety Index'* <exida.com/SSI)>

Appendix A

Distinguishing purely random behaviour

The following example was produced by simulating the toss of a coin by using the random number function in Microsoft Excel.

The probability of heads in this simulation was set to 0.5. The outcome of heads or tails in each toss is purely random.

The average number of heads per toss stabilises as the number of tosses increases. After about 1 000 tosses the long-term average is approximately equal to the probability of tossing a head.

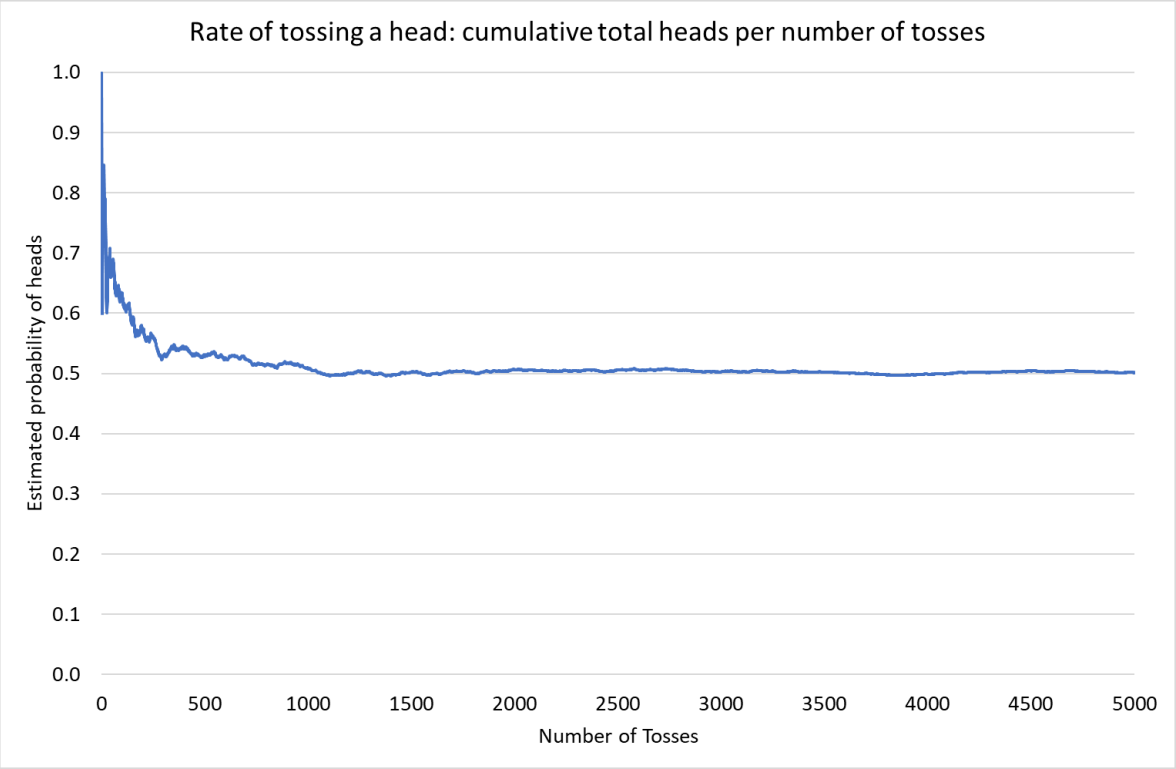

FIGURE 1. Long term average of a coin toss trial

The short-term average in any series of trials will always vary. In 50 sets of 100 tosses the short-term average proportion of heads varied from 0.33 to 0.6, a range of 1.8 : 1 (i.e. spanning one quarter of an order of magnitude. i.e. roughly $10^{0.25}$: 1).

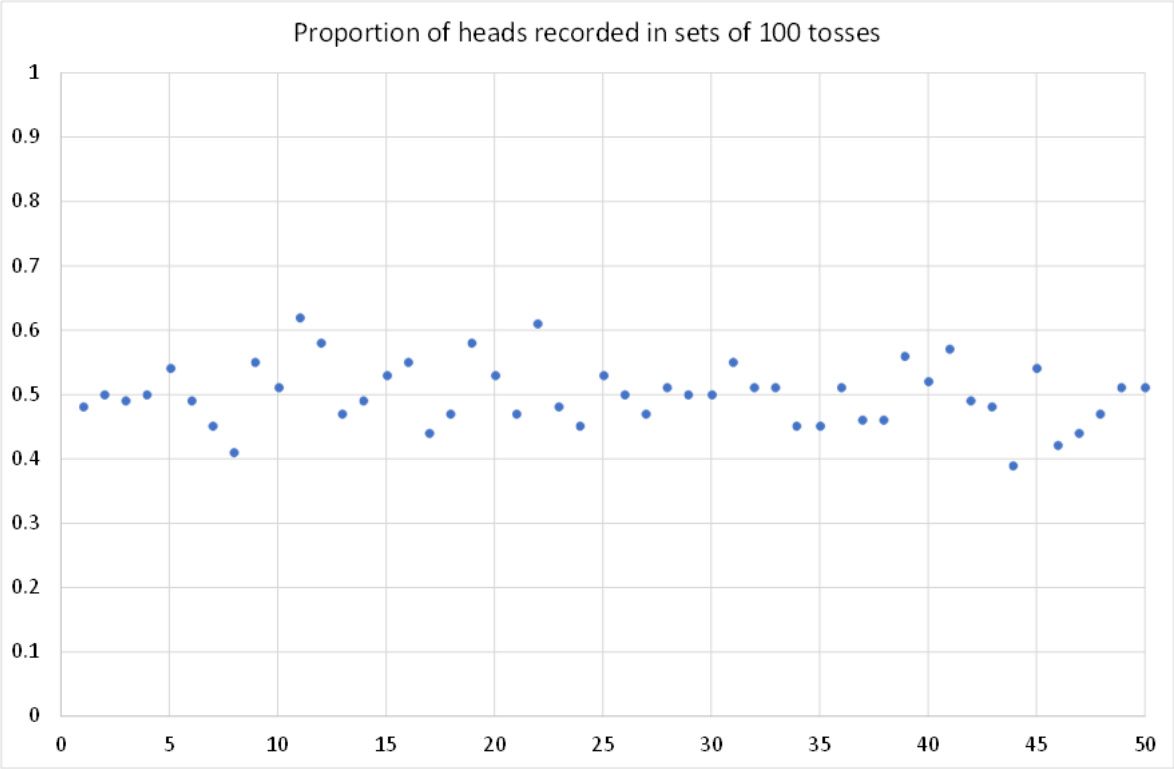

FIGURE 2. Variability in short term averages of coin toss results

Many dependant processes ***appear*** to have some degree of randomness.

Any sufficiently large collection of data can be modelled by a probability distribution as if it were purely random. We cannot assume that the behaviour is random just because we can characterise it by using probability distributions. Finding a random probability distribution to fit the behaviour does not make the process stochastic.

