

Use of Bayesian Techniques for Reliability

Table of Contents

- Introduction
- A Numerical Example
- Normal Distribution Case
- Summary
- For Further Study
- About the Author
- Other START Sheets Available

Introduction

The world of statistics is traditionally divided into two mutually exclusive camps: Classical and Bayesian. They are divided on basically a principles issue, as will be explained.

The Classical statistician believes that all distribution parameters are fixed values (i.e., Mean Life is 100 hours). On the other hand, Bayesians believe that parameters are random variables and have a distribution of their own (i.e., Mean Life can range from 90 to 110 hours following some pattern). These two principles exclude each other.

In this START sheet, we present the Bayesian approach and some applications to reliability, leaving the subjective evaluation and interpretation of this concept to the reader. Perhaps the best way to introduce this concept is via an example. Let's assume that the life X of a device follows the Exponential distribution with parameter (Mean) θ . The density $f(x)$ of X is:

$$f(x) = \frac{1}{\theta} \text{Exp}\left(-\frac{x}{\theta}\right)$$

For the Classical statistician, the parameter Mean Life (θ) is a fixed number, say 100 hours. For the Bayesian statistician, in turn, the parameter Mean Life (θ) is a random variable having a (possibly different) statistical distribution with density $f(\theta)$.

For example, assume its distribution is Uniform (90, 110):

$$f(\theta) = \frac{1}{110 - 90} = \frac{1}{20}; \text{ for } 90 < \theta < 110$$

$$= 0 \text{ for all other values of } \theta$$

In practice, this means that the Mean Life θ of the device is not fixed but varies about value 100 as a Uniform random variable does, being as low as 90 hours and as high as 110. See Figure 1.

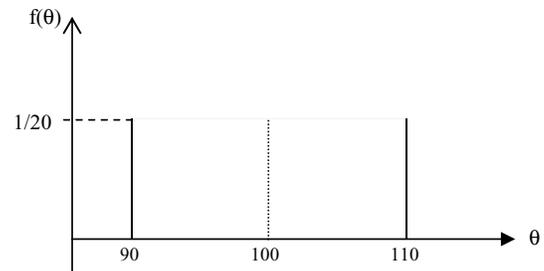


Figure 1. Representation of the Prior Function $f(\theta)$, known as the Prior Distribution. The probabilities associated with the Prior are referred to as "subjective probabilities" because they quantify some belief as to where the parameter Mean Life θ lays and how it's value varies.

Then, given a random sample of "n" device lives (denoted x_1, \dots, x_n) the conditional distribution of these data, given the parameter Mean Life (θ) is defined and denoted:

$$f(x_1, \dots, x_n | \theta) = f(x_1 | \theta) \dots f(x_n | \theta) = \frac{1}{\theta^n} \text{Exp}\left(-\frac{\sum x}{\theta}\right)$$

for some value of θ

We then calculate the Joint Distribution $f(x_1, \dots, x_n, \theta)$ of both the data and the Mean parameter θ , by multiplying their Conditional and Prior distributions:

$$f(x_1, \dots, x_n | \theta) = f(x_1, \dots, x_n | \theta) \times f(\theta)$$

$$= \frac{1}{\theta^n} \text{Exp}\left(-\frac{\sum x}{\theta}\right) \times \frac{1}{20} = \frac{1}{200^n} \text{Exp}\left(-\frac{\sum x}{\theta}\right)$$

for: $90 < \theta < 110$

We now integrate the Joint Distribution, for all possible values of the parameter Mean Life (θ) to obtain the “marginal” distribution of the sample data:

$$f(x_1, \dots, x_n) = \int f(x_1, \dots, x_n, \theta) d\theta$$

$$= \int_{90}^{110} \frac{1}{200^n} \text{Exp}\left(-\frac{\sum x}{\theta}\right) d\theta$$

We can now obtain the “Posterior” distribution of parameter θ , which is the distribution of the Mean Life θ , given the sample x_1, \dots, x_n that we have drawn:

$$f(\theta | x_1, \dots, x_n) = \text{Posterior} = \frac{\text{Joint Distribution}}{\text{Marginal Distribution}}$$

$$= \frac{f(x_1, \dots, x_n, \theta)}{f(x_1, \dots, x_n)} = \frac{\frac{1}{\theta^n} \text{Exp}\left(-\frac{\sum x}{\theta}\right)}{\int \frac{1}{\theta^n} \text{Exp}\left(-\frac{\sum x}{\theta}\right) d\theta}$$

