A Hierarchical Bayesian Approach to IEC 61511 Prior Use

Stephen L. Thomas, P.E.
SIS Engineering Team Lead (OASIS)
Chevron Corporation
Houston, TX
Email: stephen.thomas@chevron.com

Prepared for Presentation at
American Institute of Chemical Engineers
2018 Spring Meeting and 14th Global Congress on Process Safety
Orlando, Florida
April 22 – 25, 2018

A Hierarchical Bayesian Approach to IEC 61511 Prior Use

Stephen L. Thomas, P.E.
Chevron Corporation
Houston, TX
Email: stephen.thomas@chevron.com

Keywords: Safety Instrumented Systems, failure rate, Prior Use, hierarchical Bayes, confidence intervals, uncertainty, Markov Chain Monte Carlo, Gamma distribution

Abstract

Understanding the quality of failure rate data is vital to successful application of SIS, as emphasized in the latest IEC 61511 standard, including quantifying the relevant uncertainties of the inputs and communicating the confidence metrics in SIL verification calculations. This paper discusses the application of Bayesian credibility intervals to determining prior use failure rates for components in SIS service. This approach compares favorably to traditional frequentist approaches because it can incorporate diverse types of prior knowledge. Guidelines for developing Bayesian prior distributions are given, including practical examples of prior distributions based on industry data and a demonstration of the Bayesian updating process. The concept of hierarchical prior distributions is introduced and used to develop a practical model for managing enterprise failure rate data. The advantages (and potential pitfalls) of the Bayesian approach are discussed, including the inherent handling of uncertainty, as well as the potential to significantly reduce the total service hours required for prior use justification.

1 Introduction

Since the advent of the ANSI / ISA 84 standard in 1996 and the international IEC 61511 standard in 2004, the concept of performance-based design of Safety Instrumented Systems (SIS) and the related concept of probabilistic risk assessment (PRA) have steadily been gaining widespread acceptance within the process industries. However, it is a recurring challenge to ensure that the probabilistic calculations in the PRA and SIS design are relevant and meaningful. Too often, these calculations are based on reliability data of questionable quality and unknown uncertainties. Safety Integrity Level (SIL) verifications may calculate an average Probability of Failure Demand (PFDavg) to three decimal places based on data with uncertainties of two orders of magnitude.

The standard (herein referring to IEC 61511 unless otherwise noted) recommends a method called “prior use” justification, involving a combination of qualitative assessments as well

as field failure data analysis, to ensure that reliability data is relevant and that uncertainties have been minimized. Further, the standard requires that procedures be implemented to evaluate the SIS performance during operation, i.e. the model needs to be validated against reality. In practice, these requirements have proven quite difficult for end-users of the standard to effectively implement.

In this paper, we will briefly discuss why these requirements are so difficult using traditional statistical methods. We will then explore how Bayesian statistical inference offers significant advantages for the analysis of SIS reliability, with an emphasis on component prior use justification.

Bayesian methods have enjoyed widespread use in the fields of nuclear safety, aerospace risk management, as well as many other fields. They have received only limited exposure in the process industries, largely due to a perception that the underlying mathematics is complex and onerous. This paper will demonstrate that the present generation of software tools makes even advanced Bayesian analysis readily accessible to non-specialists, with significant benefits.

Starting with a simple example of valve reliability, we first demonstrate the construction of a prior distribution using generic industry data sources. Then the concept of Bayesian updating is introduced using simple conjugate prior distributions. Next, the problem of non-homogenous populations is addressed using hierarchical Bayesian analysis. Finally, these methods are tied together to show how they can be used to build step-by-step a logically coherent, highly efficient system for analyzing and managing SIS (or any other) reliability data across a plant or enterprise.

2 Literature

This paper concentrates on practical applications and does not attempt to provide an in-depth introduction to Bayesian statistics. However, there are many papers, textbooks, and handbooks that provide either brief introductions or in-depth treatments of the topic(s). The brief literature review below is by no means comprehensive but provides an overview of key literature for the interested reader.

As mentioned above, the nuclear industry has made broad use of Bayesian methods as documented in several handbooks [1][2]. In particular, NUREG/CR-6823 provides a useful conceptual introduction to Bayesian concepts. Similarly, NASA offers a handbook [3] for using Bayesian methods in risk assessment. While the NASA handbook is very oriented to practical application, it is also very focused on modelling software and does not dwell on concepts.

Many textbooks are available covering Bayesian statistics in general, such as Gelman et al. [4]. Of particular interest to readers of this paper would be texts concentrating on Bayesian applications in risk assessment and reliability, including Kelly and Smith [5] and Hamada et al. [6].

Notable examples of process industry applications of Bayesian methods include several papers from the SINTEF and NTSU organizations [7][8][9] as well as an application of failure rate estimation using Bayes conjugate priors by Shafaghi [10]. Khan et al. [11] have published

a comprehensive literature review of research in process safety that includes references to other relevant Bayesian literature.

A version of the hierarchical Bayesian procedure was initially proposed as a "two-stage" procedure by Kaplan [12], although the procedure format is quite different since it preceded the widespread availability of computerized Bayesian algorithms. More modern treatments of hierarchical Bayes are covered by Pörn [13], Droguett and Groen [14], and Kelly and Curtis [15]. Hierarchical Bayesian models may also be viewed as a special case of a Bayesian Belief Networks of which Weber et al. [16] provides a useful overview of recent applications.

