DISCREPANCY IN THE RESULTS GENERATED FROM BAYESIAN UPDATE OF LOGNORMAL DISTRIBUTIONS USING DIFFERENT CALCULATION TECHNIQUES

James C. Lin[1] and David H. Johnson[2]
ABSG Consulting Inc., 300 Commerce Drive, Suite 200, Irvine, California, 92602
[1]jlin@absconsulting.com and [2]djohnson@absconsulting

Since the posterior of a lognormal prior distribution is not a member of the lognormal distribution family, Bayesian update of a lognormal prior distribution is most commonly performed either numerically or analytically using a converted Beta/Gamma distribution. However, neither of these approaches is perfect because the numerical Bayesian update of the lognormal distribution involves discretization error and the use of a surrogate Beta/Gamma distribution for analytical update involves distribution conversion error. In selected applications, significant differences are observed in the mean values of the posterior distribution updated from a lognormal prior using the above two approaches. When the central tendency of the experience data approaches or exceeds (from the standpoint of graphical display or visual observation) the end of the distribution tails, the result of the Bayesian update may become very sensitive to the distribution discretization or conversion errors in this region. Therefore, for cases in which the experience data is not reasonably or fully encompassed by the prior probability distributions, Bayesian update using different PRA computer software with different calculation techniques can lead to very large discrepancy in the results of the Bayesian update, e.g., posterior mean. It is recommended that all computer codes that perform Bayesian data update should include sufficient diagnostics to alert the users when the input experience data is not fully compatible with the input prior probability distribution.

Keywords: Bayesian Update, Beta Distribution, Gamma Distribution, Lognormal Distribution, Probabilistic Risk Assessment.

I. INTRODUCTION

Bayesian analysis has been applied extensively in nuclear plant probabilistic risk assessments (PRA) to combine generic data and plant-specific experience for data update. When little plant-specific data is available, the results of the Bayesian update will be dominated by the generic prior distribution for the parameters analysed. If, however, substantial plant-specific experience (i.e., very strong evidence) exists, the influence of the prior distribution could become negligible and the updated distribution could be driven primarily by the plant-specific data.

Typical PRA data updated using the Bayesian analysis technique include failure rates, initiating event frequencies, etc. The generic prior distributions for these parameters are often characterized by lognormal, gamma, and beta distributions. Gamma distributions can be used to represent the operating failure rates and initiating event frequencies. Beta distributions are generally used to characterize the demand failure rates or any other parameters with values between 0 and 1.0. Lognormal distributions can, however, be used to represent initiating event frequencies and both operating and demand failure rates. As such, lognormal distribution is used more widely in PRA than any other analytical distributions.

The Bayesian updates of the Gamma and Beta prior distributions are quite straightforward because the parameter values (e.g., mean) of their conjugate posterior distributions can be calculated by closed-form expressions based on the prior distribution parameter values and the experience data. Since the posterior of a lognormal prior distribution is not a member of the lognormal distribution family, Bayesian update of a lognormal prior distribution is most commonly performed either numerically or using a converted Beta/Gamma distribution. To perform Bayesian update numerically, the lognormal distribution can be discretized into a number of bins each represented by a probability and a value. The probabilities of this set of bins can then be updated with the plant-specific experience data through the Bayes theorem. Another method to perform Bayesian update is to first convert the lognormal prior distribution to a Beta or Gamma distribution by the moments matching technique; i.e., matching the mean and variance. The converted Beta/Gamma distribution is then updated to its conjugate posterior distribution, and the mean of which can be calculated using closed-form expressions and is used to represent the mean of the posterior distribution. Numerically, neither of these approaches is perfect because the numerical

Bayesian update of the lognormal distribution involves discretization error and the use of a surrogate Beta/Gamma distribution involves distribution conversion error.

In some applications, significant differences are observed in the mean values of the posterior distribution updated from a lognormal prior using the above two approaches. This paper investigates the possible cause of these differences.

II. THREE EXAMPLE CASES WITH SIGNIFICANT DIFFERENCE

Table I shows the results of three example Bayesian update cases using the following two different calculation techniques:

- Numerical update using discretized bins to represent the lognormal distribution.
- Conjugate update using the converted Beta/Gamma distribution with matching mean and variance.