The performance of a football team is a useful example to consider. The performance of a football team can be measured by the rate at which it wins matches.

Obviously, performance in football is not purely random. It depends on factors such as:

- Availability of resources and equipment
- Skill, ability and experience of the team players
- Leadership and coaching
- Environment and weather conditions
- Ability of the opposing team
- Fairness of the competition – being free of undue influences, corruption or negligence.

Measuring the long-term cumulative totals of games and wins reveals the average win rate. Results measured over long periods may appear to be consistent with purely random processes, such as the rate of tossing a head.

The long-term average becomes more stable as the size of the dataset increases. The long-term average is not related to the probability the team winning any individual match.

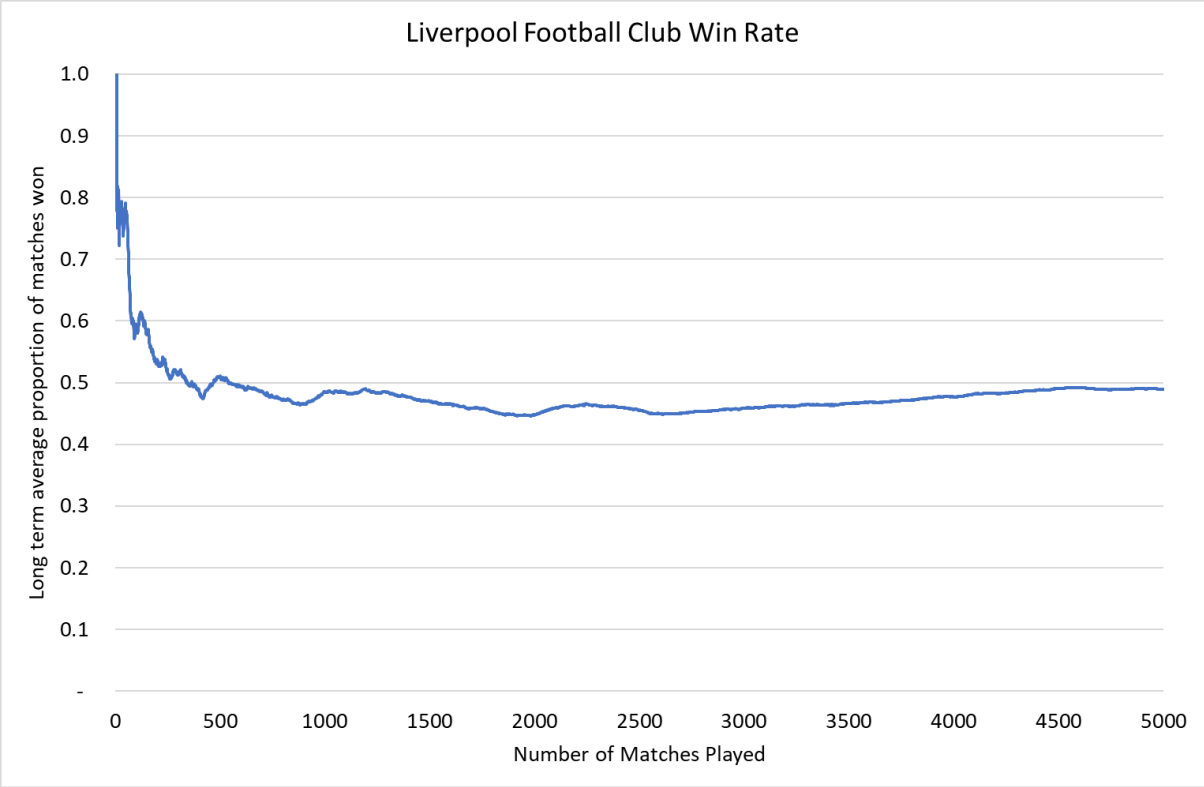

FIGURE 3. Long term average win rate – Liverpool Football Club from 1892 to 2000

The Liverpool Football Club (LFC) win rate shows more variation from season to season than would be expected in a purely random process (https://www.lfchistory.net/SeasonArchive/BySeason).

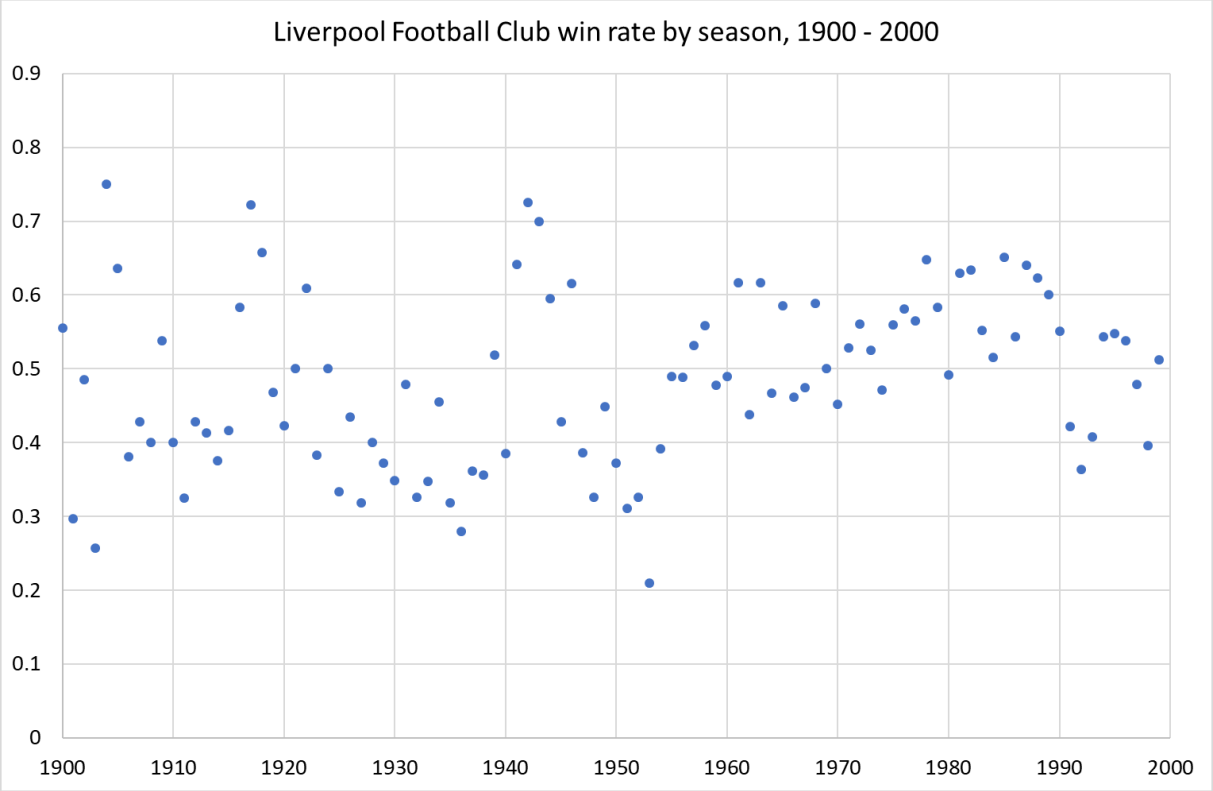

FIGURE 4. Variability in the Liverpool Football Club's win rate from season to season

The range of **variability** becomes evident from comparison of short-term average rates with the long-term average rate.

