Performance of Empirical Bayes Estimation Techniques Used in Probabilistic Risk Assessment on Failure Data collected in U.S NRC Reactor Operating Experience Database

Andrei Gribok [a*], Vivek Agarwal [a], and Vaibhav Yadav [a]
[a] Idaho National Laboratory, P.O. Box 1625, MS 3818, Idaho Falls, ID 83415

Abstract: A question of numerical credibility and statistical verification of PRA results against industry-wide historical data is often raised. The historical data is our best, most objective, and reliable source of information about true probability of failure. The Office of Nuclear Regulatory Research of the U.S Nuclear Regulatory Commission (NRC) characterizes industry-average performance for components and initiating events at U.S. commercial NPPs using data available typically since 1998. Characterization of current industry-average performance is an important step in maintaining up-to-date risk models. Results (beta distributions for failure probabilities upon demand and gamma distributions for rates) are used as inputs to the U.S. NRC's Standardized Plant Analysis Risk (SPAR) models covering U.S. commercial NPPs. For this study, the U.S NRC Reactor Operating Experience Database (NROD) was investigated. To compare performance of different parameter estimation methods on real data we analysed component reliability data available from the NROD database. We used five different components and studied their failure-to-run based on a Gamma-Poisson compound distribution model. The data were analysed using nine different Empirical Bayes parametric methods as well as nonparametric estimation techniques. The goal of the analysis was to estimate the probability of failure to run for each component and each plant using available data. The obtained results demonstrate that while Empirical Bayes techniques outperform maximum likelihood estimation, improvement produced by EB methods is not uniform; while improving estimates for some plants, they also degrade estimates for others. The degradation can be particularly severe for genuinely atypical parameters when a plant has an unusually high frequency of failures for a given component. Uniform improvement cannot be achieved for EB methods.

Keywords: PRA, Bayesian estimate, parameter estimation, compound distribution model.

1. INTRODUCTION

Specification of prior distribution is one of the most important methodological as well as practical problems in Bayesian inference. Although several approaches have been proposed, none of them is completely satisfactory from both theoretical and practical points of view. The problem is twofold: how to specify our knowledge about the parameter and data in the most succinct and tractable form, and how to transfer prior knowledge of observable variables onto prior knowledge of parameters, which are generally unobservable.

Several approaches were put forward; among them the most notable are conjugate priors, Jeffreys noninformative priors, hierarchical Bayes approach, constrained noninformative priors, and empirical Bayesian methods. Conjugate priors, although widely used, can be justified only if enough information is available to believe that the true prior distribution belongs to the specified family; otherwise the main justification for using conjugate priors is their mathematical tractability. Jeffreys noninformative prior uses the Fisher information matrix to place a maximally noninformative prior on the parameters, exploiting the fact that the Fisher information matrix is widely considered to be an indicator of the accuracy of a parameter estimate. However, this approach can only be effectively used in one dimensional cases, and does not satisfy the Likelihood principle [1].

The empirical Bayesian methods use historical data to infer the prior distribution either parametrically or nonparametrically. Parametric techniques assume a functional form of the prior, usually conjugate, and infer its parameters from the available data. Nonparametric methods are free of such assumptions and can identify prior distributions of any functional form [2]. Most empirical Bayes (EB) estimations used in Probabilistic Risk Analysis (PRA) use parametric techniques; nonparametric approaches receive

significantly less attention. The hierarchical Bayes method is a fully Bayesian approach, and is currently the method of choice in Bayesian statistics. Empirical Bayes approximates hierarchical Bayes when the first level prior belongs to the same family of distributions. Otherwise, empirical Bayes may produce significantly different results. We analyze performance of a new nonparametric empirical Bayes method (NPEB) [2,3] which infers a prior distribution from a priori information that may be available from observations. The method consists of specifying or estimating a predictive distribution for the value of interest and then working backward to identify a prior distribution for the parameter. The method requires the solution of the Fredholm integral equation of the first kind [5,8], which can be effectively solved using Tikhonov regularization [4,9]. The performance of NPEB is compared with other parametric empirical Bayes techniques as well as with maximum likelihood estimations. We present results obtained on plant reliability data obtained from a Nuclear Regulatory Commission (NRC) data base.