The numerical update approach discretizes the lognormal distribution into a significant number of bins; e.g., 100 bins. A Bayesian update calculation is then performed for each of these 100 bins using the Binomial (for demand failure rate) or Poisson (for run failure during operation) likelihood function to derive the posterior bin probabilities. The mean value of the posterior distribution is calculated by summing the products of bin value and bin probability for these distribution bins.

The conjugate update process is to first convert the lognormal distribution into a corresponding Beta (for demand failure) or Gamma (for run failure during operation) distribution with the identical mean and variance. The converted Beta/Gamma distribution is then updated analytically with the experience data. The posterior mean is calculated using a closed-form expression for the conjugate posterior distribution.

The first percent difference shown in the last column of Table I (ranging from 38% to 130%) was calculated using the average of the two updated mean values as the base value; i.e., by dividing the difference in the updated mean values by the average of the two updated mean values. If the Beta/Gamma updated mean values are used as the base values, the numerically updated mean values in these three cases are actually about 47% to 373% greater than the updated mean values for the conjugate posterior distributions. Of course, this significant difference only occurs for selected cases between the Bayesian updates of the lognormal distribution using the two different calculation techniques described in the preceding.

TABLE I. Three Example Cases of Bayesian Update Using Different Calculation Techniques

Component Failure Rate	Lognormal Distribution		No. of Start/Run Failures	No. of Starts/Run Hours	Numerically Updated Mean	Beta/Gamma Updated Mean	Percent Difference, %
	Prior Mean	Prior EF					
MPFS – Standby MDP Fails to Start	3.29E-03	7.7	1	3451	5.28E-04	3.59E-04	38/47
MPTS – Standby MDP Fails to Run during Operation	3.42E-05	6.1	0	54887.98	1.18E-05	6.32E-06	60/87
GTTS – Gas Turbine Fails to Run during Operation	8.48E-03	11.04	0	493.36	1.25E-03	2.64E-04	130/373

As can be seen in Table I, the experience data for standby motor-driven pump (MDP) fails to start (MPFS) centers around 2.9E-04, which is located toward the lower tail of the prior distribution. Due to the low likelihood of observing the experience data in the upper region of the prior distribution, the posterior distribution pulls back the upper tail of the prior distribution to be more centered around the mean value of the experience data as shown in Figure 1. However, the lower tail of the posterior distribution is somewhat bounded by the lower tail of the prior distribution (with the 5^{th} percentile at around 2.0E-04). This results from the low prior probabilities for the bins in this region. As such, the lower tail of the posterior distribution does not extend further lower than that in the prior distribution. With the experience data locating near the lowest tail of the prior distribution, it is expected that, due to the lower prior bin probabilities in this region, the result of the Bayesian update may be relatively sensitive to any errors in distribution discretization or conversion.

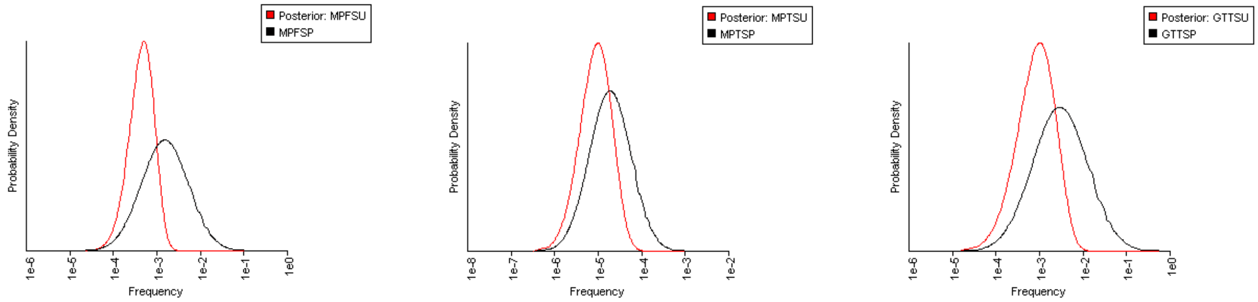

Figure 1. Comparison of Prior and Posterior Distributions

For standby MDP fails to run during operation (MPTS), the experience data indicates a central tendency likely to be significantly lower than 1.8E-5 (which is lower than both mean and median of the prior distribution). As such, the upper tail of the prior distribution is also pulled back in the posterior distribution to be more consistent with the experience data. Again, the lower tail of the posterior distribution is limited by the lower tail of the prior distribution. Similarly, for gas turbine fails to run during operation (GTTS), the lower tail of the posterior is pretty much bounded by the lower tail of the prior distribution and the upper tail of the prior distribution is pulled back to create the posterior distribution because the experience data shows a central tendency significantly lower than 2.0E-3 (which is lower than both mean and median of the prior distribution).