The practical use of all of this theory in reliability studies is as follows. Assume that we believe that a parameter θ , the Mean Life of a device, lies in some range of values, with some given pattern. We can then define such range and pattern (its distribution) via the Prior $f(\theta)$ of this parameter. We can draw a sample of actual device lives and obtain via the above approach, an expression for the Posterior distribution of θ . Finally, this Posterior distribution can be continually refined with more incoming data and used for estimation. We can summarize this as a step-wise process:

1. Decide which distribution parameter (e.g., Mean θ) you want to study.
2. Define a Prior Distribution $f(\theta)$ for this parameter (e.g., Uniform (90, 110)).

3. Establish the Conditional Distribution $f(X | \theta)$ for the variable X (e.g., for the device life).
4. Obtain the Joint distribution for the sample data (x_1, \dots, x_n) and θ by multiplying their Prior and Conditional distributions, obtained in steps 2 and 3.
5. Obtain the Marginal distribution by integrating out the θ on the above Joint distribution.
6. Calculate the Posterior Distribution $f(\theta | x_1, \dots, x_n)$ of θ , by dividing the Joint by the Marginal distributions and use it for estimation and testing.

In applying this theory, the analyst encounters some practical problems. First, as seen in the preceding example, the closed form of the Posterior distribution $f(\theta | x_1, \dots, x_n)$ is not always easy to obtain. Sometimes the Prior distributions are defined using mathematical arguments, rather than statistical ones. This allows the analyst to obtain a closed form solution for the resulting Posterior distribution. However, if the Prior distribution is postulated just for analytical convenience and it is inappropriate, then the inferences made about the parameter, based on such inappropriate Priors, will also be erroneous.

A Numerical Example

The following numerical example helps to illustrate the Bayesian approach. Let's assume that a device has a mission time, $T = 100$ hours, and that past experience has shown that it fails with probability ρ (and survives with probability $= 1 - \rho$). Let's assume that such prior experience has also taught us that there are only three different device failure probabilities, ρ_i , and that these ρ_i , can occur with the probabilities shown in Figure 2 (denoted $f(\rho)$). This means that a device failure probability of 0.01 occurs 20% of the time, that a failure probability of 0.02 occurs 50% of the time and that a failure probability of 0.03 occurs 30% of the time.

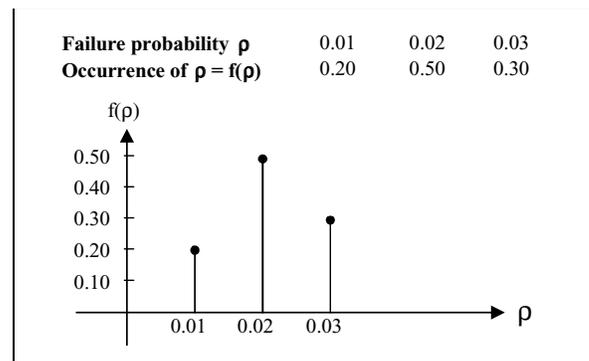


Figure 2. Prior Distribution of ρ

Assume we want to find the Bayes estimate for the true failure probability, ρ , of this device. To achieve this, we test a random sample of five operating devices and observe one failure.

Let's define the number of failed devices in our sample as X . All ($n = 5$) five test devices are randomly selected and each fails independently, with the same probability ρ . Hence, each device outcome is a Bernoulli trial and their sum (i.e., total number of failures $X = x$) is distributed as a Binomial, with $n = 5$ and probability ρ (denoted $B(x; n, \rho)$):

$$P(X = x | \rho) = f(x | \rho) = B(x; n, \rho)$$

$$= C_x^n \rho^x (1 - \rho)^{n-x}; x = 0, 1, 2, 3, 4, 5$$

For example, when the true failure probability is $\rho = 0.01$, the probability that we observe one failure in the random sample of five devices is:

$$f(x = 1 | \rho = 0.01) = B(1, 5, 0.01)$$

$$= C_1^5 0.01^1 (1 - 0.01)^{5-1} = 0.04803$$

By the same token, and applying the above formula with $\rho = 0.02$ and 0.03 , respectively, we get:

$$f(1 | 0.02) = B(1; 5, 0.02) = 0.09224$$

$$f(1 | 0.03) = B(1; 5, 0.03) = 0.13279$$

From the previous section we know that the Joint distribution $f(x, \rho) = f(x | \rho) * f(\rho)$, where the occurrence of the device failure probability is denoted $f(\rho)$. Therefore, we can write:

| ρ | 0.01 | 0.02 | 0.03 |
|--------------|----------|---------|---------|
| $f(1, \rho)$ | 0.009606 | 0.04612 | 0.03984 |

For example, for the case of one failure ($x = 1$) in the sample (our case of interest) when the true failure rate is $\rho = 0.01$ and its probability of occurrence is 0.2 , we obtain:

$$f(x = 1, \rho = 0.01) = f(1 | 0.01) * f(0.01) = 0.04803 * 0.2 = 0.009606$$

The Marginal distribution $g(1)$ of observing one failure ($x = 1$) in a random sample of five (irrespective of whether the true failure probability $\rho = 0.01, 0.02$, or 0.03) is the sum of the probability $f(\rho, x)$ for the number of failures observed ($x = 1$):

$$g(1) = \sum_{\rho} f(1, \rho) = 0.0096 + 0.0461 + 0.0398 = 0.095563$$

We can then obtain the Posterior distribution $f(\rho | x = 1)$ of the proportion of defectives ρ , in the case of testing and observing a

single failure ($x = 1$) in a random sample of five devices ($n = 5$). We do this via the Posterior distribution formula (i.e., dividing the Joint distribution by the Marginal distribution):

$$f(\rho | x = 1) = \frac{f(x = 1, \rho)}{g(x = 1)}; \rho = 0.01, 0.02, 0.03$$

We tabulate the values of the probabilities $f(\rho | x = 1)$:

| ρ | 0.01 | 0.02 | 0.03 |
|-------------------|-------|-------|-------|
| $f(\rho x = 1)$ | 0.100 | 0.483 | 0.417 |

By definition, the Bayes Point Estimate (ρ^*) of a parameter ρ is nothing else but the weighted average (expected value) of the Posterior Distribution of this parameter:

$$\rho^* = \sum \rho f(\rho | x)$$

Hence, the Bayes point estimate of the true device probability of failure ρ , based on a test that resulted in only one failure, from an experiment with five operating devices, when postulating the prior distribution $f(\rho)$ for the probability of failure ρ , is obtained by the following equation.

$$\rho^* = \sum \rho x f(\rho | x = 1)$$

$$= 0.01 \times 0.1005 + 0.02 \times 0.4826 + 0.03 \times 0.4169$$

$$= 0.023164$$

We can then compare the Bayes point estimator with the “classical” estimator (the sample proportion $x/n = 1/5 = 0.2$).

To provide an illustrative contrast, assume now that no failures were observed in the above-mentioned test ($x = 0$). In such case, the corresponding probability of obtaining no failures, given the failure probabilities $\rho = 0.01$ is:

$$P(X = 0 | \rho) = B(x = 0; n = 5, \rho = 0.01) = 0.9501$$

Applying the same formula for $\rho = 0.02$ and 0.03 , we obtain:

$$f(x = 0, \rho = 0.02) = 0.9039$$

$$f(x = 0, \rho = 0.03) = 0.8587$$

The corresponding Joint distribution for zero failures is now:

| ρ | 0.01 | 0.02 | 0.03 |
|--------------|--------|--------|--------|
| $f(0, \rho)$ | 0.1902 | 0.4520 | 0.2576 |

The Marginal distribution for zero failures is then:

$$g(0) = \sum_{\rho} f(0, \rho) = 0.1902 + 0.4520 + 0.2576 = 0.8998$$

Finally, the Posterior distribution for zero failures $f(\rho | x = 0)$ is:

| | | | |
|-------------------|--------|--------|--------|
| ρ | 0.01 | 0.02 | 0.03 |
| $f(\rho x = 0)$ | 0.2114 | 0.5023 | 0.2863 |

One interpretation of these results is that the original Prior probabilities $f(\rho)$ have been “updated” via the experiment, and the “new Prior” is now the Posterior probability $f(\rho | x)$.