2. GAMMA-POISSON COMPOUND DISTRIBUTION MODEL

We consider a compound distribution model defined as the data distribution marginalized over parameters of the mixing distribution. An example of a compound distribution is the Bayesian prior predictive distribution shown in Eq. 1.

$$\pi(x/\alpha) = \int_{\Theta} L(x/\theta)\, g(\theta/\alpha)\, d\theta \tag{1}$$

where $\pi(x/\alpha)$ is the prior predictive distribution of future data x, conditioned on the hyperparameter α; $L(x/\theta)$ is the likelihood of the future data x or an aleatory model conditioned on the value of the parameter of interest θ; and $g(\theta/\alpha)$ is the prior distribution of the parameter θ, conditioned on the hyperparameter α, which is sometimes called a nuisance parameter since it is of no direct interest. The prior distribution *g* is the mixing distribution. Both θ and α could be scalars or vectors. In a compound model, the value of the hyperparameter α is supposed to be known and fixed. The compound distribution model (1) has been the backbone of empirical Bayesian methods for decades. The integral (1) appears in the denominator of Bayes theorem and is known under multiple names such as evidence, the marginal data distribution, and the prior predictive distribution. In PRA applications, the compound distribution model is known as the plant population model, plant variability model or source-to-source variability model. In the case of the Gamma-Poisson compound distribution model, the prior predictive distribution is a negative binomial (NB) and can be written as:

$$\pi(x/\alpha,\beta,T) = \mathrm{NB}\left(\mathrm{x};\alpha,\frac{T}{\beta+T}\right) = \frac{\Gamma(x+\alpha)}{x!\Gamma(\alpha)} \cdot \left[\frac{T}{\beta+T}\right]^{x} \cdot \left[\frac{\beta}{\beta+T}\right]^{\alpha} = \int_0^{\infty} \underbrace{\frac{(\lambda\cdot T)^x}{x!} e^{-\lambda\cdot T}}_{Likelihood} \cdot \underbrace{\frac{\beta^{\alpha}\cdot\lambda^{\alpha-1}\cdot e^{-\lambda\cdot\beta}}{\Gamma(\alpha)}}_{Prior} d\lambda \tag{2}$$

where x is the number of failures, T is the time of operation, λ is the failure rate, α and β are parameters of Gamma distribution, and Γ is the gamma function. Graphically, the Gamma-Poisson compound distribution model is shown in Fig. 1. The failure rate λ is generated from a parametrized Gamma prior distribution, and is used in a Poisson probabilistic model along with the time of operation t to generate the number of failures x. The number of failures and operation times are the only data available for observation. Using these data, the goal is to estimate "true" failure rates for each plant.

The focus of our analysis is the prior predictive distribution (1). Under the assumption that $\pi(x/\alpha)$ and $L(x/\theta)$ are known, Eq. (1) represents a linear Fredholm integral equation of the first kind [6]. The future likelihood in this case represents the kernel and prior distribution over parameter is the sought solution. It should be stressed that a predictive distribution is a function of the observable variable x while a prior distribution is a function of the unobservable variable θ. In many practical engineering applications, the range of future observations is known from physical considerations. For example, the range of temperature, pressure, and flow rate measurements in nuclear power plants is known if plants operate under normal conditions. Also, we can use expert knowledge about the number of failures for a particular component as well as its distribution. Hence, we can place a rather informative restriction on the predictive distribution of future observations. This information can come from physical and engineering judgments as well as from historical observations of the parameter of interest. Once we deduce what predictive distribution of future observations could be we can solve the integral equation (2) to get a

prior distribution for the parameter. Doing this we effectively transform prior information about observable variables into prior information about unobservable parameters.