Based on the preceding observations, it appears that the results of the Bayesian update of the three examples above indicates that, if the central tendency of the experience data is very close to the very end of the lower tail of the prior distribution, the Bayesian update results could be very sensitive to any numerical errors in the Bayesian update calculation process. As different types of numerical errors are involved in the two different Bayesian update techniques described in the preceding, significant discrepancies in the posterior mean could result in these cases.

III. ADDITIONAL TEST CALCULATION CASES

To more clearly identify the significant discrepancies in the Bayesian updated mean values the lognormal distributions using the two different calculation approaches (e.g., numerical update is used in the RISKMAN™ software[1] and conjugate update is used in the CAFTA software[2]), additional test cases were calculated and shown in Table II and Figure 1.

TABLE II. Additional Test Cases of Bayesian Update Using Different Calculation Techniques

Case	Lognormal Distribution		No. of Run Failures	No. of Run Hours	Numerically Updated Mean	Gamma Updated Mean	Percent Difference, %
	Prior Mean	Prior EF					
1	1.00E-05	10	0	10	1.00E-05	9.99E-06	0/0
2	1.00E-05	10	0	100	9.96E-06	9.94E-06	0/0
3	1.00E-05	10	0	1,000	9.52E-06	9.43E-06	1/1
4	1.00E-05	10	0	10,000	7.31E-06	6.21E-06	16/18
5	1.00E-05	10	0	100,000	3.38E-06	1.41E-06	82/140
6	1.00E-05	10	0	1,000,000	9.05E-07	1.61E-07	139/461
7	1.00E-05	10	0	10,000,000	1.70E-07	1.64E-08	165/938
8	1.00E-05	10	0	100,000,000	3.81E-08	1.64E-09	183/2,222
9	1.00E-05	10	0	1,000,000,000	3.49E-08	1.64E-10	198/21,170
10	1.00E-05	10	1	10	6.50E-05	7.09E-05	9/8
11	1.00E-05	10	1	100	6.41E-05	7.05E-05	10/9
12	1.00E-05	10	1	1,000	5.60E-05	6.69E-05	18/16
13	1.00E-05	10	1	10,000	2.79E-05	4.41E-05	45/37
14	1.00E-05	10	1	100,000	8.03E-06	1.00E-05	22/20
15	1.00E-05	10	1	1,000,000	1.57E-06	1.15E-06	31/37

TABLE II. Additional Test Cases of Bayesian Update Using Different Calculation Techniques

Case	Lognormal Distribution		No. of Run Failures	No. of Run Hours	Numerically Updated Mean	Gamma Updated Mean	Percent Difference, %
	Prior Mean	Prior EF					
16	1.00E-05	10	1	10,000,000	2.47E-07	1.16E-07	72/113
17	1.00E-05	10	1	100,000,000	4.06E-08	1.16E-08	111/249
18	1.00E-05	10	1	1,000,000,000	3.49E-08	1.16E-09	187/2,898
19	1.00E-05	10	2	10	2.26E-04	1.32E-04	53/71
20	1.00E-05	10	2	100	2.23E-04	1.31E-04	52/70
21	1.00E-05	10	2	1,000	1.96E-04	1.24E-04	45/58
22	1.00E-05	10	2	10,000	7.08E-05	8.19E-05	15/14
23	1.00E-05	10	2	100,000	1.44E-05	1.86E-05	25/23
24	1.00E-05	10	2	1,000,000	2.33E-06	2.13E-06	9/9
25	1.00E-05	10	2	10,000,000	3.31E-07	2.16E-07	42/53
26	1.00E-05	10	2	100,000,000	4.46E-08	2.16E-08	69/106
27	1.00E-05	10	2	1,000,000,000	3.49E-08	2.16E-09	177/1,513

In all of the above calculations, the 5th percentile, median, and 95th percentile of the prior distribution are about 3.8E-7, 3.8E-6, and 3.8E-5, respectively.

