Bayesian Confidence Intervals for Means and Variances of Lognormal and Bivariate Lognormal Distributions

J. Harvey[a,b],
&
A.J. van der Merwe[b]

[a]Centre for Statistical Consultation
Department of Statistics and Actuarial Science
University of Stellenbosch
Private Bag X1
Stellenbosch
South Africa
jharvey@sun.ac.za (corresponding author)

[b]Department of Mathematical Statistics and Actuarial Science
University of the Free State
P.O. Box 339
Bloemfontein
9300
South Africa
Matheeme.sci@ufs.ac.za

Abstract

The lognormal distribution is currently used extensively to describe the distribution of positive random variables. This is especially the case with data pertaining to occupational health and other biological data. One particular application of the data is statistical inference with regards to the mean of the data. Other authors, namely Zou, Taleban and Huo (2009), have proposed procedures involving the so-called "method of variance estimates recovery" (MOVER), while an alternative approach based on simulation is the so-called generalized confidence interval, discussed by Krishnamoorthy and Mathew (2003). In this paper we compare the performance of the MOVER-based confidence interval estimates and the generalized confidence interval procedure to coverage of credibility intervals obtained using Bayesian methodology using a variety of different prior distributions to estimate the appropriateness of each. An extensive simulation study is conducted to evaluate the coverage accuracy and interval width of the proposed methods. For the Bayesian approach both the equal-tail and highest posterior density (HPD) credibility intervals are presented. Various prior distributions (independence Jeffreys' prior, the Jeffreys-rule prior, namely, the square root of the determinant of the Fisher Information matrix, reference and probability-matching priors) are evaluated and compared to determine which give the best coverage with the most efficient interval width. The simulation studies show that the constructed Bayesian confidence intervals have satisfying coverage probabilities and in some cases outperform the MOVER and generalized confidence interval results. The Bayesian inference procedures (hypothesis tests and confidence intervals) are also extended to the difference between two lognormal means as well as to the case of zero-valued observations and confidence intervals for the lognormal variance. In the last section of this paper the bivariate lognormal distribution is discussed and Bayesian confidence intervals are obtained for the difference between two correlated lognormal means as well as for the ratio of lognormal variances, using nine different priors.

Keywords: Bayesian procedure; Lognormal; Highest Posterior Density; MOVER; Credibility intervals; Coverage probabilities; Zero-valued observations; Bivariate Lognormal; Lognormal variance

1 Introduction

Lognormally distributed data presents itself in a number of scientific fields. According to Limpert et al. (2001), the distribution may be used to approximate right skewed data that arises in a wide variety of scientific settings. Particularly in the area of health costs the log distribution has been extensively used by other authors and numerous statistical methods have been developed. The literature dealing with the analysis of the means of lognormal data has therefore increased substantially. This includes procedures for a single sample case, a difference between two sample means and additional zero values for each of the cases. Recently researchers have also looked at confidence intervals for the lognormal variance and for the ratio of two lognormal variances.

Many authors have proposed methods based on the simulation of pivotal statistics, commonly referred to as generalized confidence intervals. These methods have generated a series of papers on the means of lognormal data (see for example Krishnamoorthy and Mathew (2003), Tian (2005), Chen and Zhou (2006), Krishnamoorthy, Mathew and Romachandren (2006), Tian and Wu (2007 [a],[b]) and Bebu and Mathew (2008)). For further details refer to Zou, Taliban and Huo (2009). The simulation of pivotal statistics is a frequentist method that is in effect rather similar to some of the Bayesian procedures that will be proposed in the next sections.

Instead of adapting a simulation approach for making inferences on the lognormal mean, Zou et al (2009) proposed procedures involving the so-called “method of variance estimates recovery” (MOVER). The MOVER method was designed in order to apply to a general scenario and also to provide adequate coverage rates in estimation procedures relating to lognormally distributed data. The advantage of the MOVER is therefore that it is easily applicable to many different settings with little more than a basic knowledge of introductory statistical text.