As will be discussed later, practical solutions to many Bayesian problems rely on numerical solutions using the Markov Chain Monte Carlo (MCMC) technique implemented using various software algorithms. Several software packages are available for this analysis, but we have chosen to use the free JAGS software within the popular free R / RStudio statistics software. There are a wide variety of free resources available online related to R and JAGS. There is also a notable textbook by Krushke [17] that provides many example models using R and JAGS.

3 Background and Motivation

3.1 Data quality in IEC 61511

Consistently, one of the major challenges of implementing PRA and performance-based SIS has been finding applicable failure data for different process industry applications. The IEC 61511 standard itself notes that the lack of high quality reliability data reflective of the operating environment has been a significant shortcoming of PRA and SIS probabilistic calculations [18].

To fill this gap, various sources of reliability data have been made available, including SIL certifications from various providers, Reliability Data Collection Projects, Reliability Handbooks, etc.

However, these resources arguably do not fully meet the intent of the IEC 61511 standard, which requires that:

> *"[reliability data] shall be credible, traceable, documented, justified and shall be based on field feedback from similar devices used in a similar operating environment."*

and further:

> *"[data] uncertainties shall be assessed and taken into account when calculating the failure measure."* [18]

Many of the above data sources are based on generic devices (i.e. not model specific) or generic process services, or both. Only a few of them explicitly provide estimates of the uncertainty in their data, but even that may be suspect since it is likely aggregated based on generic devices, generic service, or both.

3.2 Static vs. Dynamic Approaches to SIL Verification

As will be discussed later in the paper, the lack of information related to the quality and uncertainty of generic industry reliability data makes it difficult to effectively monitor whether the actual performance of the SIS is consistent with the reliability parameters assumed during the design. This difficulty has resulted in what I refer to as a static approach to SIL verification, as shown in Figure 1 below.

Figure 1. Static Approach to SIL Verification

Figure 1 shows what is essentially "open loop control" of SIS performance. The system is originally modelled and designed using static inputs to a static SIL verification model. Reliability data may be monitored and gathered on the operations phase, but only ad hoc methods are available to determine bad actors on a case-by-case basis.

What the standard actually recommends is a dynamic approach where the performance of the SIS is monitored and *evaluated* versus the design assumptions. The design assumptions should be updated based on actual performance (in our case, using Bayesian inference), leading to a dynamic approach to SIL verification, as shown in Figure 2 below.

Figure 2. Dynamic (Bayesian) Approach to SIL Verification

As will be demonstrated in this paper, the dynamic Bayesian approach inherently provides "closed loop" evaluation of performance and also addresses all of the data quality requirements of the standard, namely:

- Real-word data based on field feedback
- Credible, traceable, documented data
- Uncertainties assessed
- Performance monitored

The Bayesian framework offers several advantages over traditional frequentist methods, but the most important advantage is feasibility. Traditional methods are often just not feasible for SIS analysis, as briefly discussed next.

3.3 Difficulties of the Frequentist Approach

The core limitation of traditional frequentist methods versus Bayesian methods is that frequentist methods typically do not consider prior knowledge. They make the implicit assumption that prior to taking a sample, all possible outcomes are equally likely. This concept is best illustrated with an example.

Example question: *We have several valves in similar service. We have prior use data for these valves where we have experienced one failure in 871,620 service hours (100 service years). What failure rate should we use in our SIL calculations?*

The point estimate for the failure rate is simply:

$$\hat{\lambda} = \frac{1\ failure}{871{,}620\ hr} = 1.15 \times 10^{-6}\ /hr = \frac{1}{100\ yr} \qquad (1)$$

However, IEC 61511 calls for a 70% upper confidence limit for prior use data. A frequentist approach would have us calculated the 70% confidence limit based on our sample data using the χ^2 distribution with 4 degrees of freedom.

$$\lambda_{70\%} = \frac{\chi^2_{70\%,4}}{2\times 871620} = \frac{4.88}{1743240} = 2.78 \times 10^{-6}/hr \qquad (2)$$

Note that this result is 2.8x times higher than the point estimate, even with 100 service years of sample data. This result begs the question; how many years of sample data are required to have 70% confidence the failure rate is less than 1/100 yr (i.e. the point estimate)?

$$T = \frac{\chi^2_{70\%,4}}{2\times\hat{\lambda}} = \frac{4.88}{0.02} = 244\ yr \qquad (3)$$

This result demonstrates the fundamental issue with using frequentist methods for SIS performance analysis. It simply requires too much data to build the required confidence because there is no incorporation of prior knowledge.

To drive home this point, a similar frequentist analysis for an entire Safety Instrumented Function (SIF) indicates that at least 120 successful tests (or demands) are required to achieve 70% confidence that SIL 2 performance (i.e. $PFD_{avg} < 0.01$) has been achieved. This amount of testing is clearly beyond the lifetime of a single SIF and is probably infeasible for all but the largest populations of identical SIFs.