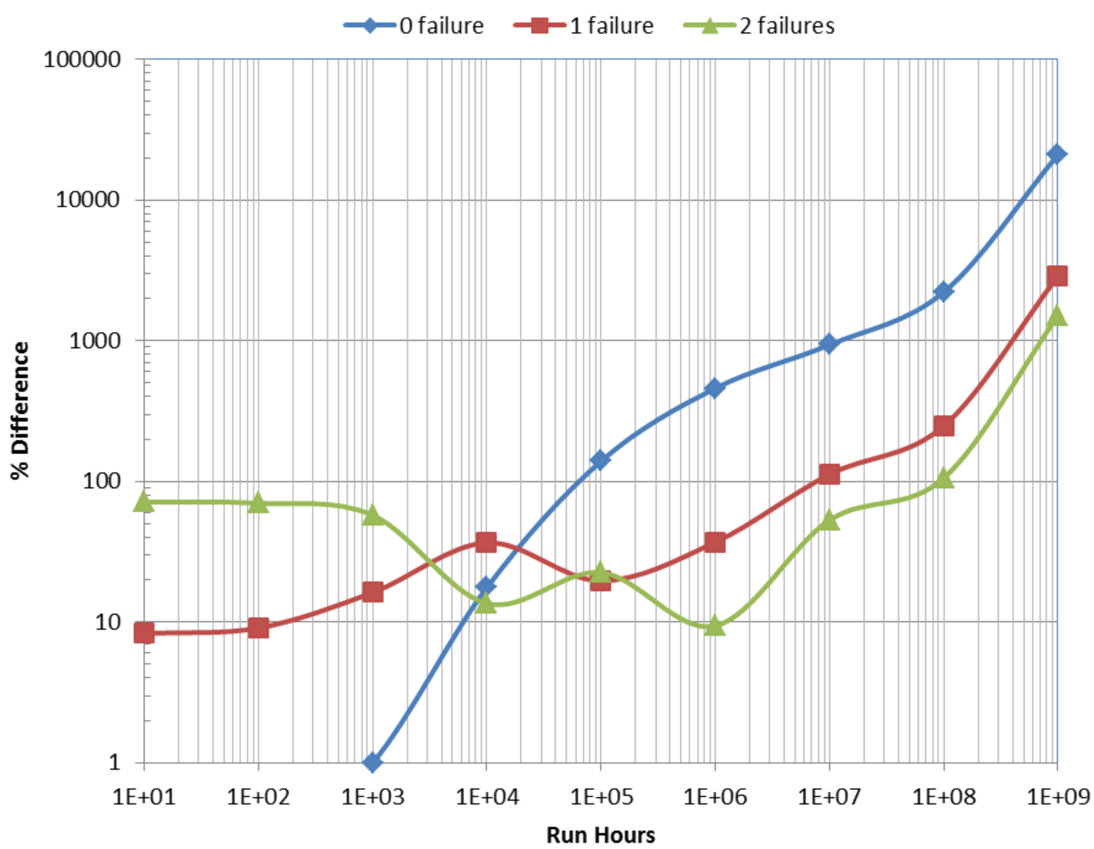

Figure 2. Percent Difference in Bayesian Updated Mean Using Two Different Approaches

For the cases involving two operating failures, the lowest discrepancy in the posterior means between the two calculation approaches occurs when the central tendency of the experience data is closest to the central tendency of the prior distribution. As the central tendency of the experience data deviates from the central tendency of the prior distribution, the percent difference in the posterior mean of the posterior distributions increases. For Case 27, the central tendency of the experience data (around 2.0E-9) lies essentially outside the lower tail of the prior distribution. As a result, the percent difference in the

Bayesian updated means of the posterior distributions resulting from the two different calculation approaches is in excess of 1,500%.

Figure 3 shows the prior/posterior distributions for Cases 27, 26, and 25 with 2 run failures in 1,000,000,000, 100,000,000, and 10,000,000 operating hours, respectively. After accounting for both the prior distribution and the operating experience, the posterior distributions in these three cases deviate substantially from the prior distributions. The bulk of the posterior distributions were generated by the weak prior probabilities in the region either near or beyond the lowest tail of the prior distributions and by the strong operating experience data in the same region. Incidentally, the percent inaccuracy in the low bin probabilities in this tail region may also be the greatest, which thus may contribute the most to the discrepancy in the Bayesian update posterior mean using different calculation techniques.