The proportion of matches that the LFC won in 5085 games between 1892 and 2000 varied from 0.21 to 0.78, a range of 3.7 : 1 (i.e. spanning more than half of an order of magnitude, $10^{0.57}$: 1).

The range of uncertainty or variability in the rate estimate can be assessed with either frequentist or Bayesian methods.

Appendix B

Confidence intervals have limited applicability

The chi-square function can be used in a frequentist approach to determine a confidence interval for the estimate of a constant rate that characterises behaviour of a purely random process.

The basic assumption is that the event rate is reasonably constant. Times to events will be exponentially distributed. If we model the event rate with the constant 'λ', the cumulative probability distribution function for time to event is $1- e^{-\lambda t}$.

The chi-square function may then be used to estimate confidence intervals for variance from the mean time to event [7]. The chi-square distribution is the distribution of the sum of squared standard Normal deviates. The mean and standard deviation of the probability density distribution of times to events can be inferred from the number of events in the number of trials (i.e. hours or cycles).

If it is true that the event rate is constant, then there is a 90% chance that a 90% confidence interval will include the true value of the event rate.

Confidence intervals for a series of coin toss trials confirm that in at least 90% of the trials, the 90% confidence interval includes the long-term average rate of 0.5 heads per toss.

In a set of 50 trials conducted for this study, only 1 out of 50 of the confidence intervals excluded the long-term average rate.

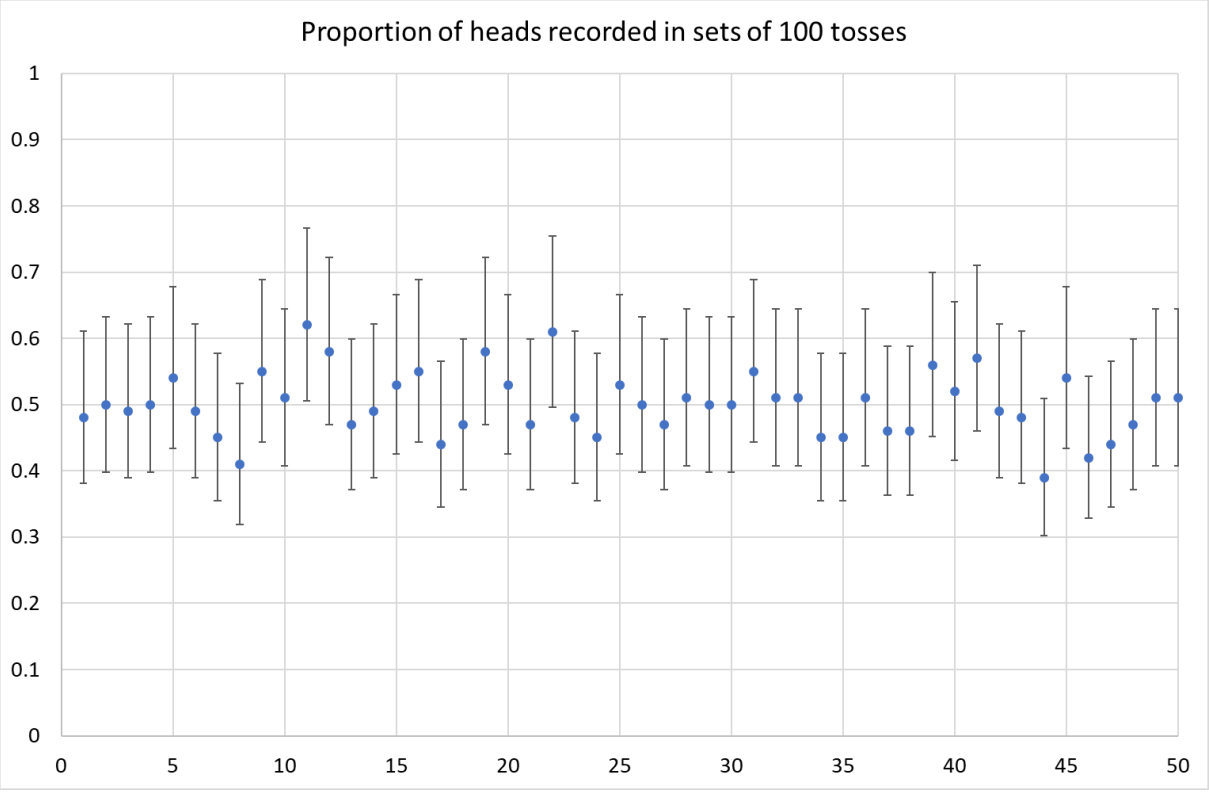

FIGURE 5. Short term averages of coin toss results with 90% confidence intervals

Over 5000 trials 2486 heads were recorded, giving a long-term average of 0.497 heads per toss.

The span of the confidence interval depends only on the number of events measured. The 90% confidence interval for the rate of heads calculated over all 50 sets of 100 tosses spans from 0.481 to 0.514. The long-term confidence interval is evenly distributed around 0.5, and it is consistent with

each of the confidence intervals for the individual sets of 100 tosses. It could be concluded that the Microsoft random number generation algorithm is effective in producing random numbers.

The long-term average rate of heads provides a dependable prediction of the rate of heads in any set of 100 tosses.

Football matches show a wider range of variation. Long-term average win rates are not so dependable in predicting future performance.

Over 5045 matches played between 1892 and 2000, the Liverpool Football Club scored 2468 wins. The club's long-term average was 0.489 wins per match.

The 90% confidence interval for the rate of wins calculated over all 5045 matches spans from 0.473 to 0.506, centred evenly around the long-term average.

The long-term average does not serve as a good prediction of performance within each season. The club's performance from season to season varies more widely than would be expected if the win rate were constrained by some underlying mathematical relationship.

The confidence intervals calculated for each season played by the Liverpool Football Club show that only 80% of the 90% confidence intervals include the long-term average rate.