We can also try to compare the point estimator ρ^* (from this second test, that has no failures) with the “Classical” point estimator (x/n) . But since now there are no failures, the sample proportion is zero and will not provide a useful estimate of ρ .

In Table 1, we complete the Bayes analyses results for all the other cases ($x = 2, 3, 4$, and 5 failures) in this experiment.

Table 1. Table of Joint and Marginal Probabilities for the Experiment

| Fail. Prob. | Occur-FP | $f(0, \rho)$ | $f(1, \rho)$ | $f(2, \rho)$ | $f(3, \rho)$ | $f(4, \rho)$ | $f(5, \rho)$ |
|------------------------|----------|--------------|--------------|--------------|--------------|--------------|--------------|
| 0.01 | 0.200000 | 0.190198 | 0.009606 | 0.000194 | 0.000002 | 0.000000 | 0.000000 |
| 0.02 | 0.500000 | 0.451961 | 0.046119 | 0.001883 | 0.000039 | 0.000001 | 0.000000 |
| 0.03 | 0.300000 | 0.257620 | 0.039838 | 0.002464 | 0.000076 | 0.000000 | 0.000000 |
| Marginal Probs. | | 0.899779 | 0.095563 | 0.004541 | 0.000117 | 0.000001 | 0.000000 |

We can see, from Table 1, that $X = 4$ and 5 cases are nil. Table 2 shows the posterior distributions and the Bayes estimates of ρ .

Table 2. Table of Posterior distributions and Bayes Estimates of ρ

| F.P. | $f(\rho x = 0)$ | $f(\rho x = 1)$ | $f(\rho x = 2)$ | $f(\rho x = 3)$ |
|------------------------------------|-------------------|-------------------|-------------------|-------------------|
| 0.01 | 0.211383 | 0.10052 | 0.042725 | 0.017138 |
| 0.02 | 0.502302 | 0.482599 | 0.414584 | 0.329906 |
| 0.03 | 0.286315 | 0.41688 | 0.542692 | 0.652956 |
| Bayes (ρ^*) | 0.020749 | 0.023164 | 0.025000 | 0.026358 |

Summarizing: in the absence of any particular information about the parameters in question, the Classical estimate is not only the best one but the only one available. When the Prior distribution is an accurate reflection of reality, Bayes estimate is more efficient than the Classical. However, if the Prior distribution is inappropriate or inaccurate, the Bayes estimate may also be significantly in error.

Normal Distribution Case

In practice, the Prior distribution $f(\theta)$ used to describe the variation pattern (or belief) of the unknown parameter θ , is often the Normal distribution. When applicable, the Normal greatly simplifies all the preceding derivations, as we will show through another example.

Assume that the unknown parameter of interest θ is now the mean of a Normal (θ, σ^2) population. Assume that the Prior distribution $f(\theta)$ of parameter θ is, itself, also Normal with mean μ and known standard deviation δ . Assume that we have obtained the sample average \bar{x} from a random sample of size “ n ” from the population of interest (i.e., of the Normal (θ, σ^2)).

Then, the Posterior Distribution of the population mean θ is also the Normal, and its parameters μ^* and σ^* , respectively, will be given by the formulas:

$$\mu^* = \frac{n\bar{x}\delta^2 + \mu\sigma^2}{n\delta^2 + \sigma^2} \text{ and } \sigma^{*2} = \frac{\delta^2\sigma^2}{n\delta^2 + \sigma^2}$$

Graphical depictions of the Prior and Posterior distributions of θ are shown in Figure 3.

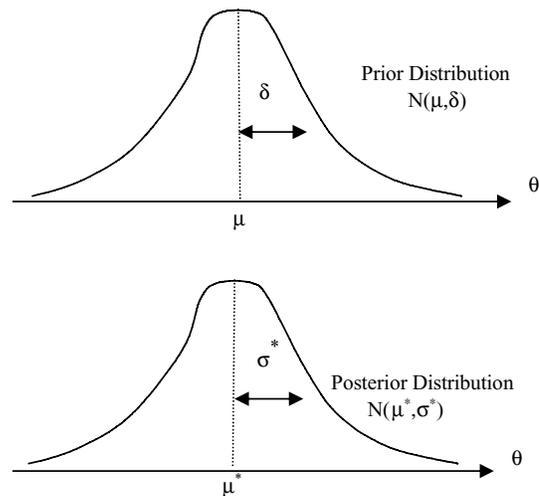


Figure 3. Prior and Posterior Distributions for the parameter Population Mean θ

Therefore, a $100(1 - \alpha)\%$ Bayesian confidence interval (CI) for the true mean θ (say, an 95% CI for an $\alpha = 0.05$ and $z_{\alpha/2} = 1.96$)

can be constructed from the Posterior distribution of θ using the following equation.