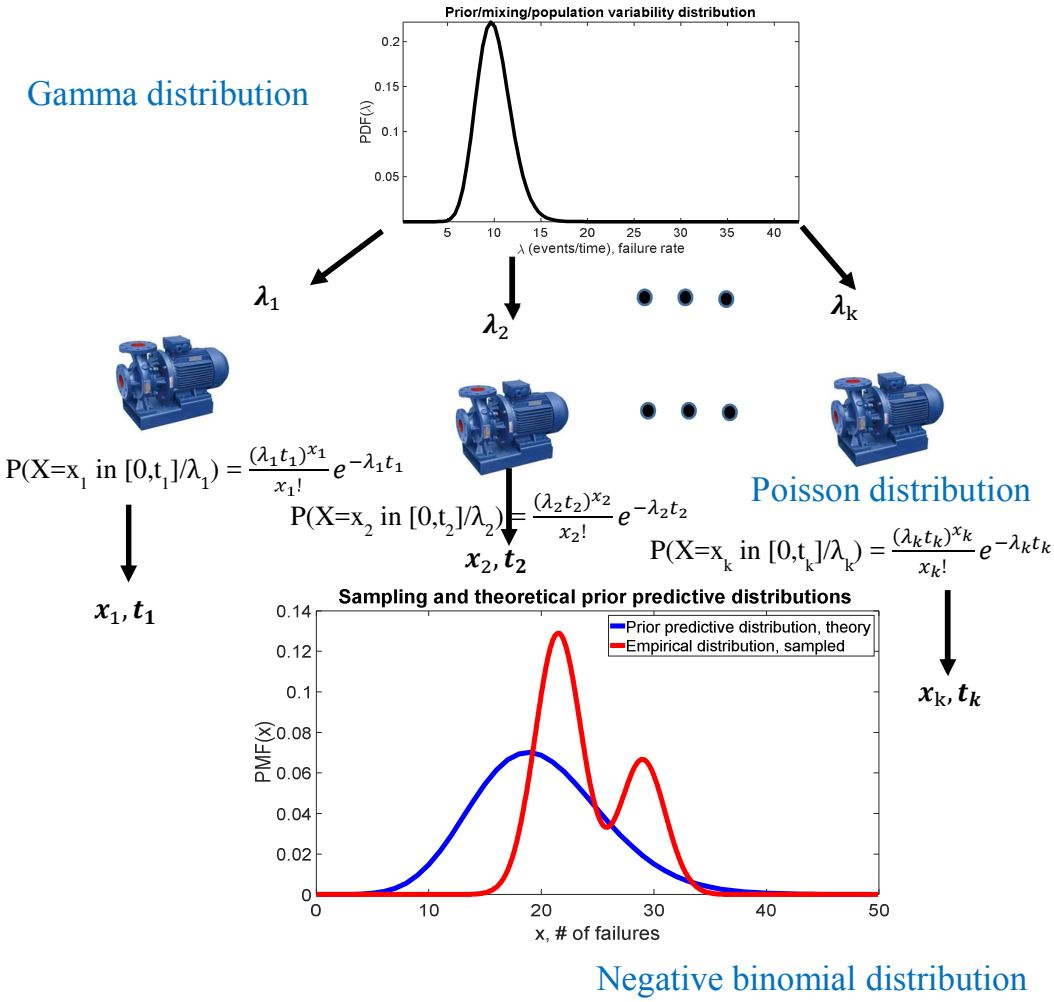

Figure 1 Gamma-Poisson compound distribution model

The solution of integral equations (1) or (2) will require the use of regularization because of the ill-posed nature of the problem. It should be pointed out that predictive distribution of future observation $\pi(x/\alpha)$ will always contain uncertainty or noise because of its empirical nature. Solving integral equation (1) by numerical methods will effectively transform ill-posedness into ill-conditioning of the matrix $L(x/\theta)$. We apply Tikhonov regularization to solve this ill-conditioned system of equations. Tikhonov regularization imposes smoothness constraints on the sought solution which is in our case a probability density function. Imposing smoothness constraints on the probability density function (pdf) is a very natural restriction because many known and practically used pdfs are smooth and differentiable.