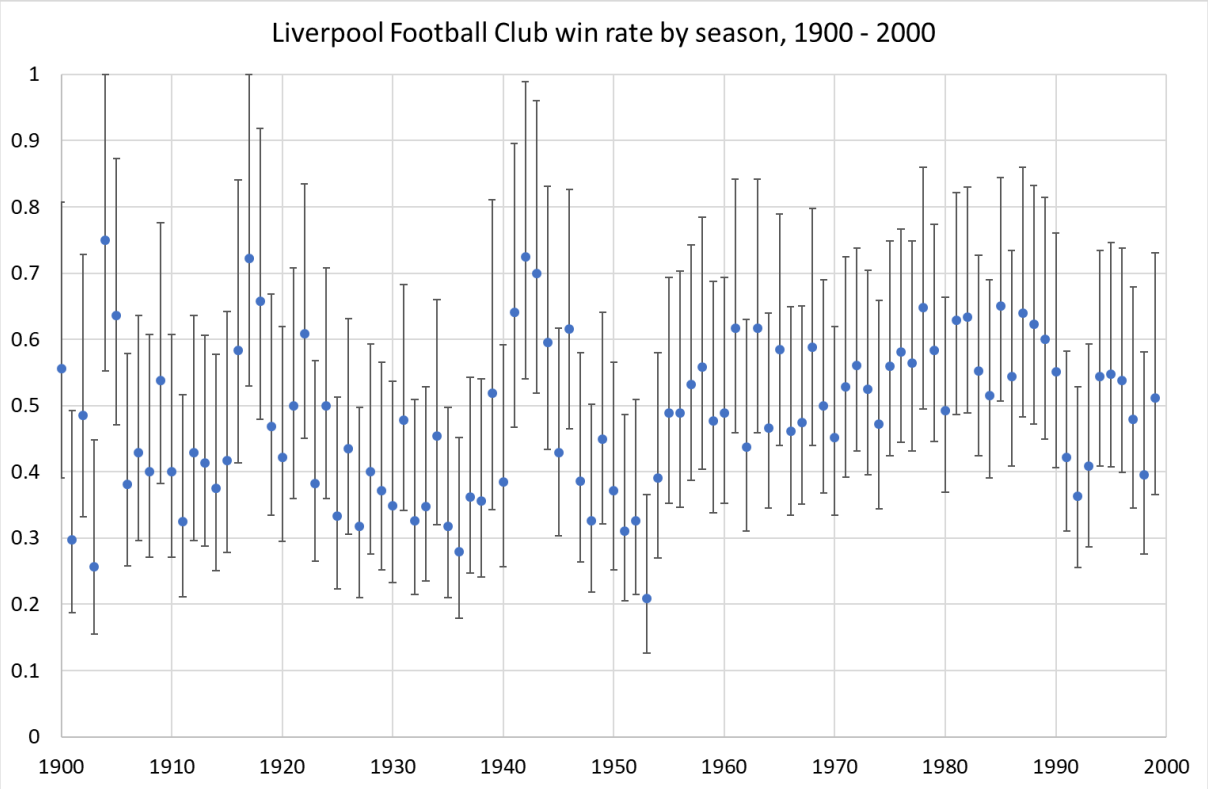

FIGURE 6. Liverpool Football Club seasonal win rates with 90% confidence intervals

Performance varies widely between teams as well as between seasons. Clubs such as Norwich or Southend have lower win rates.

The assumption of a constant win rate is not valid. There is no reason why any individual football club's performance could be characterised by a constant win rate.

The measured win rate over any season is only an indication of a team's performance in that past season. It cannot be used on its own to predict future performance with accuracy or precision.

Appendix C

Variability in safety system device failures

Some failures of electronic **components** are purely random [7, 9].

Purely random behaviour results from a stochastic process or Bernoulli process, a series of independent events. The impact of cosmic radiation on individual electronic components is stochastic. The rate of collisions and resulting damage is reasonably constant.

Purely random failures of components can never be eliminated or prevented.

Failures are classed as **systematic** rather than random if they have a known cause and can be eliminated, avoided or controlled to some extent [6]. The probability of systematic failures can be changed.

Failure rates might be reduced through radiation hardening or through fault detection and correction. Devices relying on electronic components can be designed to be fault tolerant. Eventually device failure rates increase over time as damage accumulates. The failure behaviour would no longer be purely random if failure rates increase. Purely random behaviour is characterised by constant rates.

Purely random failure in safety system applications is limited to devices such as:

- Logic solver electronics
- Sensor electronics
- Signal conditioners
- Variable speed motor drives

Most failures in field-mounted electronic components are not purely random because they depend on the strength of the component with respect to applied stresses. They may also depend on various physical and chemical phenomena, almost none of which are purely random (see Chapter 11 in Smith, D.J. *Reliability, Maintainability and Risk* [7]).

Stress/strength failures are *quasi*-random. The failure rates appear to be reasonably constant, but the failures are not due to stochastic processes and can be changed.

These failures can be treated mathematically as if they were random if they cannot be eliminated and if a reasonably constant failure rate can be achieved in the target application environment.

Quasi-random failure rates vary widely between environments. The failure rates may vary with an Arrhenius characteristic. Strength degrades as damage accumulates over time. The failure rates depend on the magnitude and duration of the stress, and on the design strength limits. Design strength depends on manufacturing methods, materials, tolerances, inspection and testing.

Stress-related failure rates can be controlled by design, testing, and by reliability-centred preventive maintenance.

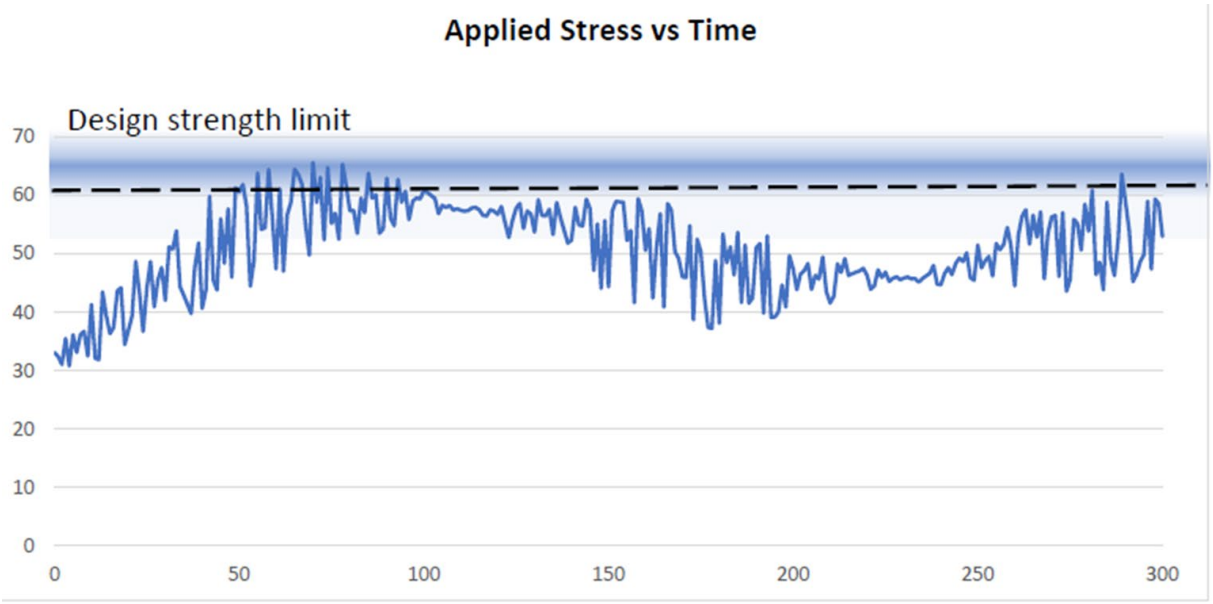

FIGURE 7. Stress exceeding strength causes quasi-random failure behaviour (after Figure 11.2 in *Reliability, Maintainability and Risk* by Dr D J Smith [7])

Typical stress factors include:

- Temperature
- Shock
- Vibration
- Cyclic loading
- Voltage surge.