The preceding section has attempted to frame the problem. The remainder of the paper will describe the Bayesian analysis process and develop several examples illustrating how the process addresses the issues outlined above.

4 A Simple Bayesian Approach with Conjugate Priors

The Bayesian updating process is quite straightforward, but Bayes theorem relates conditional probabilities and can be generally written as:

$$P(A|B) = \frac{P(B|A)\times P(A)}{P(B)} = \frac{P(B|A)\times P(A)}{P(\mathrm{A|B})\times P(A)+P(B|\bar{A})\times P(\bar{A})} \qquad (4)$$

The above equation (4) can be re-written for continuous probability distributions as:

$$P(A|B) = \frac{P(B|A)\times P(A)}{\int_{-\infty}^{\infty} P(\mathrm{B|A})\times P(A)dA} \quad (5)$$

We are interested in the probability that a failure rate (λ) is a certain value conditional on some observable evidence (E), such as failures observed over some time period. Substituting into equation (5) yields:

$$P(\lambda|E) = \frac{P(E|\lambda)\times P(\lambda)}{P(E)} = \frac{P(E|\lambda)\times P(\lambda)}{\int_{0}^{\infty} P(\mathrm{E|\lambda})\times P(\lambda)d\lambda} \quad (6)$$

The following terminology is commonly used to break down the equation:

- P(λ) Prior Distribution – represents prior knowledge about the failure rate λ
- P(E|λ) Likelihood – the likelihood of observing the evidence given λ
- P(E) Marginal Likelihood - likelihood of observing the evidence conditioned over all values of λ
- P(λ|E) Posterior Distribution – updated knowledge of λ based on the evidence

Equation (6) is the form that will be used for the initial analysis, but first the Likelihood and Prior Distribution must be defined.

4.1 Likelihood

SIL calculations per IEC 61511 are typically based on the exponential distribution, which is a special case (i.e. where x=0) of the more general Poisson distribution. Given a constant failure rate (λ), the Poisson distribution gives the probability of failures (x) per time period (t), as shown below.

$$P(\mathrm{x,t|\lambda}) = \mathrm{e}^{-\lambda \mathrm{t}}\frac{(\lambda t)^x}{x!} \quad (7)$$

In the completed model, the variables x and t will take the place of the evidence (E) in equation (6).

4.2 Prior Distribution

The prior distribution represents a quantification of our prior knowledge about the failure rate (λ). Prior knowledge may come in many forms, including but not limited to:

- Industry data
- Expert opinion
- Testing
- Engineering analysis (e.g. FMEDA)

There are many techniques and considerations when selection a prior distribution. For purposes of this paper, our criteria include simplicity and mathematical tractability. For these reasons, a Gamma distribution was chosen as the prior distribution. The Gamma distribution is a "conjugate prior" of the Poisson likelihood function which enables equation 6 to be solved analytically and elegantly, as will be show below.

The Gamma distribution for the failure rate (λ) can be written as a function of a scale parameter (α) and a rate parameter (β):

$$P(\lambda|\alpha,\beta) = \frac{\beta^{\alpha}\lambda^{\alpha-1}e^{-\beta\lambda}}{\Gamma(\alpha)} \quad (8)$$

$$Mean = E(\lambda) = \frac{\alpha}{\beta} \quad (8a)$$

$$Variance = Var(\lambda) = \frac{\alpha}{\beta^2} \quad (8b)$$

The task of selecting the prior becomes selecting values for α and β, which will be called α_0 and β_0. To make this selection, data for valve failure rates was gathered from a variety of industry data sources, as shown in Table 1 below.

Table 1: Industry Failure Rate Data for Valves

Source Type	Failure Rate hr^{-1} (λ_{DU})		
	Low	*Mean*	*High*
Website	8.2E-07	--	2.3E-06
Certificate	--	9.0E-07	--
Book	1.1E-06	--	4.6E-06
Certificate		1.3E-06	
Certificate	--	1.5E-06	--
Handbook	1.3E-07	1.9E-06	5.4E-06
Report	--	2.5E-06	--
Average		***2.0E-06***	
Minimum	***1.3E-07***		
Maximum			***5.4E-06***

Using the data in Table 1 and the relations in equations (8a) and (8b), a Gamma distribution can easily be fit that represents our belief about the range of probable failure rates based on available industry data. Our subjective criteria are that the mean match the industry data mean and that all of the industry data falls within a 90% confidence interval. The resulting Gamma prior distribution has parameters:

$$\alpha_0 = 0.9 \quad (9a)$$

$$\beta_0 = 441{,}000 \quad (9b)$$

The same approach may be employed to build prior distributions for other common SIS equipment types. For information, Table 2 provides a summary of other proposed prior distributions.

Table 2: Other Prior Distributions

Component	Gamma Parameters	
	Alpha	*Beta*
Diff. Pressure Transmitter	0.6	623,000
Solenoid Valve	0.6	403,000
SIL 3 Logic Solver	0.5	2,300,000
SIL 2 Logic Solver	0.5	840,000

Figure 3 shows the resulting prior distribution plotted with the average, minimum, and maximum of the industry data. The cumulative distribution is shown for clarity.