In this paper an extensive simulation study will be conducted to ascertain the coverage accuracy of the proposed methods on the mean and variances of lognormally distributed data as well as for the bivariate lognormal distribution. These proposed methods are all based on a Bayesian framework, where the choice of prior distribution is the factor of interest. Specifically, the choice of different prior distributions in different parameters settings and the appropriateness of each is of primary importance. This will be compared to the MOVER algorithm developed by Zou et al (2009) and the generalized confidence interval procedure proposed by Krishnamoorthy and Mathew (2003) by means of simulation studies. Another setting that is also addressed in this paper is that of zero values, particularly zero costs in medical health costs. The focus was also on inferences about the overall population means, including zero costs. In the last section of this paper the bivariate lognormal distribution will be discussed and Bayesian confidence intervals will be obtained for the difference between two correlated lognormal means as well as for the ratio of two lognormal variances using nine different priors.

As mentioned a Bayesian approach to the problem will be taken. Depending on the choice of prior distribution it will be shown that the Bayesian procedure has equal or better coverage accuracy than both the MOVER method and generalized confidence interval procedure. In the next section we begin with a formulation of the model and a specification of all parameters and distributions of interest. In further sections we compare the performance of the Bayesian methods for different prior distributions by conducting a simulation study to assess specific quantities of the proposed credibility intervals in pre-defined finite sample sizes.

2 Bayesian Confidence Intervals for the Mean of Lognormally Distributed Data

Suppose that random variable X follows a lognormal distribution so that $\mathrm{Y} = \ln(X) \sim N(\mu, \sigma^2)$. The mean of X is defined as

$$M = E(X) = exp\left(\mu + \frac{1}{2}\sigma^2\right) \tag{1}$$

Let $X_1, X_2, \dots, X_n$ be a random sample from the lognormal distribution, then for the likelihood function the Fisher Information matrix follows as

$$I(\mu, \sigma^2) = diag\begin{bmatrix}\frac{n}{\sigma^2} & \frac{n}{2\sigma^4}\end{bmatrix} \tag{2}$$

As mentioned, one of the objectives of our research is compare the Bayesian procedure for different choices of prior distributions for $\boldsymbol{\theta} = (\mu, \sigma^2)$, the unknown parameters. In the following sections different choices of prior distributions will be discussed in an effort to eventually compare the results.

2.1 Independence Jeffreys Prior:

Since $\boldsymbol{\theta} = (\mu, \sigma^2)$ is unknown the prior

$$p(\boldsymbol{\theta}) \propto \frac{1}{\sigma^2} \tag{3}$$

will be specified for the unknown parameters. This is known as the independence Jeffreys prior. Equation (3) follows from the Fisher Information matrix by assuming μ and $\sigma^2$to be independently distributed, *a priori*, with μ and $log\sigma^2$ each uniformly distributed. Combining the likelihood function and the prior density function (3) the joint posterior density function can be written as:

$$P(\boldsymbol{\theta}|data) = \left(\frac{2\pi\sigma^2}{n}\right)^{-1/2} exp\left(-\frac{n}{2\sigma^2}(\mu - \hat{\mu})^2\right)\left(\frac{\nu\hat{\sigma}^2}{2}\right)^{1/2\nu}\left(\frac{(\sigma^2)^{-1/2(\nu+2)} exp\left[-\frac{\nu\hat{\sigma}^2}{2\sigma^2}\right]}{\Gamma\left(\frac{\nu}{2}\right)}\right) \tag{4}$$

where $\hat{\mu} = \frac{1}{n}\sum_1^n y_i$, $\nu = n - 1$, $y_i = ln(x_i)$

$$\hat{\sigma}^2 = \frac{1}{n-1}\sum_1^n (y_i - \hat{\mu})^2.$$

From (4) it follows that the posterior distribution can be defined as the conditional posterior distribution of μ, which is normal:

$$\mu|\sigma^2, data \sim N\left(\hat{\mu}, \frac{\sigma^2}{n}\right) \tag{5}$$

and for σ^2, the posterior density function is an Inverted Gamma density, specifically:

$$P(\sigma^2|data) = \left(\frac{\nu\hat{\sigma}^2}{2}\right)^{1/2\nu}\left(\frac{(\sigma^2)^{-1/2(\nu+2)}exp\left[-\frac{\nu\hat{\sigma}^2}{2\sigma^2}\right]}{\Gamma\left(\frac{\nu}{2}\right)}\right). \quad (6)$$