Electronic **devices and systems** are usually composed of many individual components. The overall failure rates of devices and systems cannot be fixed and invariant because the failure rates of most components are dependent on applied stress. The failure rate of each component can vary over at least an order of magnitude [7, 10, 11]. Reasonably constant failure rates can be achieved in environments where the stress factors are controlled.

Failures in non-electronic components do not occur at fixed rates because the failure mechanisms are all dependent on design and on environmental factors. The failures are not associated with any stochastic processes. Safety system elements that are not subject to purely random failures include:

- Chemical and physical process interfaces
- Mechanical components
- Electro-mechanical components
- Pneumatic components
- Hydraulic components
- Wiring systems
- Communication networks
- Human decisions.

Causal factors can be identified for most failures

Performance in safety systems is much like performance in football. Safety performance depends on a team effort. It is affected by many factors such as:

- Effective management
- Equipment design quality, including suitability for service
- Manufacturing quality

- Environmental influences
- Installation design
- Competence of the designers, installers and maintainers
- Accessibility for maintenance
- Availability of resources for maintenance
- Effectiveness of maintenance, inspection and testing.

Failures in safety systems are classed as systematic because they have causes that can be determined and to at least some extent the failures can be eliminated, avoided or controlled.

Appendix D

Introduction to Bayes' theorem

Bayes' theorem allows us to modify an estimated probability of an event by considering factors that are known to affect the probability.

Refer to the text '*Risk Assessment and Decision Analysis with Bayesian Networks*' by Fenton and Neil [13] for an explanation of Bayes' theorem.

In this summary we assume that the 'probability of the hypothesis' is $P(H)$ and the 'probability of the evidence' is $P(E)$.

The 'probability of the hypothesis **and** the evidence both happening' is $P(H \cap E)$.

That can be expressed as the 'probability of the hypothesis given the evidence' multiplied by the 'probability of the evidence', i.e. $P(H \cap E) = P(H|E) \times P(E)$

It can also be expressed the other way around as the 'probability of the evidence given the hypothesis' multiplied by the 'probability of the hypothesis': $P(H \cap E) = P(E|H) \times P(H)$.

These three terms are therefore equal: $P(H|E) \times P(E) = P(H \cap E) = P(E|H) \times P(H)$,

Bayes' theorem is that the 'probability of the hypothesis given the evidence' can be estimated if we have prior knowledge of 'probability of the evidence given the hypothesis':

$$P(H|E) = \frac{P(E|H) \times P(H)}{P(E)}$$

This allows us to modify the estimate probability of an event if we have prior knowledge of factors that are known to influence that probability.

A simple example of how this can be applied is given below in Appendix F.

Appendix E

Bayesian analysis using probability distributions

Bayes' theorem can be applied with probability distributions as well as with discrete probabilities.

Probability distributions can be used to model the occurrence of series of events. Many different types of distribution are available. Distributions commonly used in safety system engineering include exponential, Poisson, binomial, Normal and Gaussian distributions.

Each type of distribution has one or more parameters that determine the scale, extent and shape of the distribution. For example: exponential distributions are characterised by a rate; Normal or Gaussian distributions are characterised by mean and variance or mean and standard deviation.

The conventional method of modelling is to collect data and then calculate the parameters that make the chosen distribution fit the data best. Few processes can be modelled by a fixed set of parameters. The parameters for most processes will vary over same range.

Bayesian analysis provides a technique that evaluates the relative likelihoods that different sets of parameter values accurately represent a system's behaviour. It can reveal the expected range over which the parameters will vary, and the most likely values. The Bayesian approach allows for subjectivity and uncertainty in assessing event probability distributions [4, 5].

This type of Bayesian analysis starts by identifying the expected value or a range of possible values for the parameters. A range of parameter values may be chosen for comparison and an expected probability distribution can be assigned to the range of values.

Bayesian analysis may be useful in gaining insights into variability in the performance of systems.

For example, the prior knowledge can be a selection of assumed range of failure rates for the equipment.

Bayesian analysis with a flat prior to estimate uncertainty

Upper and lower bounds of an uncertainty interval for failure can be inferred from a posterior distribution when a uniform prior distribution is used.

In this study Bayesian analysis was applied to the number of tosses needed to toss the first 20 heads in each set of 100 tosses, and to the number of matches between wins for the first 20 wins by the Liverpool Football Club in each season between 1990 and 2000.

The posterior distributions indicate the likely range of heads per toss or wins per match.

The calculation method followed the example described by Recchia in the presentation *'Bayesian Methods in Reliability Engineering'* [5].

The number of tosses to obtain a head or number of matches needed to score a win was taken to be exponentially distributed. Random behaviour was assumed with the win rate denoted by λ.

The probability of the win event occurring at time t given win rate λ is $\lambda . e^{-\lambda . t}$

The probability of a series of win events occurring successively at times t_i given win rate λ is $\prod_{i=1}^{n} \lambda . e^{-\lambda . t_i}$

The assumed prior distribution is $g(\lambda_j)$, with a series of possible values of λ numbered from 1 to *j.*

The posterior distribution $g(\lambda_j|t_i)$ gives the likelihood of λ taking the value of λ_j given the evidence of a series of events with measured times t_i:

$$g(\lambda_j|t_i) = \frac{\prod_{i=1}^{n} \lambda . e^{-\lambda . t_i} \times g(\lambda_j)}{\sum_{j=1}^{n} \left(\prod_{i=1}^{n} \lambda . e^{-\lambda . t_i} \times g(\lambda_j) \right)}$$

In this study an 'uninformed' or 'weak' prior distribution was assumed in order to estimate an uncertainty interval. The prior distribution is flat or uniform. Each value is equally likely, though the likelihood of the two extreme values is increased slightly to bring the total probability to 1.:

Win rate:	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Expected likelihood:	0.15	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.15

The resulting posterior distributions vary slightly between each set of tosses in the same way that short-term average win rates vary.