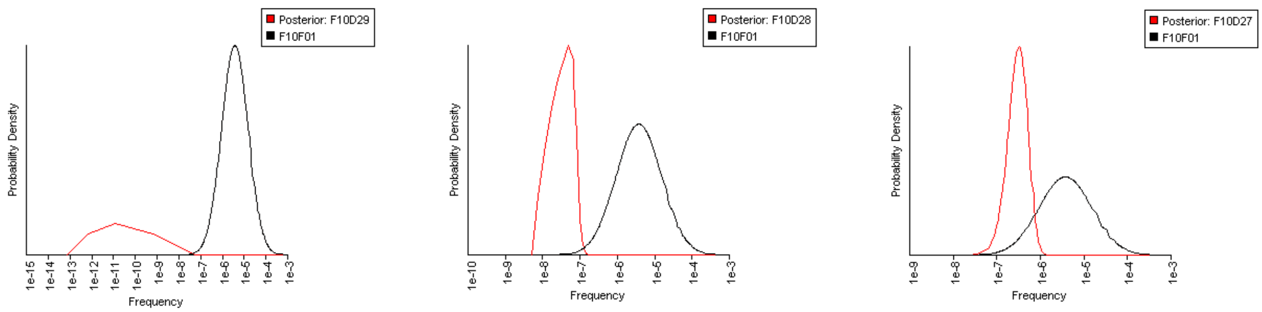

Figure 3. Prior (1E-5, EF=10) and Posterior Distributions (2 failures in 1E9/1E8/1E7 hours)

Shown in Figure 4 are the prior/posterior distributions for Cases 19, 20, and 21 with 2 failures in 10, 100, and 1,000 operating hours, respectively, where the discrepancies have also increased as the central tendency of the experience data moves higher beyond the region with significant bin probabilities in the upper tail of the prior distribution. However, the upper tail of the posterior distribution is still to some extent influenced by the prior distribution.

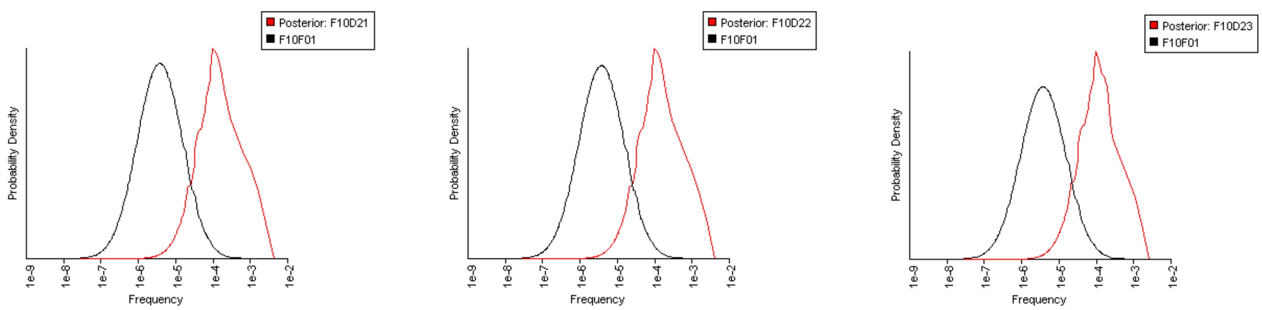

Figure 4. Prior (1E-5, EF=10) and Posterior Distributions (2 failures in 10/100/1,000 hours)

The cases involving 1 operating failure exhibit similar pattern to those of the 2 run failure cases as more operating hours is accumulated in the experience data; i.e., the discrepancy increases significantly as the central tendency of the experience data approaches and exceeds the very end of the lower tail of the prior distribution. However, on the other hand, as the central tendency of the experience data approaches the upper tail, the upper tail of the posterior distribution is extended to higher values. Yet, the upper tail of the posterior distribution for the one-failure cases is more limited by the upper tail of the prior distribution due to the weaker evidence than the two-failure cases. Figure 5 shows the prior/posterior distributions for Cases 10, 11, and 12 with 1 failure in 10, 100, and 1,000 operating hours, respectively.