From equation (6) it follows that $\tau^* = \frac{\nu \ \hat{\sigma}^2}{\sigma^2}$ has a chi-square distribution with ν degrees of freedom. From classical statistics (if $\hat{\sigma}^2$ is considered to be random) it is well known that τ^* is also distributed chi-square with ν degrees of freedom. This agreement between classical and Bayesian statistics is only true if the prior $p(\sigma^2) \propto \sigma^{-2}$ is used. If some other prior distributions are used, for example $p(\sigma^2) \propto \sigma^{-3}$ or $p(\sigma^2) \propto constant$, then the posterior of τ^* will still be a chi-square distribution but the degrees of freedom will be different.

Since $lnM = \mu + \frac{1}{2}\sigma^2$ standard routines can be used in the simulation procedure.

2.2 Simulation Procedure:

The following simulation was obtained from the preceding theory using the MATLAB® package:

1. For given values of μ, σ^2 and n, sample values $\hat{\mu}$ and $\hat{\sigma}^2$ were drawn from $\hat{\mu} \sim N\left(\mu, \frac{\sigma^2}{n}\right)$ and $\hat{\sigma}^2 \sim \frac{u\sigma^2}{\nu}$, where $u \sim \chi^2_\nu$. This was repeated 10000 times. It is only necessary to simulate the sufficient statistics $\hat{\mu}$ and $\hat{\sigma}^2$ for the sample data and not the complete sample.
2. For every pair of $\hat{\mu}$ and $\hat{\sigma}^2$ values, μ and σ^2 can be simulated from their respective posterior distributions, as given in (5) and (6).
3. For the 10000 generated values, $lnM = \mu + \frac{1}{2}\sigma^2$ can be calculated and ordered. The equal-tailed and shortest 95% Highest Posterior Density Intervals can then be obtained.
4. The coverage and average widths of the intervals from the 10000 samples can be calculated.

2.3 Jeffreys-Rule Prior:

The Jeffreys'-Rule prior is defined as

$$p(\boldsymbol{\theta}) \propto \sigma^{-3} \quad (7)$$

This is derived from the square root of the determinant of the Fisher Information matrix. The posterior distribution of σ^2 is as defined in (6), but with $\nu + 1$ instead of ν and the variance component is therefore simulated as $\sigma^2 = \frac{\nu\hat{\sigma}^2}{\chi^2_{\nu+1}}$.

2.4 Reference and Probability Matching Priors

In addition to the two previously mentioned Jeffreys' prior distributions the following prior distributions were all tested:

$$p(\boldsymbol{\theta}) \propto \frac{1}{\sigma}\sqrt{1 + \frac{2}{\sigma^2}} \quad (8)$$

$$p(\boldsymbol{\theta}) \propto \frac{1}{\sigma^2}\sqrt{1 + \frac{2}{\sigma^2}} \quad (9)$$

Prior distributions (8) and (9) are the reference and probability-matching priors and the derivations of these priors are given in Harvey, Groenewald and van der Merwe (2010).

As in the case of the independence Jeffreys' prior and the Jeffreys-rule prior, 10000 samples were simulated and for each sample 10000 Bayesian simulations were made to obtain credibility intervals.

The reference prior method is also derived from the Fisher information matrix. Reference priors depend on the specification of a parameter of interest, specification of the grouping of parameters and ordering of the groupings. According to Berger and Bernardo (1992) the reference prior theory has been the most successful technique for deriving objective priors.

Datta and Ghosh (1995) on the other hand derived the differential equation which a prior must satisfy so that its posterior distribution will have good frequentist properties. Probability matching priors are especially suitable for constructing accurate one-sided credibility intervals.