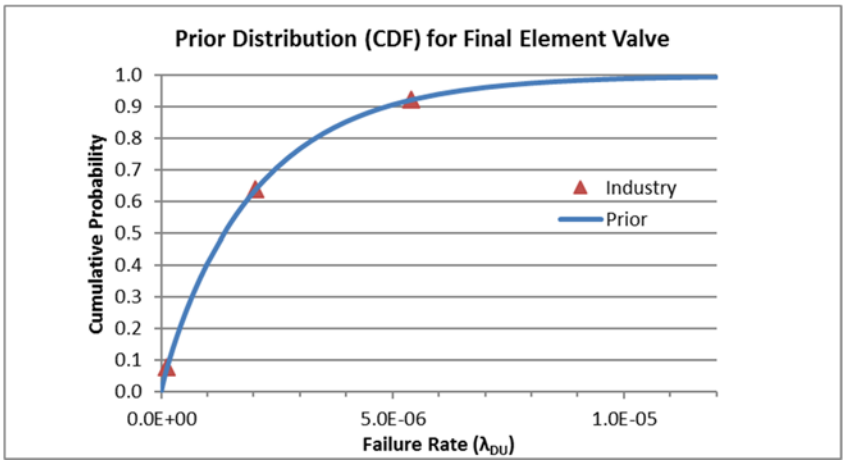

Figure 3: Prior Distribution for Valve Failure Rate (λ)

4.3 Bayesian Updating Process

Now that the likelihood function and the prior distribution have been defined, we can return to equation (6) and begin the updating process. Substituting equations (7) and (8) into equation (6) yields the following:

$$P(\lambda|\alpha,\beta,x,t) = \frac{\frac{(\lambda t)^x}{x!}e^{-\lambda t}\frac{\beta_0^{\alpha_0}}{\Gamma(\alpha_0)}\lambda^{\alpha_0-1}e^{-\beta_0\lambda}}{\int_0^{\infty}\frac{(\lambda t)^x}{x!}e^{-\lambda t}\frac{\beta_0^{\alpha_0}}{\Gamma(\alpha_0)}\lambda^{\alpha_0-1}e^{-\beta_0\lambda}d\lambda} \qquad (10)$$

Despite the appearance of complexity, this equation can actually be integrated and solved analytically. The earlier choice of a Poisson likelihood and a conjugate Gamma prior yields an elegant solution, as follows:

$$P(\lambda|\alpha,\beta,x,t) = \frac{\beta_P^{\alpha_P}}{\Gamma(\alpha_P)}\lambda^{\alpha_P-1}e^{-\beta_P\lambda} \qquad (11)$$

where $\alpha_P = \alpha_0 + x$ (11a)

and $\beta_P = \beta_0 + t$ (11b)

Note that the posterior distribution in equation (11) is also a Gamma distribution like the prior distribution in (8), but the parameters have been "updated" with the failures (x) and time (t) data from the evidence. This property conveniently allows us to cyclically repeat the updating process each time there is new evidence to be incorporated. Each cycle, the current posterior distribution becomes the prior for the next cycle. This Bayesian updating process is illustrated in Figure 4 below.

Figure 4: Cyclical Bayesian Updating Process

4.4 Example One-Stage Bayesian Updating

Recall from the earlier example question that for a certain set of valves, a plant had experienced one failure in 871,620 service hours (100 service years). We can now apply the Bayesian updating process to this data. Using equations (11a) and (11b) and the prior distribution parameters in (9a) and (9b) yields:

$$\alpha_P = \alpha_0 + x = 0.8 + 1 = 1.8 \quad (12a)$$

$$\beta_P = \beta_0 + t = 441{,}000 + 871{,}620 = 1{,}312{,}620 \quad (12b)$$

The prior and posterior distributions from this updating cycle are shown in Figure 5 below.

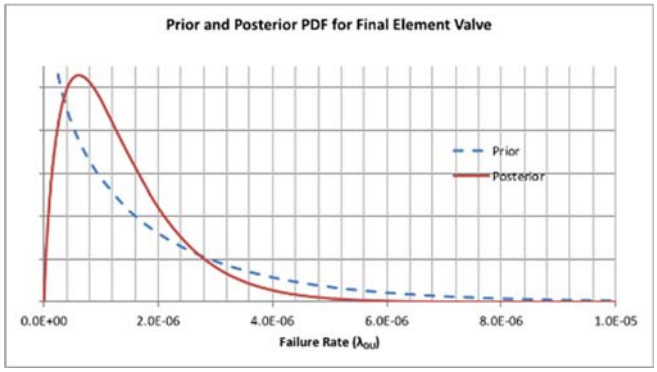

Figure 5: Example Prior and Posterior Distributions

Note that the new evidence causes the posterior to be narrower than the prior (i.e. more volume of evidence decreases uncertainty) and the posterior shifts slightly lower because the new evidence suggests a lower failure rate.

The upper 70% credibility limit (as per IEC 61511 prior use requirements) may be directly calculated based on the quantiles of the new posterior distribution

$$\lambda_{70\%} = 1.68 \times 10^{-6} \quad (13)$$

Note that this value is 40% lower than the earlier frequentist result due to the impact of the prior industry knowledge captured in the prior distribution. In this case and potentially many others, the simple Bayesian updating process substantially reduced the number of service years required to justify lower prior use-based failure rates.