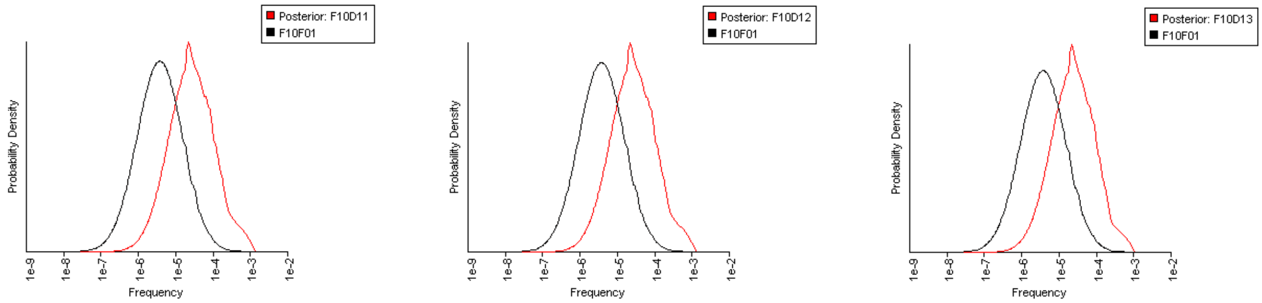

Figure 5. Prior (1E-5, EF=10) and Posterior Distributions (1 failures in 10/100/1,000 Hours)

For the cases of 0 failure, Cases 1, 2, and 3 (0 failure in 10, 100, and 1,000 operating hours, respectively) result in almost identical posterior distribution as the prior distribution. This is because 0 failure in 10, 100, and 1,000 hours implies very, very low likelihood of observing failure in region below 1.0E-3, which is completely consistent with the prior probability distribution. As the experience data changes to 0 failures in 10,000 hours, the central tendency of the operating data is still compatible and consistent with the central tendency of the prior distribution. As such, the posterior distribution does not differ much from the prior distribution. As the operating hours continue to increase with no failures present, the central tendency of the experience data begins to deviate significantly from that reflected in the prior distribution. Eventually, it not only approaches but also exceeds the end of the lower tail of the prior distribution. Since the experience data lies essentially outside of the prior distribution (as far as graphical display or visual observation is concerned), the very low prior probabilities in this region results in very sensitive response of the Bayesian updated means to the numerical discretization or distribution conversion errors. This then leads to the large percent difference in the Bayesian updated means obtained using the two different Bayesian update calculation techniques.

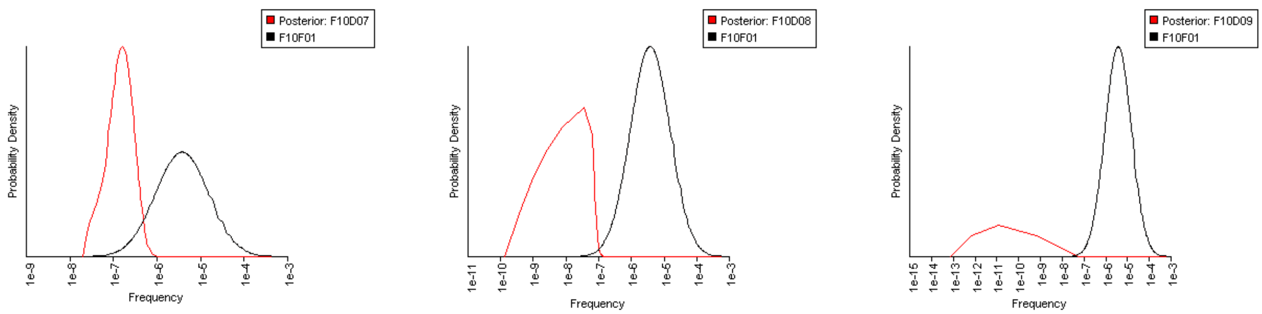

Figure 6. Prior (1E-5, EF=10) and Posterior Distributions (0 failures in 1E9/1E8/1E7 hours)

The above discussion is also reflected in Figure 2. As the central tendency of the experience data approaches or exceeds the tails of the prior distribution, the results of the Bayesian updated distribution and its mean value could be very sensitive to any numerical errors involved in the Bayesian update calculation techniques, thus resulting in large discrepancy in the results between different calculation techniques.

IV. COMPUTER IMPLEMENTATION OF BAYESIAN UPDATE PROCESS

Bayesian data update calculation is included in many PRA computer codes. However, most of these programs do not provide sufficient information for the users to discover any incompatibility in the input and abnormality in the results of the Bayesian update calculations.