Notice, once T is specified, the Poisson aleatory model is a function of two variables, λ and x, and can be regarded as a kernel of the integral equation (2). If the hyperparameters α and β of the Gamma distribution were known, the integral in the right-hand side of Equation 2 could be calculated to obtain the prior predictive distribution $\pi(x|\alpha,\beta)$. However, in practice, having historical data on the number of component failures, and assuming that the data generation process is homogeneous, we can replace the prior predictive distribution $\pi(x|\alpha,\beta)$ with the sampling distribution$\hat{g}(x|\alpha,\beta)$ and solve Eq. 2 for the prior distribution.

2.1. Numerical Solution of Integral Equation

Since the sampling distribution $\hat{g}(x|\alpha,\beta)$ will inevitably contain some degree of noise, the solution to Equation 2 should be done through the method of least squares. To obtain least squares solution, Equation (2) needs to be discretized using, for example, a midpoint rule. The discretization leads to the overdetermined system of linear equations:

$$\underbrace{\begin{pmatrix}\hat{g}(x_1)\\ \hat{g}(x_2)\\ \vdots\\ \hat{g}(x_m)\end{pmatrix}}_{predictive} = \underbrace{\begin{pmatrix}\frac{(\lambda_1 T)^{x_1}}{x_1!}e^{-\lambda_1 T} & \frac{(\lambda_2 T)^{x_1}}{x_1!}e^{-\lambda_2 T} & \dots & \frac{(\lambda_n T)^{x_1}}{x_1!}e^{-\lambda_n T}\\ \frac{(\lambda_1 T)^{x_2}}{x_2!}e^{-\lambda_1 T} & \frac{(\lambda_2 T)^{x_2}}{x_2!}e^{-\lambda_2 T} & \dots & \frac{(\lambda_n T)^{x_2}}{x_2!}e^{-\lambda_n T}\\ \vdots & \vdots & \dots & \vdots\\ \frac{(\lambda_1 T)^{x_m}}{x_m!}e^{-\lambda_1 T} & \frac{(\lambda_2 T)^{x_m}}{x_m!}e^{-\lambda_2 T} & \dots & \frac{(\lambda_n T)^{x_m}}{x_m!}e^{-\lambda_n T}\end{pmatrix}}_{kernel\,matrix} \cdot \underbrace{\begin{pmatrix}\tilde{f}_1\\ \tilde{f}_2\\ \vdots\\ \tilde{f}_n\end{pmatrix}}_{prior} \tag{3}$$

which can be solved via regularized least squares [4]. In compact form, the system (3) can be written as:

$$\underset{m\times 1}{\hat{g}} = \underset{m\times n}{P} \cdot \underset{n\times 1}{\tilde{f}} \tag{4}$$

$$\underbrace{argmin}_{\tilde{f}} \left\{ \left\| P\cdot \tilde{f} - \hat{g} \right\|_2^2 + \lambda^2 \cdot \left\| L\cdot \tilde{f} \right\|_2^2 \right\}, \tag{5}$$

where

$$L = \begin{pmatrix}1 & -2 & 1 & \dots & 0\\ \ddots & \ddots & \ddots & \ddots & \ddots\\ 0 & \dots & 1 & -2 & 1\end{pmatrix} \in R^{(n-2)\times n} \tag{6}$$

is the scaled approximation to the second derivative operator.