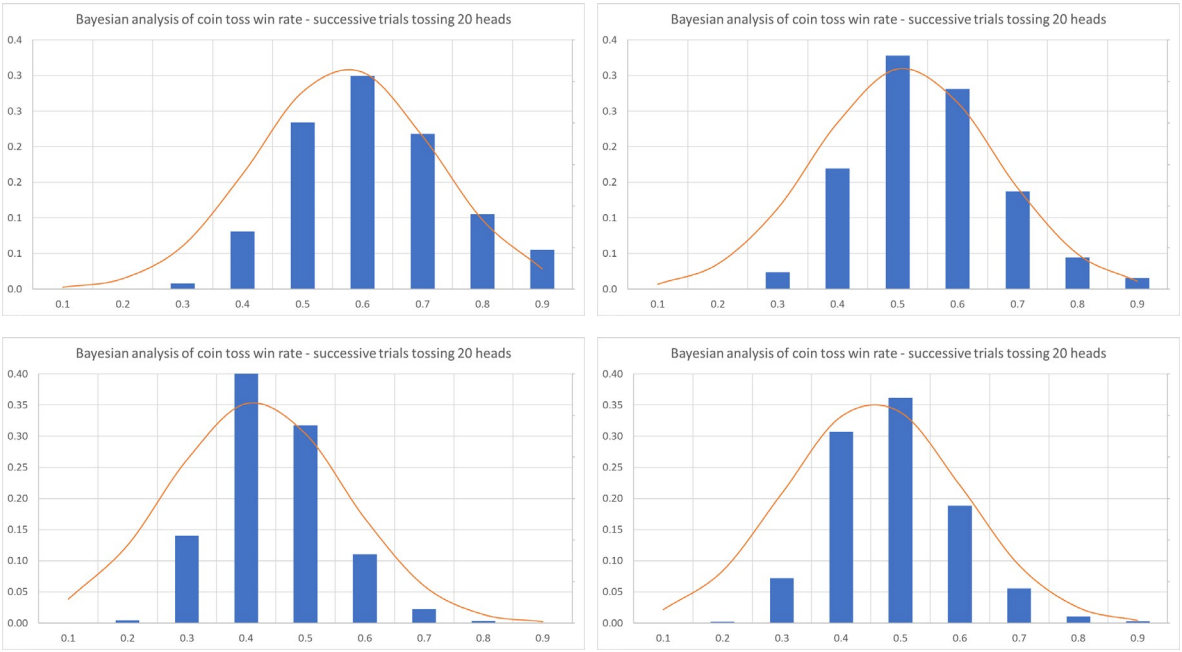

FIGURE 8. Bayesian posterior distributions of coin toss results based on a flat prior

The same flat prior was applied for the first 20 wins by the Liverpool Football Club in each season.

The posterior distributions for the football win rates are similar to the posterior distributions calculated for the coin toss trials. The difference between coin toss trials and football matches is not as obvious as it is if short-term averages are compared.

Uncertainty intervals based on Bayesian analysis are not effective in revealing that the performance of a football team is variable rather than random. Similarly, uncertainty intervals do not reveal that safety system failures are variable rather than random.

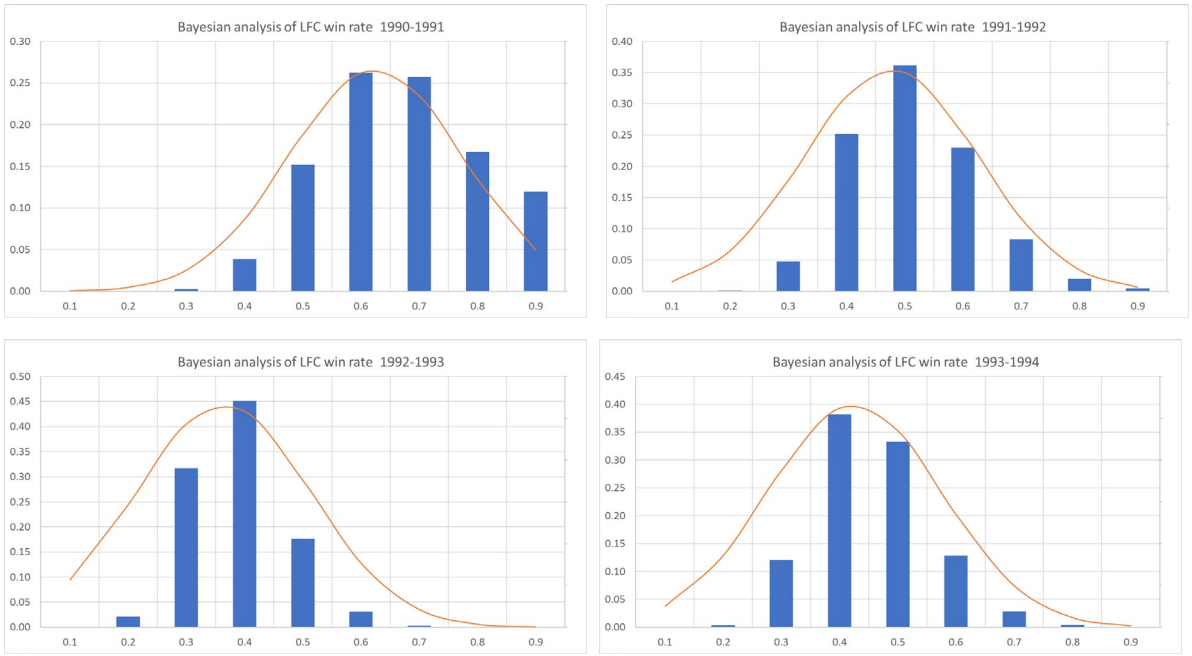

FIGURE 9. Bayesian posterior distributions of football win rates based on a flat prior

Bayesian analysis with a narrow ('strong') prior masks variability

Repeated Bayesian analysis might use prior distributions based on the posterior distributions from previous analysis. Successively updating the prior distributions masks the extent of variability in performance. The resulting posterior distributions become progressively narrower and eventually 'shrink' to indicate only the long term mean values of the parameters. A narrow prior is called a 'strong' prior.

After only 10 iterations the posterior distribution becomes so narrow that it shows only the long term mean values of the parameters. It no longer indicates how widely the parameters vary over time.