For the Bayesian update of lognormal distributions, the Electric Power Research Institute CAFTA software includes a feature that takes as the input the mean and variance of a lognormal distribution. For component failure rates or initiating event frequencies, the software calculates the values of parameters α and β for a Gamma distribution with the same mean and variance as the lognormal distribution. Given the values for parameter α and β as well as the experience data in terms of the number of failures and the associated operating hours, the program performs the conjugate Bayesian update of the Gamma distribution. The updated mean and variance of the posterior Gamma distribution are then used as the mean and variance of the posterior lognormal distribution. The user interface in the CAFTA program for this feature includes only a dialogue table for the input of mean and variance of the lognormal prior, the number of failures and operating hours for the experience data, and the output mean and variance of the Gamma posterior distribution. For component demand failure probabilities or other parameters with values between 0 and 1, the Bayesian update calculation process and user interface are similar except that the Beta distribution is used.

The CAFTA software does not provide as a user interface the graphical display of the prior distribution, the central tendency of the experience data, and the posterior distribution. The user just inputs the data, obtains the posterior mean, and uses the Bayesian update results as is. As such, the user is not alerted with any information if and when the central tendency of the experience data approaches or exceeds the tails of the prior probability distribution, which would lead to an inaccurate posterior distribution and mean if the prior distribution is not adjusted to more fully encompass the central tendency of the experience data. Additional information should be provided to the user to enable the user to identify any incompatible input in which the experience data is not reasonably encompassed by the prior distribution.

The RISKMAN software includes both the one-stage and two-stage Bayesian update features. Similar to some of the Excel spreadsheet developed for the one-stage Bayesian update calculations, the prior distribution is discretized into a set of bins. Each of these bins is associated with a parameter value (e.g., component failure rate) and a probability. The probabilities of these bins are then updated with the experience data using a likelihood function to obtain the posterior probabilities for their corresponding bin values. The user interface includes the bin probabilities for both the prior and the posterior distributions. In addition, both the prior and posterior distributions are displayed graphically to show the users how the distribution changes by the experience data. However, since the numerical discretization technique is used, the tails of the prior and posterior distributions are bounded by the values of the end bins. If the experience data exceeds the bounds of these end bins, the posterior distribution will still be bounded by the values of the end bins in the prior distribution. The graphical display of the prior and posterior distributions may not fully alert the users that the central tendency of the experience data has exceeded the tails of the prior distribution. Further, if the central tendency of the experience data does not exceed but approaches the values of the end bins, the results would be very sensitive to the discretization error because the prior probabilities associated with the bins close to the tails are relatively small. Therefore, although the prior and posterior probabilities of the discretized bins are graphically displayed in the output, the program can still be further improved to better alert the users about the cases in which the experience data entered in the input approaches or exceeds the tails of the prior distributions.

V. CONCLUSION

Because the lognormal distribution is not a conjugate function, Bayesian update of the lognormal distribution cannot be performed analytically. Numerical method or alternate technique used for the Bayesian update of the lognormal distributions typically involves such inaccuracy as distribution discretization or distribution conversion error. The percent inaccuracy of these errors is often most prominent near the tails of the distribution.

One of the most important considerations in Bayesian data update analysis is that the prior probability distribution should reasonably encompass the central tendency of the experience data. However, most of the PRA computer programs that include the capability for Bayesian data update do not provide sufficient information to alert the users when the input information for the experience data is not fully compatible with the parameters for the prior probability distribution. In particular, when the central tendency of the experience data approaches or exceeds (from the standpoint of graphical display or visual observation) the end of the distribution tails (i.e., the experience data is not reasonably encompassed by the prior probability distribution), the result of the Bayesian update may become very sensitive to the distribution discretization or conversion errors in this region. Since different PRA computer programs use different calculation techniques for Bayesian update, cases in which the experience data is not reasonably or fully encompassed by the prior probability distributions can lead to very large discrepancy in the results of the Bayesian update (e.g., posterior mean) produced by different PRA computer software.

It is recommended that all computer codes that perform Bayesian data update should include sufficient diagnostics to alert the users when the input experience data is not fully compatible with the input prior probability distribution; i.e., when the central tendency of the experience data is not reasonably encompassed by the prior probability distribution.

REFERENCES

1. ABSG Consulting Inc., "RISKMAN™ for Windows User Manual," Version 14.2, December 2013.
2. Electric Power Research Institute, "CAFTA Fault Tree Analysis System Software Manual," Version 5.4, January 2009.