5 Hierarchical Bayes

The simple updating process outlined above makes the important assumption that the evidence, and thus the failure rate, is from a homogenous population. In other words, the distribution will eventually converge to an underlying "true" failure rate once sufficient evidence is incorporated (i.e. as $t \rightarrow \infty$). This assumption is often not entirely true in practice, expect perhaps for very small populations.

Instead of a homogenous population (i.e. the same device in the same service), what if an analyst is presented with only *similar* devices in *similar* services? Assuming these are homogenous populations and using the Bayesian updating process above may lead to overly-optimistic results because it ignores the underlying variability in the sub-populations.

One solution to this problem is to only use small, homogenous populations in the analysis. However, this does not address the "sparse data" problem outlined above. Intuitively, an ideal approach would allow us to incorporate the imperfect information coming from these similar services and use them to improve our knowledge of other similar services.

The Bayesian framework can address this problem, but the simple single-stage model above needs to be expanded to what was original called a two-stage model but is now more commonly known as *Hierarchical Bayes*.

5.1 Expanding the Model to Non-Homogenous Data

Recall that equation (6) gave a general form of Bayes theorem relating the evidence (E) and failure rate (λ):

$$P(\lambda|E) = \frac{P(E|\lambda)\times P(\lambda)}{P(E)} = \frac{P(E|\lambda)\times P(\lambda)}{\int_0^\infty P(E|\lambda)\times P(\lambda)d\lambda} \quad (6)$$

This equation may be expanded to incorporate multiple groups of evidence (e.g. E_1,E_2,E_3,…) and failure rates (λ_1,λ_2,λ_3,…) corresponding to multiple sub-populations. We also incorporate the fact that the distribution of failure rates (λ_i) is dependent on the joint distribution of the gamma parameters (α, β).

Others cover this derivation in detail [14] [15], so only the two key results are provided here:

$$P\left(\alpha,\beta\middle|(E_i)_{i=1,\dots,N}\right)=\frac{P((E_i)_{i=1,\dots,N}|\alpha,\beta)\times P(\alpha,\beta)}{\int_\alpha\int_\beta P((E_i)_{i=1,\dots,N}|\alpha,\beta)\times P(\alpha,\beta)\,d\alpha\,d\beta} \quad (14)$$

Note that the α and β parameters are now variables that can take on different values for different sub-populations and have an associated probability distribution. Equation (14) allows these parameters to be updated based on new evidence. Also note that this equation can generally not be solved analytically, so numerical methods will be employed. Both of these topics are discussed in the proceeding sections.

The other key result is the definition of an *expected distribution* for the failure rate (λ) which incorporates all available evidence from the non-homogenous population.

$$P\left(\lambda\middle|(E_i)_{i=1,\dots,N}\right)=\int_\alpha\int_\beta P(\alpha,\beta|(E_i)_{i=1,\dots,N})\times P(\lambda|\alpha,\beta)\,d\alpha\,d\beta \quad (15)$$

The expected distribution includes the variability between the sub-populations, so no amount of evidence would make that uncertainty completely go away. In other words, although the devices or services are similar, there may be inherent differences that do not go away with time. This allowance for population variability addresses the issue of over-optimism resulting from the previous assumption of homogenous data.

It is helpful to illustrate the new model using a hierarchical diagram, as shown in Figure 6 below.

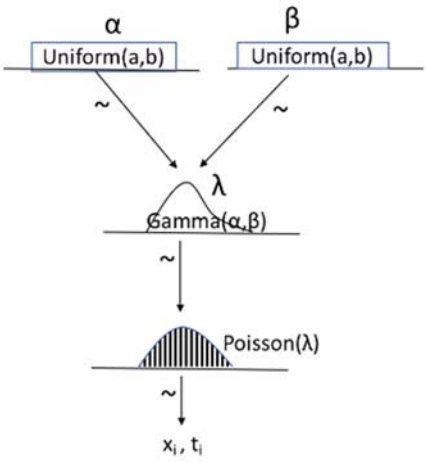

Figure 6: Hierarchical Bayesian Model for Failure Rates

The arrows in the hierarchy indicate probabilistic dependencies, but it is important to realize that inference is *bidirectional*, so the evidence (x_i, t_i) at the bottom is used to update upward in the hierarchy.

5.2 Quantifying Population Uncertainty with Hyperpriors

In the single-stage model, the prior distribution parameters (α,β) were selected as constants based on the available industry data. In the hierarchical model, the uncertainty in these parameters is acknowledged and captured in a prior distribution. When a prior distribution

is placed on model parameters that have no direct physical meaning, they are commonly called *hyperpriors*.

Industry data does not provide much information to inform our choice of the hyperprior distributions. General guidelines used for this paper were:

- The mean of the new distribution should match the constant parameters (α_0,β_0)
- The variance of P($\lambda|\alpha,\beta$) should match the original prior P($\lambda|\alpha_0,\beta_0$)
- The prior distributions should be diffuse to reflect the lack of knowledge

Since numerical methods will be employed, there is no need to consider conjugate prior relationships. Based on these rules of thumb, simple Uniform distributions were chosen for α and β as follows:

$$\alpha \sim Uniform(0.10, 0.90) \tag{16a}$$

$$\beta \sim Uniform(220000,960000) \tag{16b}$$

The literature provides additional guidance and highlights potential problems with technical aspects of choosing hyperpriors [15], but we omit those considerations for clarity.