2.5 A Comparison of the Bayesian Methods with the Generalized Confidence Interval Procedure and the MOVER in obtaining Confidence Intervals for $\boldsymbol{\mu+\frac{1}{2}\sigma^2}$

The object of the paper by Zou et al (2009, page 3760) was to demonstrate the MOVER in different scenarios, i.e. for a few different combinations of n and σ^2, where $\mu = -\frac{1}{2}\sigma^2$. The same scenarios will be presented here for the Bayesian confidence intervals and these will be compared to the results from the MOVER method and generalized confidence interval procedure to evaluate the performance of the Bayesian confidence intervals for both the equal-tailed and HPD intervals. The following characteristics are reported:

- coverage probabilities
- average interval lengths

A nominal significance level of $\alpha = 0.05$ will be used for each parameter setting.

The confidence limits for the MOVER, as given by Zou et al (2009) on page 3758, are:

$$l = \hat{\mu} + \frac{\hat{\sigma}^2}{2} - \sqrt{Z^2_{\alpha/2}\frac{\hat{\sigma}^2}{n} \left\{\frac{\hat{\sigma}^2}{2}\left(1 - \frac{\nu}{\chi^2_{1-\alpha/2,\nu}}\right)\right\}^2} \tag{10}$$

$$u = \hat{\mu} + \frac{\hat{\sigma}^2}{2} + \sqrt{Z^2_{\alpha/2}\frac{\hat{\sigma}^2}{n} \left\{\frac{\hat{\sigma}^2}{2}\left(\frac{\nu}{\chi^2_{\alpha/2,\nu}} - 1\right)\right\}^2} \tag{11}$$

where $Z_{\alpha/2}$ is the upper $\alpha/2$ quantile of the standard normal distribution and $\chi^2_{\alpha/2,\nu}$ is the $\alpha/2$ percentile of the chi-squared distribution with ν degrees of freedom.

The generalized confidence interval procedure discussed by Krishnamoorthy and Mathew (2003) involves simulating the pivotal statistic:

$$lnM^* = \hat{\mu} - \frac{Z^*}{\sqrt{\frac{\tau^*}{\nu}}}\frac{\hat{\sigma}}{n} + \frac{\hat{\sigma}^2}{\frac{2\tau^*}{\nu}}$$

where $Z^* \sim N(0,1)$ and $\tau^* \sim \chi^2_\nu$ denote simulated values. This method is exactly the same as the Bayesian simulation procedures in the case of the independence Jeffreys'prior.

2.6 Discussion of Results for the Simulation Studies

The objective of this section is to compare these results against results obtained from a Bayesian-based simulation study using a specifically chosen set of prior distributions and to evaluate the performance of each prior distribution against both the other distributions and the results obtained by Zou et al (2009), overall. The generalized confidence intervals will also be discussed.

The following table presents the summary statistics of the results in both the Zou et al (2009) simulation study and the Bayesian simulation study using Jeffreys' Independence prior.

The same designs are used as considered by Zou et al (2009), where (< , >)% refers to the proportion of cases where the interval is below or above the true value respectively:

Table 1
Comparison of the MOVER and Independence Jeffreys' Prior for Constructing Two-sided 95% Confidence Intervals for $\mu + \frac{1}{2}\sigma^2$

		MOVER		GCI		Jeffreys' (Equal-tail)		Jeffreys' (HPD)	
n	σ^2	Cover (< , >)%	Width	Cover (< , >)%	Width	Cover (< , >)%	Width	Cover (< , >)%	Width
5	0.5	93.47 (3.19, 3.34)	2.54	93.99 (2.00, 4.01)	2.69	94.31 (1.87, 3.82)	2.68	95.67 (3.01, 1.32)	2.27
	2.0	95.10 (2.99, 1.91)	8.76	93.77 (2.36, 3.87)	8.82	94.13 (2.11, 3.76)	8.68	95.70 (3.66, 0.64)	6.64
	3.0	95.35 (2.60, 2.05)	12.72	93.90 (2.09, 4.01)	12.74	94.35 (2.07, 3.58)	12.71	95.81 (3.68, 0.51)	9.50
20	0.5	94.24 (3.37, 2.39)	0.74	94.56 (2.35, 3.09)	0.76	94.56 (2.29, 3.15)	0.77	95.06 (3.00, 1.94)	0.75
	2.0	94.94 (2.89, 2.17)	2.04	94.39 (2.39, 3.22)	2.06	95.16 (2.30, 2.54)	2.05	95.77 (3.30, 0.93)	1.96
	3.0	95.41 (2.82, 1.77)	2.86	94.97 (2.37, 2.66)	2.87	95.59 (2.36, 3.05)	2.86	95.06 (3.62, 1.32)	2.71