For many applications, it is computationally more convenient to convert the penalized least squares problem Eq. (5) into ordinary linear least squares using alternative formulation [8] with a nonnegativity constraint imposed since the sought solution *f* is a pdf:

$$\underbrace{argmin}_{\tilde{f}} \left\| \underbrace{\begin{pmatrix}P\\ \lambda L\end{pmatrix}}_{m+n-2\times n} \underset{n\times 1}{\cdot \tilde{f}} - \underbrace{\begin{pmatrix}\hat{g}\\ 0\end{pmatrix}}_{m+n-2\times 1} \right\|_2^2 , subject\ to\ \tilde{f} \geq 0 \tag{7}$$

The regularized solution $f_{\lambda,L}$ is subject to both nonnegativity and smoothness constraints imposed by the regularization parameter λ and the smoothing matrix *L*. Throughout this paper, we used this formulation (see Eq. 7) to obtain nonparametric empirical Bayes estimation.

Along with NPEB technique, we also applied eight other parameter estimation techniques to the same data for comparison. We investigated the following parameter estimation techniques: plant-specific maximum likelihood estimation (MLE) [10], industry-average maximum likelihood estimation [10], James-Stein Empirical Bayes Estimator (JS) [11], Jeffreys non informative prior (Jeffreys) [1], constrained Jeffreys non informative prior (CNI) [10], method of moments (Moments) [1,10], Gamma maximum likelihood (Gamma MLE), and Maximum Likelihood II (MLE2) [1,7].

3. REACTOR OPERATING EXPERIENCE COMPONENT FAILURE DATA

To compare performance of different prior selection methods on real data we analyzed component reliability data available from the NRC Reactor Operating Experience Data (NROD) database. We selected five different components which can be modelled with Poisson-Gamma compound distribution and analyzed their failure to run data. The five specific components and their failure modes that were selected are:

1. ACX_FR: Accumulator Fail to Run (ACX)
2. AHU_NR_FTR: Air Handling Unit Normally Running Fails to Run (AHU)
3. CHL_FR: Chiller Unit Fails to Run (CHL)
4. MDC_FR: Motor Driven Compressor Fail to Run (MDC)
5. MDP_FR: Motor Driven Pump Fail to Run (MDP)

The data represent individual number of component failures for the plants along with time intervals during which the failures occurred. The goal of the analysis is to estimate component failure rates for

each plant using the available data. The component data are currently available for 18 years of plants operation: 1998-2015. Since for the plant data the true prior distribution is not available we used a cross-validation approach to evaluate the performance of the different estimation techniques. The data were divided into two time periods: 1998-2000 and 2001-2015. The initial three-year period between 1998 and 2000 has been used to estimate component failure rates for each plant using different techniques, while the second much longer period between 2001 and 2015 was used as a validation or test data set to evaluate performance of the estimates obtained on the training data set. Since the validation set is significantly longer, the maximum likelihood estimation of failure rates for each plant is assumed to be closer to the true value and is used as ground truth. The overall error for N plants was estimated using the following formula:

$$\mathbf{RMSE} = \sqrt{\frac{\mathbf{1}}{\boldsymbol{N}}\sum_{\boldsymbol{i}=\mathbf{1}}^{\mathbf{N}}\left(\boldsymbol{\lambda}_{\boldsymbol{est}}^{\boldsymbol{i}} - \boldsymbol{\lambda}_{\boldsymbol{"true"}}^{\boldsymbol{i}}\right)^{\mathbf{2}}} \tag{8}$$

where $\boldsymbol{\lambda}_{\boldsymbol{est}}^{\boldsymbol{i}}$ is the plant-specific estimate of the probability of failure for i-th plant, $\boldsymbol{\lambda}_{\boldsymbol{"true"}}^{\boldsymbol{i}}$ is the "true" failure rate for i-th plant obtained from maximum likelihood estimate on the validation set spanning 2001–2015. For the empirical and hierarchical Bayesian techniques, the posterior mean was used as the probability of failure estimate. It should noted that for each component, the NROD database contains data for a different number of plants, so N depends on component type.