5.3 Markov Chain Monte Carlo

The complete model is now defined such that equation (15) can be used to calculate the expected distribution of the failure rate (λ). However, as noted above, the problem cannot be solved analytically and requires numeric methods.

The technique for solving this type of model is called Markov Chain Monte Carlo (MCMC). A complete discussion of MCMC is beyond the scope of this paper, but the topic is well covered elsewhere. [4][17] MCMC generally involves taking large numbers of random samples from the different distributions in the model, which ultimately allows the user to sample from the posterior distribution. Various efficient algorithms are available for MCMC simulation, and several free software packages are available. For this study, the popular free MCMC software JAGS was used running inside the popular statistics software R.

The code for implementing the JAGS model in R is shown in Figure 7 below. The code basically consists of three steps:

1. Load necessary libraries and load the failure data for different units from a file
2. Setup the JAGS model based on Figure 6 and including the hyperpriors (14a/b)
3. Run the JAGS model for the specified number of sampling iterations (30,000)

```
## preliminaries
library(rjags)
rjags::load.module('dic')
setwd("/Users/Stephen/Desktop/NewModel")
dat = read.table(file="HPData2.txt", header=TRUE)

## setup the JAGS model
mod_string = " model {
for (i in 1:length(Failures)) {
  Failures[i] ~ dpois(Lambda[Unit[i]] * Hours[Unit[i]])
}
for (j in 1:max(Unit)) {
  Lambda[j] ~ dgamma(alpha, beta)
}
alpha ~ dunif(0.10, 0.90)
beta ~ dunif(220000, 960000)
} "

## run the jags model
set.seed(107)
data_jags = as.list(dat)
params = c("Lambda", "alpha", "beta")
mod = jags.model(textConnection(mod_string), data=data_jags, n.chains=3)
update(mod, 2e3)
mod_sim = coda.samples(model=mod, variable.names=params, n.iter=1e4)
mod_csim = as.mcmc(do.call(rbind, mod_sim))
```

Figure 7: R/JAGS Code for Hierarchical Bayesian Model

5.4 Example Hierarchical Bayesian Updating

With the model built and coded, all that remains is to update the model with actual field data (i.e. evidence). To continue with the valve example, sample data has been collected from several other units that are believed to represent similar (but not identical) valves in similar (but not identical) services. Table 3 below shows the collected data, with Unit 1 representing the data discussed in the previous example.

Table 3: Sample Data from Non-Homogenous Units

Unit	Failures	Hours
1	1	871620
2	0	525600
3	1	1576800
4	0	175200
5	1	1752000
6	0	96360
7	0	700800

Running the model in R/JAGS using the data in Table 3 yields posterior distributions for the failure rate (λ_i) of each unit as well as the posterior hyperparameters (α, β) for the overall population. The results are shown in Figure 8(a-d) below.