$$\mu^* \pm z_{\alpha/2} \times \sigma^* = (\mu^* - z_{\alpha/2} \times \sigma^*, \mu^* + z_{\alpha/2} \times \sigma^*)$$

We will now illustrate the use of all this information with a practical reliability example. Assume that a manufacturer of an electronic amplifier wants to estimate the length of the life X of this device. Assume that, for historical reasons, it is believed that the distribution of the life is Normal, with unknown mean θ and known standard deviation $\sigma = 150$ hours.

Assume that, based on the same historical reasons, this electronic amplifier manufacturer believes that the Prior Distribution for the unknown mean life θ is also Normal, with known mean $\mu = 1000$ hours and standard deviation $\delta = 100$ hours.

Assume now that the manufacturer measures the lives from a random sample of $n = 25$ electronic amplifiers (if the sample size is $n \geq 30$, then the distribution of life X doesn't even need to be Normal and the unknown variance σ^2 can now be approximated by the sample variance s^2). Assume that this sample yields an average life $\bar{x} = 900$ hours. Using all the information, we derive a $100(1 - \alpha)\%$ Bayesian CI for the amplifier mean life θ , as:

$$\mu^* - z_{\alpha/2} \times \sigma^* < \theta < \mu^* + z_{\alpha/2} \times \sigma^*$$

According to the preceding, the resulting Normal Posterior Distribution will have the following parameters.

$$\mu^* = \frac{n\bar{x}\delta^2 + \mu\sigma^2}{n\delta^2 + \sigma^2} = \frac{25 \times 900 \times 100^2 + 1000 \times 150^2}{25 \times 100^2 + 150^2} = 908.26$$

$$\sigma^{*2} = \frac{\delta^2\sigma^2}{n\delta^2 + \sigma^2} = \frac{100^2 \times 150^2}{25 \times 100^2 + 150^2} = 825.69$$

Therefore, for $\alpha = 0.05$ and $z_{\alpha/2} = 1.96$, a 95% Bayesian CI for the unknown mean life θ is:

$$\begin{aligned} \mu^* \pm z_{\alpha/2} \times \sigma^* &= 908.26 \pm 1.96 \sqrt{825.69} = 908.26 \pm 56.32 \\ &= (851.94, 964.58) \end{aligned}$$

Had we chosen to ignore (or had no access to) this information about the Prior Distribution, we could have still obtained a 95% Classic CI for θ as:

$$\bar{x} \pm z_{\alpha/2} \times \sigma / \sqrt{n} = 900 \pm 1.96 \times \frac{150}{\sqrt{25}} = 900 \pm 58.8$$

$$= (841.2, 958.8)$$

As we have advised several times in the previous sections, if the assumed Prior Distribution is appropriate and the Prior's parameters are accurate, then the Bayesian CI is more efficient. Such efficiency becomes evident in the fact that the Bayesian CI is narrower (notice its smaller half-width, 56.32) than the Classical CI (half-width is 58.8). This reduction in the CI width occurs because the Bayesian approach uses more information than the Classical (e.g., the shape of the Prior distribution as well as its parameters).

In addition, the Bayesian point estimator (908.26) of the parameter (θ) differs from the Classical point estimator, given by the sample average (900). This difference is due to Bayes weighing the data via the Prior standard deviation and mean.

However, if any or all of the information regarding Bayes Prior (either the distribution or its parameters) which are used for deriving the Bayesian CI, is inaccurate or incorrect, then the CI obtained with such information will also be incorrect.

In such case, the Bayes point estimator (908.26) will be biased. In addition, the resulting CI, with its narrower half width and centered on this biased point estimator, will most likely not cover the true parameter.