The equal-tailed results from the Jeffreys' Independence prior should match the generalized confidence interval (GCI) results as found in Zou et al (2009). This is due to the formulation of the model as described in previous sections. From the above simulations it is apparent that the equal-tail intervals above do indeed match the GCI results, for all practical purposes.

In comparison to the MOVER confidence intervals, the equal-tail intervals also seem to compare reasonably well, with insignificant differences in both the proportion of confidence intervals above and below the true parameter, but the width of the intervals are larger than those of the MOVER. Naturally, as the sample size increases the width of the interval tends to decrease. On the other hand, if the variance increases the width of the intervals will also increase.

Thus, the equal-tail intervals do not offer an improvement on the MOVER method. However, when considering the HPD intervals, which are only possible through the Bayesian framework in this setting, a large improvement on the MOVER can be gained, particularly when n is small. Of particular interest to note is that these HPD intervals result in considerable reductions in interval width. Also of note is that the proportion of intervals above the true parameter is considerably less than both the MOVER and the equal-tail intervals.

So, the performance of the Jeffreys' Independence prior is comparable (or improved for HPD intervals) to the MOVER. However, in terms of the literature, Box and Tiao (1973), this would be the natural choice of prior distribution in this setting and thus, its accuracy is an expected result. As mentioned previously though, other prior distributions were also chosen to examine their effectiveness in this situation. These other prior distributions were mentioned in sections 2.3 and 2.4. The table below represents the results of these distributions, once again compared to the MOVER. However, these were only performed for the extreme values of σ^2 in Table 1. Also, only the coverage and interval widths are presented.

Table 2
Comparison of the MOVER and Other Prior Distributions for Constructing Two-sided 95% Confidence Intervals for $\mu + \frac{1}{2}\sigma^2$

			Equal-Tail / MOVER		HPD Intervals	
n	σ^2	Prior / Method	Cover %	Width	Cover %	Width
20	0.5	MOVER	94.24	0.74	N/A	N/A
		$p(\boldsymbol{\theta}) \propto \sigma^{-3}$	94.39	0.74	94.14	0.72
		Reference Prior	94.88	0.77	95.26	0.76
		Probability-Matching Prior	94.74	0.74	94.81	0.73
20	3	MOVER	95.41	2.86	N/A	N/A
		$p(\boldsymbol{\theta}) \propto \sigma^{-3}$	94.59	2.66	94.21	2.54
		Reference Prior	94.71	2.99	95.66	2.84
		Probability-Matching Prior	94.78	2.78	94.95	2.65

It appears as though the coverage of the other prior is not as good as the Independence Jeffreys' prior, particularly for small sample sizes. However, as the sample size increases the effect of the prior distribution seems to decrease and the results are comparable.

From Table 2 it is also clear that the reference prior seems to have better coverage than the probability-matching prior. It must however, be remembered that the probability-matching prior is derived for one-sided credibility intervals. This might be the reason for the undercoverage if n is small.

2.7 Comparing the Means of Two Lognormal Distributions

The ratio of means from two different lognormal populations can also be examined. Due to the problem specification the ratio between these two means can be written as $lnM_1 - lnM_2$ where $lnM_j = \mu_j + \frac{1}{2}\sigma_j^2, (j = 1,2)$. Thus it is an easy matter to extend the simulation study applied earlier to this difference. Using the simulation methods mentioned before, first one population mean is simulated, then the next and finally they are subtracted from each other. The credibility intervals can be calculated from this "differenced" data.