Since all Bayesian methods elicit prior distributions, it makes sense to first verify that the prior distribution is not degenerate (i.e., the probability of failure is not a constant and is not the same for all plants). For this data set, the null hypothesis of constant θ has been rejected based on the Pearson chi-square test with $p<<0.01$. Since in this case, the data from different plants are not "pullable," modelling population variability makes sense.

4. PERFORMACE OF DIFFERENT PARAMETER ESTIMATION TECHNIQUES

The results of applying different parameter estimation techniques to the five selected components are shown in Fig. 2.

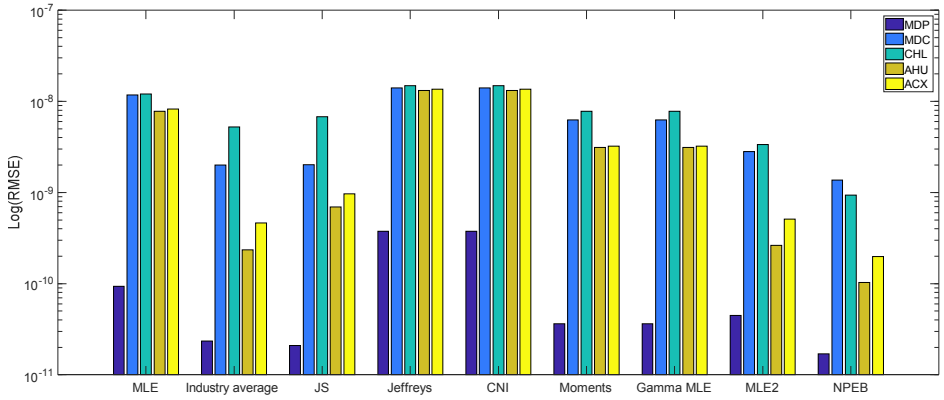

Figure 2. Average performance of different parameter estimation methods for different components on NROD data set.

As can be seen in Fig. 2, the lowest RMSEs are achieved by three methods: NPEB, the James-Stein estimator, and the industry-average MLE. Their RMSEs sometimes are smaller by a factor of about two than plant-specific MLE. The second group of methods comprised of different EB methods and plant-specific MLE. The largest RMSE was obtained using two noninformative priors applied to plant-specific data. While good performance of NPEB and JS is not surprising, the performance of industry-average MLE is unexpected particularly taking into account that this data set failed the "pullability" test. The difference between NPEB and JS is significant with $p<<0.01$ and the difference between JS and industry-average MLE is also significant with $p<<0.01$. While the difference in RMSEs between NPEB and JS seems to be small, this small difference compounded over hundreds of thousands of components in a commercial nuclear power plant may lead to significant differences in the total probability of a serious accident. In PRA, individual probabilities of failure are consolidated through

fault trees and rules of probability into a probability of a nuclear accident. Any inaccuracies in parameter estimation for individual components will be aggregated and amplified in the final result. The largest RMSE was obtained using noninformative priors, which can be explained by the fact that the priors do not provide strong enough information to improve performance in comparison with maximum likelihood estimations.

Analysing Fig. 2, a question arises—why EB methods have not displaced the principle of maximum likelihood in statistical practice? The answer can be offered by the following observation presented in Table 1:

Component	# of plants in NROD database	# plants for which NPEB has larger RMSE than plant-specific MLE
ACX	34	13
AHU	34	14
CHL	22	3
MDC	29	12
MDP	99	33

Table 1 Number of plants for which plant-specific MLE outperforms NPEB

The NPEB estimator for MD{s dominates plant-specific MLEs for 66 plants, while it degrades the performance for 33 others. A similar observation is true for all other components. There is not a component for which NPEB would improve the estimate for all plants. The EB methods concentrate attention on the total expected squared loss for the whole population, without concern for the effects on individual cases. Improvement produced by EB methods is therefore not uniform; while improving estimates for some plants, they also degrade estimates for others. The degradation can be particularly severe for genuinely atypical parameters when a plant has an unusually high frequency of failures for a given component. Uniform improvement cannot be achieved for EB methods. Since it is impossible *a priori* to know which estimate will be more accurate than MLE, we have to accept that some of our estimates will be less accurate than MLE. It is also impossible to guarantee that an EB technique is "no worse" than MLE for particular cases.