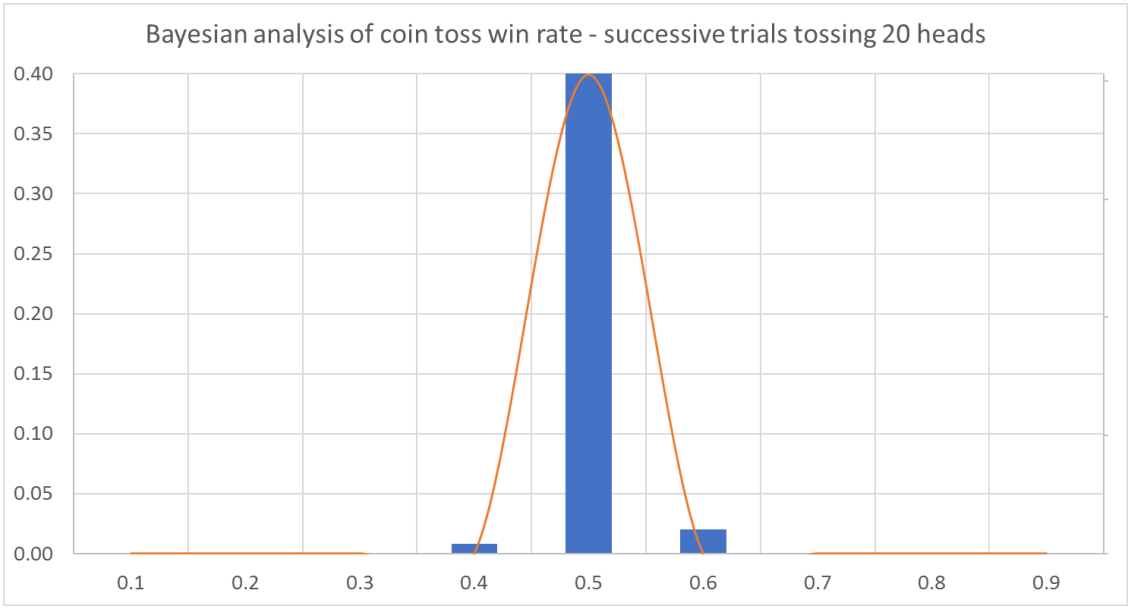

FIGURE 10. Bayesian posterior distribution of coin toss results with successively updated prior distributions over 10 sets of trials

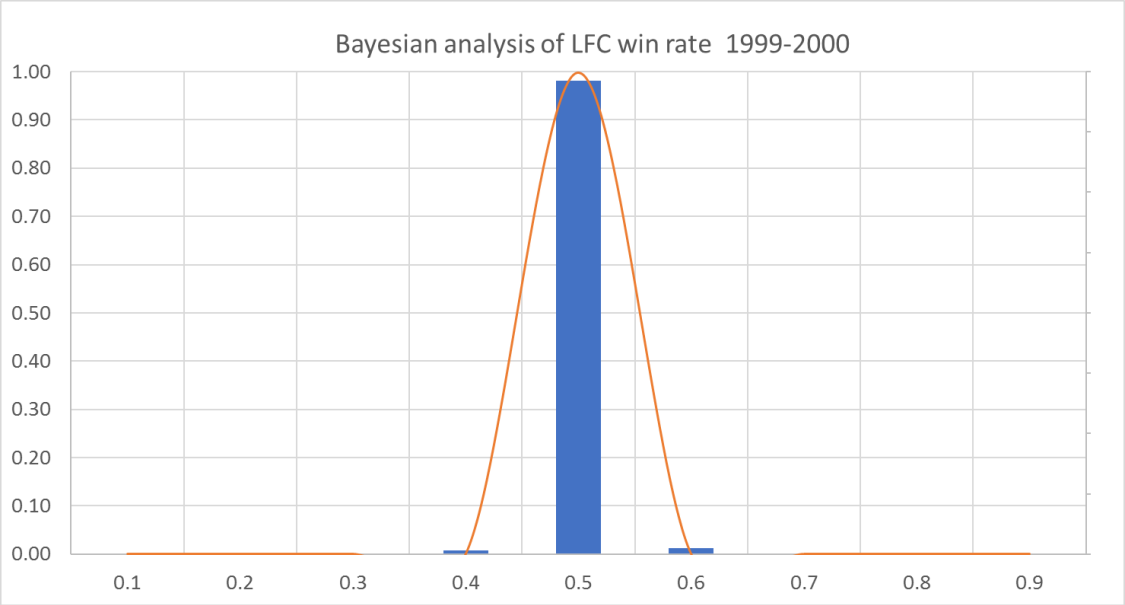

FIGURE 11. Bayesian posterior distribution of football win rates with successively updated prior distributions over 10 seasons

Again, the analysis does not reveal that the performance of a football team is variable rather than random.

Bayesian analysis with a biased prior supports confirmation bias

One potential advantage of Bayesian analysis is that it allows subjective existing knowledge or beliefs to be included in the analysis.

This type of analysis needs to be treated with caution because a biased prior will skew the results to reinforce the assumption - even if that assumption is not valid.

The following example shows an analysis testing the hypothesis that the rate of heads is 0.6 per toss in a trial tossing 20 heads. A biased prior was used with 0.6 as the most likely win rate:

Win rate:	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Expected likelihood:	0	0	0	0.05	0.1	0.7	0.1	0.05	0

The resulting posterior distribution after 20 heads seems to reinforce the assumption that the true rate of heads is 0.6 per toss, even though the actual rate in the trial is exactly 0.5 heads per toss:

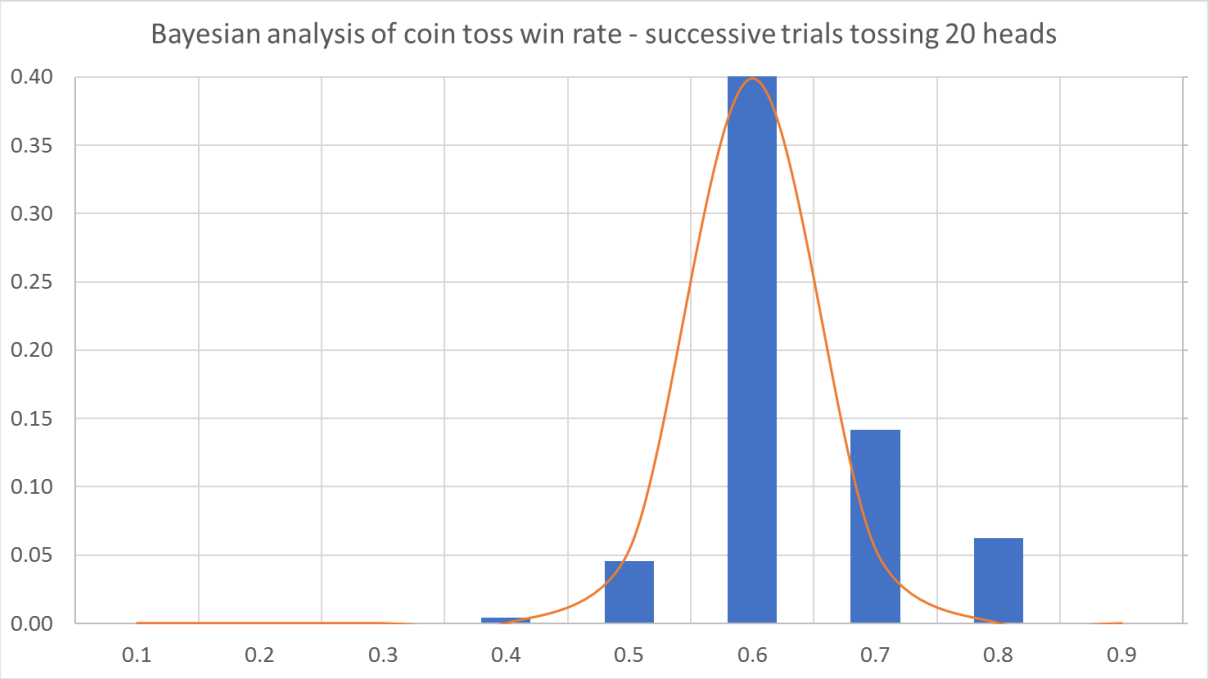

FIGURE 12. Bayesian posterior distribution of coin toss results with a biased prior

Appendix F

Bayes' theorem applied to probability of failure

The following purely hypothetical examples are intended to illustrate how Bayes' theorem can be applied to explore the impact of different maintenance strategies. They are based on similar examples given in Fenton and Neil, chapter 6 [13].