In the case of the MOVER the $(1-\alpha)100\%$ confidence limits for $lnM_1 - lnM_2$ are given by

$$L = \hat{\theta}_1 - \hat{\theta}_2 - \sqrt{\left(\hat{\theta}_1 - l_1\right)^2 + \left(u_2 - \hat{\theta}_2\right)^2}$$

$$U = \hat{\theta}_1 - \hat{\theta}_2 + \sqrt{\left(\hat{\theta}_2 - l_2\right)^2 + \left(u_1 - \hat{\theta}_1\right)^2}$$

where

$$\hat{\theta}_i = \hat{\mu}_i + \frac{1}{2}\hat{\sigma}_i^2$$

and l_i and u_i $(i = 1,2)$ are defined in equations (10) and (11).

<u>*Example 1:*</u> *Two-sided 95% Bayesian Confidence Intervals for Lognormal Means and the Difference between Two Lognormal Means*

On page 3761 of Zou et al (2009) the following example was given (as referenced in Zhou et al (1997)): the effect of race on the cost of medical care for type I diabetes was investigated using MOVER. Log transformed cost data for 119 black patients and 106 white patients. For the black patients the following log-transformed data was available: $\hat{\mu}_1 = \bar{y}_1 = 9.06694$ and $\hat{\sigma}_1^2 = s_1^2 = 1.82426$. For the white patients this was $\hat{\mu}_2 = \bar{y}_2 = 8.69306$ and $\hat{\sigma}_2^2 = s_2^2 = 2.69186$. For the MOVER the following results were obtained and this is compared to results obtained from the Bayes methods:

	MOVER	Jeffreys-Rule	Independence Jeffreys	Reference Prior	Probability Matching Prior
<u>Black Patients</u>					
Lower Limit (equal-tail)	15806.00	15970.19	15882.58	16021.01	15906.92
Upper Limit (equal-tail)	31388.77	31759.23	31752.94	32025.42	31557.68
Lower Limit (HPD)		15708.48	15626.14	15753.04	15671.29
Upper Limit (HPD)		31063.21	31121.86	31351.21	30982.03
<u>White Patients</u>					
Lower Limit (equal-tail)	14842.03	15097.15	15039.52	15081.46	14962.31
Upper Limit (equal-tail)	39722.09	40079.56	40156.06	40483.85	40019.03
Lower Limit (HPD)		14744.94	14668.91	14765.88	14461.04
Upper Limit (HPD)		38799.26	38847.41	39404.13	38289.58
<u>Difference</u>					
Lower Limit (equal-tail)	-19112.18	-19156.98	-19241.18	-19812.57	-19169.17
Upper Limit (equal-tail)	11371.14	11229.07	11381.24	11457.97	11361.90
Lower Limit (HPD)		-19161.99	-19254.62	-19820.21	-19170.15
Upper Limit (HPD)		11225.65	11367.43	11446.41	11360.83

Since the sample sizes are large the interval lengths for the different procedures are more or less the same. It is however, clear that the HPD intervals for the Independence Jeffreys' prior are somewhat shorter for black and white patients than those of the MOVER.

3 The Case of Zero-Valued Observations

3.1 Model Formulation

A further problem suggested by Zou et al (2009) and other authors (e.g. Zhou and Tu (2000)) is that of the inclusion of zero-valued observations. This setting we assume that the non-zero-valued observations are lognormally distributed, but zero valued observations are also present in the data.

We assume that the probability of obtaining a zero observation from the population is δ where $0 < \delta < 1$. This can be specified in the likelihood function by means of a binomial distribution. Furthermore, we assume that the non-zero observations are distributed lognormally in the same way as described in section 2. From this preliminary setting specification we wish to construct credibility intervals for $\widetilde{M}$, where:

$$\widetilde{M} = (1 - \delta)exp\left(\mu + \frac{1}{2}\sigma^2\right).$$

The only adjustment to the setting mentioned in section 2 is that $\hat{\delta} = \frac{n_0}{n}$, where n_0 is the number of zero-valued observations, n_1 is the number of non-zero observations and $n = n_0 + n_1$.