5. CONCLUSION

The performance of parameter estimation methods for PRA application on empirical data is an important technical issue, as well as a regulatory issue. For the numerical estimates produced by PRA models to be taken seriously, they should reflect the failure data observed in practice. Also, the multitude of parameter estimation methods used in PRA analysis raises the question of the "best" or at least most appropriate method for a given problem. Bayesian inference is widely used in modern PRA applications with three major approaches to prior selection tightly intertwined in practice: noninformative priors with subsequent updating, empirical Bayes, and hierarchical Bayes. Each one of these approaches depends on prior selection and elicitation methods. In this paper, we retrospectively validated different prior selection and estimation methods against nuclear industry data available in the NRC Reactor Operating Experience Database.

Our results suggest that parametric and nonparametric EB methods produce results that most accurately predict empirical component failure data available in NROD. The total risk of EB estimators depends on the actual value of the parameter; it is the smallest when the true values are clustered around zero, which is exactly the case for rate of failure of components in nuclear power plants. For our data, the noninformative single-stage priors were among the worst performers as they shift estimates towards their mean values. Our results suggest that MLE is a better choice for a single-stage inference than Bayesian inference with noninformative priors. In this study, we only focused on the accuracy of the estimates without analysis of their precision and credible intervals. This will constitute future work.

The major conclusion of this paper is that by applying EB methods to industry data we can improve our industry-wide accuracy significantly; however, the improvement for a specific plant cannot be guaranteed. Thus, if our goal is to improve accuracy of parameter estimates across all plants, then EB and hierarchical Bayes is the way to go; however, if the goal is improving the accuracy for a given plant, the researcher should look at all the options available for parameter estimation.

References

1. Christian P. Robert, The Bayesian Choice. A Decision – Theoretic Motivation. 1994 Springer-Verlag New York.
2. Bradley Efron, Two Modeling Strategies for Empirical Bayes Estimation, Statistical Science 2014, Vol. 29, No. 2, 285–301
3. Gribok, A.V., Urmanov, A.M., Hines, W.J., Uhrig, R.E "Backward Specification of Prior Distribution in Bayesian Inference as an Inverse Problem," *Inverse Problems in Science and Engineering,* Vol. 12 No. 3, June 2004, pp. 263-278.
4. A.N. Tikhonov (1963). Solution of incorrectly formulated problems and the regularization method. Doklady Akad. Nauk USSR, 151, 501–504.
5. Wahba, G. (1990). Spline models for observational data. Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania.
6. V.A. Morozov (1966). On the solution of functional equations by the method of regularization. Soviet Math. Dokl., 7, 414–417
7. A.Gelman, J.B. Carlin, H.S. Stern and D.B. Rubin (1995). Bayesian Data Analysis. Chapman & Hall, London,
8. Hansen PC. Rank-Deficient and Discrete Ill-Posed Problems Numerical Aspects of Linear Inversion. SIAM, 1998.
9. Tikhonov AN, Arsenin VY. Solution of Ill-Posed Problems, Winston & Sons, Washington, 1977.
10. Atwood, C. L., J. L. LaChance, H. F. Martz, D. J. Anderson, M. Englehardt, D. Whitehead, and T. Wheeler, *Handbook of Parameter Estimation for Probabilistic Risk Assessment*, NUREG/CR-6823, SAND2003-3348P, U.S. Nuclear Regulatory Commission (2003).
11. James, W., and C. Stein, "Estimation with Quadratic Loss," *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics, 361–379, University of California Press, Berkeley, California, 1961*.