Case 1 – A site with poor maintenance practices

Hypothesis: Overall, we might expect 10% of devices to fail on demand or on test. The 'probability of the hypothesis' that any individual device has failed is $P(H) = 0.1$

Evidence: We might find that around 30% of the devices that we inspect are in a degraded condition due to ineffective maintenance. The 'probability of the evidence' for ineffective maintenance is $P(E) = 0.3$

Prior knowledge: We might know that around 80% of the devices that fail on demand or on test are found to have defects that resulted from ineffective maintenance. The 'probability of the evidence' for ineffective maintenance *given* that 'the hypothesis of failure is true' (i.e. the device has already failed) is $P(E|H) = 0.8$

Posterior belief: We can estimate the increased likelihood of a device failing if we know that it has not been maintained effectively. This is the probability of the hypothesis that the device will fail on demand or on test *given* that we know that maintenance may have been ineffective:

$$P(H|E) = \frac{P(E|H) \times P(H)}{P(E)} \approx \frac{0.8 \times 0.1}{0.3} \approx 0.3$$

Therefore the probability of device failure could be expected to increase by a factor of around 3 (from 0.1 to 0.3) if it is known that the device has not been maintained. Note that the precision here is limited to a single figure because the performance is understood to be variable and the estimates of probability are subjective.

The ratio of increase is simply the ratio $P(E|H) \,/\, P(E)$.

The probabilities might be expected to vary, for instance: and $P(E|H)$ might vary between 0.7 and 0.9 and $P(E)$ between 0.2 and 0.4. The resulting factor of increase in failure probability might then vary in the range from 2 to 4. A factor of 3 is a reasonable approximation.

Case 2 – A site with effective maintenance practices

Effective maintenance can reduce the overall probability of failure of the devices by an order of magnitude. We would expect to find fewer degraded devices on inspection.

Hypothesis: Overall, we expect only 1% of devices to fail on demand, $P(H) = 0.01$

Evidence: We find that around 2% of the devices that we inspect are in poor condition due to ineffective maintenance, $P(E) = 0.02$

Prior knowledge: 80% of the devices that fail on demand or on test are found to have defects that resulted from ineffective maintenance, $P(E|H) = 0.8$

Posterior belief: The probability of the hypothesis that the device has failed *given* that we know maintenance may have been ineffective is:

$$P(H|E) \approx \frac{0.8 \times 0.01}{0.02} \approx 0.4$$

At this second site, the probability of device failure could be expected to increase by a factor of around 40 (from 0.01 to 0.4) if it is known that the device has not been maintained.

The estimated probability of failure given knowledge of ineffective maintenance, $P(H|E)$, is similar to the estimate in Case 1 above. The ratio $P(E|H)/\,P(E)$ has increased by a factor of around 10 because $P(E)$ is around 10 times lower, but $P(H)$ is also around 10 times lower.

With some variation in our assumed probabilities, the factor of increase might vary in the range from 30 to 70. We should not expect to be able to quantify these probabilities with precision.

The point of this example is that probability of failure can be increased by one or two orders of magnitude if maintenance is ineffective. We should complete the planned maintenance if we want safety devices to function as designed.

There is no value in trying to calculate the precise factor by which probability of failure will increase if the maintenance is ineffective. It is enough to know that the probability of failure can easily increase by at least a factor of 3 and possibly by as much as a factor of 30 or more. The scale of the effect is already evident from OREDA [10, 11] and from studies such as that by Bukowski and Stewart [21].

We do not need to carry out calculations to understand the causal relationship in every similar case, but we do need to understand the cause (ineffective maintenance) and its potential impact (increased probability of failure).

There is, however, value in setting and monitoring key performance indicators for maintenance. The purpose of the performance indicators is to ensure that maintenance is completed as planned.

Appendix G

Application of prior knowledge in LOPA

Here are two hypothetical cases of how a heuristic Bayesian approach might be adopted.

Case 1 – A site with excellent maintenance practices

Consider a company that has a risk-averse culture. Safety and health are valued more highly than profit, but continuity of production is also important. The company's philosophy is that a safe plant is a productive plant. Production and safety goals should not be in conflict.

The company operates with a formal reliability-centred maintenance system [28]. The production plant is of a modular and redundant design. Each of the critical parts is designed to allow inspection and maintenance to be carried out without interruption to production.

The plant operators have established a comprehensive set of leading indicators for equipment failures and process disturbances. Formal systems and procedures are in place to measure, control and monitor performance. The process safeguards are subject to regular periodic inspections and audits.

The plant owners have several years of recorded evidence to show that at this plant the event rates and failure probabilities are typically 3 to 10 times better than the rates presented in the CCPS guidelines [22]. They make sure that the event rates assumed in the LOPA studies are traceable to the failure analysis and equipment condition reports.

Case 2 – A site with compromised practices

Consider another company which runs a refinery that is in relatively poor condition. The refinery has been in operation for almost 50 years. Profit margins are tight. The opportunities and resources for equipment maintenance are restricted. Funding is limited. The equipment is not readily accessible for maintenance while the plant is operating. Production targets take priority.

The company engages an engineering contractor to design some modifications to the plant. The design team plans to carry out a LOPA study. The designers walk around the plant and notice that much of the equipment is in a degraded condition. The operating company has some records of maintenance, but the records are incomplete. The causes of failure have not been investigated; the work orders often only state that the equipment was 'broken'. The descriptions of corrective actions lack detail: 'repaired, returned to service'. Failure rates have not been analysed. There are no audit procedures, let alone any audit reports. It looks like some of the pressure relief valves might not have been serviced for several years. The designers notice what seem to be cracks in the bund walls around the storage tanks. Most of the explosion-protected electrical equipment appears to be in a degraded state due to heavy corrosion.

After talking with the control room operators and the maintenance supervisors, the design engineers suggest to the operating company that the *exida* Site Safety Index should be assumed to be 'SSI 0' [29]. They will need to assume failure rates and failure probabilities at least 3 times or even 10 times worse than those suggested in the CCPS LOPA guidelines. In many cases they will not be able to take any credit for independent protection layers because there is no evidence to demonstrate protection integrity.

After a heated argument, the operating company advises the design engineers that their services are no longer required if they can't just follow the guidelines. They decide to find another design contractor that will take a more 'cost effective' approach and will follow the guidelines without question. After all, they have been operating for almost 50 years without a major accident event (yet).