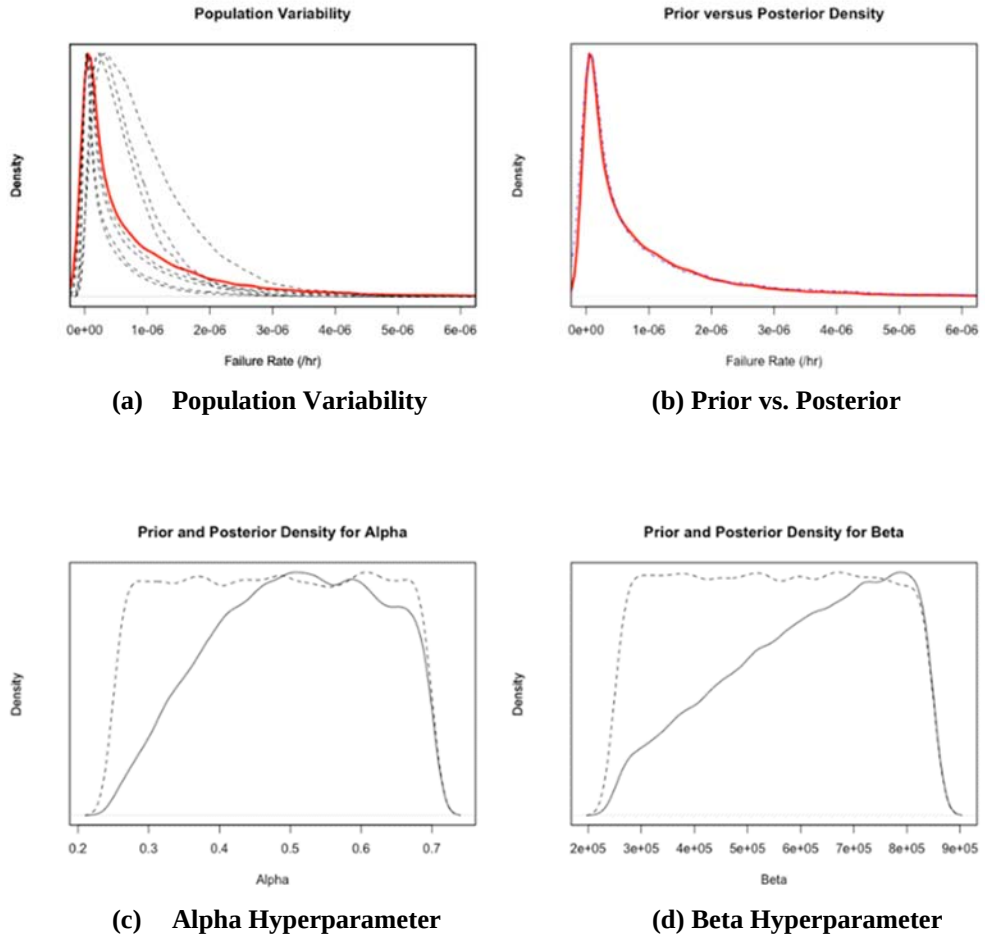

(a) Population Variability **(b) Prior vs. Posterior**

(c) Alpha Hyperparameter **(d) Beta Hyperparameter**

Figure 8: Results of Hierarchical Bayesian Model

6 Discussion

Reviewing the results in Figure 8, note that although significant variability exists among the units (8a), the posterior expected distribution for the overall population did not change much from the prior (8b). The posterior distribution is only slightly narrower than the prior, indicating that the small number and the large variance of the other units limited the "strength" of the new evidence.

Examining the behavior of the hyperparameters in (8c) and (8d), the posterior distributions are narrower than the priors, indicating that the new data has updated our very diffuse priors. However, the hyperpriors are still fairly diffuse. Additional data for either these units or new units will be required to further reduce uncertainty.

Although the small data set in this example did not produce particularly startling results, the key point is that the Hierarchical Bayesian framework positions us to incorporate all future data from all possible units. We will consider next a hypothetical example of how this framework could be deployed across an enterprise.

6.1 An Enterprise Hierarchy of Prior Use Distributions

Consider an example of a typical large chemical or refining enterprise. There would potentially be multiple plants, each with dozens of process units, each unit with many pieces of equipment, and all with many different instrument and valve services. Many of these plants, units, equipment, services have similarities, but they are not the same (i.e. they are not homogenous).

Because the hierarchical Bayes framework makes no assumptions about homogeneity, it creates an ideal framework for grouping and analyzing subjectively similar equipment. These groupings may be as simple or as complex as required to meet the objectives of the analysis. To illustrate this point, two example hierarchies are shown in Figure 9a and 9b below

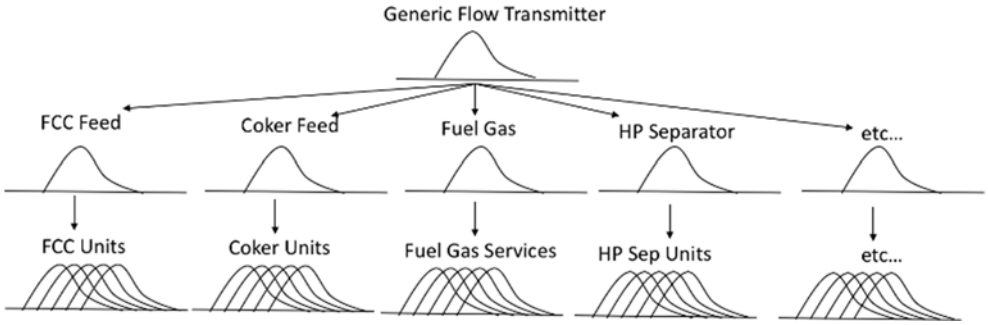

Figure 9a: Complex Hierarchy Based on Refinery Unit / Service

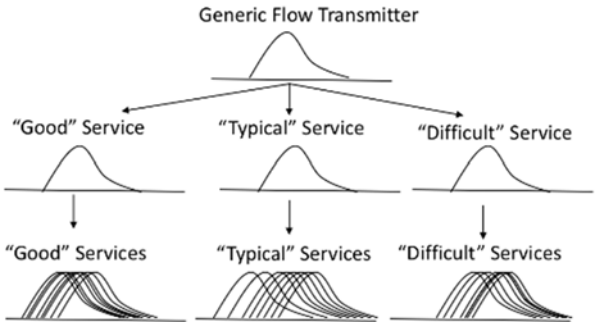

Figure 9b: Simple Hierarchy Based on Service Class

It is no coincidence that these diagrams look similar to the model in Figure 6. The subjective groupings in Figures 9a/b would be implemented using the same kind of models illustrated earlier.

The choice of grouping can depend on several factors, including but not limited to:

- The volume of data available (i.e. less data → more grouping)
- The resources available for data analysis
- How the data will be used (i.e. specific goals such as a prior use justification)

Note that the choice of analysis framework is by no means set in stone. As long as the original data is maintained, the analysis can always be repeated under a different framework as requirements and capabilities change. For example, a simple hierarchy can provide meaningful results quickly, then be modified after a larger volume of data has been accumulated.