Summary

Through the discussion and the examples presented here, we have seen how the Bayesian approach is applied to data analysis and how Bayesian point estimators and confidence intervals are obtained and updated. In particular we have seen how they are implemented to a life data analysis example.

Bayesian techniques are widely used by some practitioners in industrial statistics and can become very useful analysis tools. However, it is always necessary to bear in mind this large dependency that exists between the applicability of the postulated prior distributions and its parameters, and the results obtained, through them (e.g., our estimations).

We have just touched on a few relevant concepts of Bayesian statistics, a topic that is very extensive and parallels Classical statistics. In **For Further Study**, we list the usual background readings plus several references that can serve as starting points for those interested in studying in this topic at a more advanced level.

For Further Study

1. Probability and Statistics for Engineers and Scientists, Walpole and Myers, Prentice Hall, NJ, 1988.
2. An Introduction to Probability Theory and Mathematical Statistics, Rohatgi, V.K. Wiley, NY, 1976.

3. Methods for Statistical Analysis of Reliability and Life Data, Mann, N., R. Schafer, and N. Singpurwalla, John Wiley, NY, 1974.
4. Reliability and Life Testing Handbook, (Vol. 2), Kececioglu, D., Editor, Prentice Hall, NJ, 1994.
5. Practical Statistical Tools for Reliability Engineers, Coppola, A., RAC, 1999.
6. A Practical Guide to Statistical Analysis of Material Property Data, Romeu, J.L. and C. Grethlein, AMPTIAC, 2000.
7. Statistical Confidence, Romeu, J.L., RAC START, Volume 9, Number 4, <http://rac.alionscience.com/pdf/STAT_CO.pdf>.

About the Author

Dr. Jorge Luis Romeu has over thirty years of statistical and operations research experience in consulting, research, and teaching. He was a consultant for the petrochemical, construction, and agricultural industries. Dr. Romeu has also worked in statistical and simulation modeling and in data analysis of software and hardware reliability, software engineering and ecological problems.

Dr. Romeu has taught undergraduate and graduate statistics, operations research, and computer science in several American and foreign universities. He teaches short, intensive professional training courses. He is currently an Adjunct Professor of Statistics and Operations Research for Syracuse University and a Practicing Faculty of that school's Institute for Manufacturing Enterprises. Dr. Romeu is a Chartered Statistician Fellow of the Royal Statistical Society, Full Member of the Operations

Research Society of America, and Fellow of the Institute of Statisticians.

Romeu is a senior technical advisor for reliability and advanced information technology research with Alion Science and Technology. Since joining Alion in 1998, Romeu has provided consulting for several statistical and operations research projects. He has written a State of the Art Report on Statistical Analysis of Materials Data, designed and taught a three-day intensive statistics course for practicing engineers, and written a series of articles on statistics and data analysis for the AMPTIAC Newsletter and RAC Journal.

Other START Sheets Available

Many Selected Topics in Assurance Related Technologies (START) sheets have been published on subjects of interest in reliability, maintainability, quality, and supportability. START sheets are available on-line in their entirety at <<http://rac.alionscience.com/rac/jsp/start/startsheet.jsp>>.

For further information on RAC START Sheets contact the:

Reliability Analysis Center
 201 Mill Street
 Rome, NY 13440-6916
 Toll Free: (888) RAC-USER
 Fax: (315) 337-9932

or visit our web site at:

<<http://rac.alionscience.com>>



About the Reliability Analysis Center

The Reliability Analysis Center is a Department of Defense Information Analysis Center (IAC). RAC serves as a government and industry focal point for efforts to improve the reliability, maintainability, supportability and quality of manufactured components and systems. To this end, RAC collects, analyzes, archives in computerized databases, and publishes data concerning the quality and reliability of equipments and systems, as well as the microcircuit, discrete semiconductor, and electromechanical and mechanical components that comprise them. RAC also evaluates and publishes information on engineering techniques and methods. Information is distributed through data compilations, application guides, data products and programs on computer media, public and private training courses, and consulting services. Located in Rome, NY, the Reliability Analysis Center is sponsored by the Defense Technical Information Center (DTIC). Alion, and its predecessor company IIT Research Institute, have operated the RAC continuously since its creation in 1968. Technical management of the RAC is provided by the U.S. Air Force's Research Laboratory Information Directorate (formerly Rome Laboratory).