6.2 Toward Dynamic SIL Verification

As already discussed, the process of updating with new evidence will lead progressively to lower uncertainty and narrower posterior distributions. The posterior credibility interval can be used as part of an IEC 61511 prior use justification, but the analysis does not need to end there. The updated failure rate distributions may be used to periodically re-evaluate SIS performance, which may lead to one or more of:

- Procedural changes to eliminate systematic failure mechanisms
- Decreased proof test intervals
- Design changes to replace poor performers or add fault tolerance.

In other literature, Hauge et al [7] propose a method that combines Bayesian updating with elements of SPC concepts to monitor performance and take corrective actions. The author of the present paper proposed a simple method for performing and updating SIL calculations using Monte Carlo simulation. [19]

The framework and techniques for dynamic SIL verification are in place, but the lack of commercial tools has so far kept them from enjoying widespread application outside of academia.

7 Conclusions

The international Safety Instrumented Systems standard IEC 61511 calls for using credible, traceable field failure rate data in SIL verification and requires that actual performance in operation be monitored versus design assumptions. These requirements have proven difficult for end users in part because of the large amount of sample data required for traditional frequentist methods.

The proposed Bayesian framework addresses the requirements by providing a cyclical updating process that allows industry knowledge about failure rates to be incorporated in a prior distribution and cyclical updated with new data as it becomes available. Even a simple single-stage framework is demonstrated to reduce data requirements by 40% by leveraging this prior knowledge.

The hierarchical Bayes framework expands on the single-stage model and allows data from other similar equipment to be leveraged in the updating process. This more complex model closes the loop on SIS performance by leveraging *all* available enterprise data in the updating process. Despite the complex mathematics involved, software tools using Markov Chain Monte Carlo (MCMC) algorithms make it practical to solve and update these models in seconds.

The hierarchical Bayes framework can be implemented step-by-step as part of an enterprise hierarchical grouping of equipment types. The structure and complexity of the hierarchy may depend on several factors, but the hierarchy may grow or be modified as requirements change. The Bayesian methodology provides a flexible, coherent framework for managing failure rate data in any enterprise.

8 References

[1] J.L. LaChance, et al. “Handbook of Parameter Estimation for Probabilistic Risk Assessment”. US Nuclear Regulatory Commission, NUREG/CR-6823, 2003.

[2] T.L. Chu, et al. "Traditional Probabilistic Risk Assessment Methods for Digital Systems." US Nuclear Regulatory Commission, NUREG/CR-6962, 2008.

[3] H. Dezfuli, et al. "Bayesian inference for NASA probabilistic risk and reliability analysis.", Available at https://ntrs.nasa.gov/search.jsp?R=20090023159 accessed on Feb 11, 2018.

[4] A. Gelman, et al. *Bayesian data analysis.* CRC press, Boca Raton, FL, 2014.

[5] D. Kelly and C. Smith. Bayesian inference for probabilistic risk assessment: A practitioner's guidebook. Springer Science & Business Media, 2011.

[6] M.S. Hamada, et al. *Bayesian reliability*. Springer Science & Business Media, 2008.

[7] S. Hauge, M. A. Lundteigen, and M. Rausand. "Updating failure rates and test intervals in the operational phase: A practical implementation of IEC 61511 and IEC 61508." I Reliability, Risk and Safety-Theory and Applications, Proceedings of the European Safety and Reliability Conference, ESREL, 2009.

[8] M.W. Bjartnes. Provision and Updating of Estimates of Reliability Parameters for Use in Reliability Analyses of Safety-Instrumented Systems. MS thesis. Institutt for produksjons-og kvalitetsteknikk, 2012.

[9] H. Jin, et al. "Quantification of organizational influences on failure rate: A Bayesian approach." Industrial Engineering and Engineering Management (IEEM), IEEE International Conference on. IEEE, 2012.

[10] A. Shafaghi. "Equipment failure rate updating - Bayesian estimation." Journal of hazardous materials, 2008.

[11] F. Khan, S. Rathnayaka, and S. Ahmed. "Methods and models in process safety and risk management: past, present and future." Process Safety and Environmental Protection, 2015.

[12] S. Kaplan. "On a Two-Stage Bayesian Procedure for Determining Failure Rates from Experimental Data." IEEE Transactions on Power Apparatus and Systems, 1983

[13] K. Pörn. "The two-stage Bayesian method used for the T-Book application." Reliability Engineering & System Safety, 1996.

[14] E.L. Droguett, F.J. Groen, and A. Mosleh. "Bayesian assessment of the variability of reliability measures." Pesquisa Operacional, 2006.

[15] D.L. Kelly and C.L. Smith. "Bayesian inference in probabilistic risk assessment - the current state of the art." Reliability Engineering & System Safety, 2009.

[16] P. Weber, et al. "Overview on Bayesian networks applications for dependability, risk analysis and maintenance areas." Engineering Applications of Artificial Intelligence, 2012.

[17] J. Kruschke. Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan. Academic Press, 2014.

[18] International Electrotechnical Commission. “Functional Safety - Safety Instrumented Systems for the Process Industry”. IEC 61511-1 Ed 2, IEC, Geneva, 2016.

[19] S. Thomas. “Practical Confidence Methods for SIS Performance Assessment”. Proceedings of ISA Process Control & Safety Symposium & Exhibition, Houston, 2016.