In the first part of this section Bayesian confidence intervals will be simulated for $ln\widetilde{M}$ using different priors. These intervals will be compared to the MOVER and generalized confidence interval procedures.

Another advantage of the Bayesian procedure over the MOVER and generalized confidence interval methods is that it can also easily be used to obtain confidence intervals for

$$\tilde{\sigma}^2 = (1 - \delta)exp(2\mu + \sigma^2)[exp\sigma^2 - (1 - \delta)]$$

the variance of the lognormal distribution in the case of zero-valued observations.

In section 3.4 a simulation study is therefore conducted to evaluate the coverage accuracy and interval width of the Bayesian confidence intervals for $\tilde{\sigma}^2$.

3.2 Intervals Based on a Bayesian Procedure

In the first part of this section Bayesian confidence intervals will be simulated for $ln\widetilde{M}$ using different priors. These intervals will be compared to the MOVER and generalized confidence interval procedures. The unknown parameters can now be defined by: $\widetilde{\boldsymbol{\theta}}' = [\delta \quad \mu \quad \sigma^2]$ and the likelihood function is given by:

$$L(\tilde{\boldsymbol{\theta}}|data) \propto \delta^{n_0}(1-\delta)^{n_1} \prod_{i=1}^{n_1} \left(\frac{1}{\sigma^2}\right)^{1/2} exp\left[-\frac{(y_i-\mu)^2}{2\sigma^2}\right] \tag{12}$$

where X_i $(i = 1,2,\dots,n_1)$ is a random sample from a lognormal distribution and $Y_i = ln(X_i) \sim N(\mu,\sigma^2)$.

From equation (12) the Fisher Information matrix is obtained as:

$$I(\tilde{\boldsymbol{\theta}}) = diag\left[\frac{n}{\delta\ \ (1-\delta\ \ )}\quad \frac{n\ \ (1-\delta\ \ )}{\sigma^2}\quad \frac{n\ \ (1-\delta\ \ )}{2\sigma^4}\right] \tag{13}$$

The following prior distributions were then chosen:

Independence Jeffreys' Prior: $p(\tilde{\boldsymbol{\theta}}) \propto \sigma^{-2}\delta^{-1/2}(1-\delta)^{-1/2}$ (14)

Jeffreys' Rule Prior: $p(\tilde{\boldsymbol{\theta}}) \propto \sigma^{-3}\delta^{-1/2}(1-\delta)^{1/2}$ (15)

which is obtained from $|I(\tilde{\boldsymbol{\theta}})|^{1/2}$. The following: $p(\delta) \propto \delta^{-1/2}(1-\delta)^{-1/2}$ is the prior proposed by Jeffreys (1967) for the binomial parameter.

In addition to these prior distributions, the Reference and Probability-Matching prior will also be evaluated.

It is easy to show that the posterior distribution of δ in the case of Independence Jeffreys' prior is a Beta distribution, specifically $B\left(n_0+\frac{1}{2};\ n_1+\frac{1}{2}\right)$, while the posterior distribution of δ for the Jeffreys' Rule prior is $B\left(n_0+\frac{1}{2};\ n_1+\frac{3}{2}\right)$. The posteriors $p(\delta|data)$ and $p(\mu\ ,\sigma^2|data)$ will be independently distributed.

The simulation procedure is similar to that previously mentioned except for the simulation from the Beta distribution to obtain δ.

An accurate confidence interval for δ (Zhou et al (2009)) is given by

$$\left[\hat{\delta} + \left(Z^2_{\alpha/2}\middle/2n\right) \pm \sqrt{\frac{\hat{\delta}(1-\hat{\delta})+Z^2_{\alpha/2}/4n}{n}}\middle/\left(1+Z^2_{\alpha/2}/n\right)\right] \tag{16}$$

By combining (16) with (10) and (11) the MOVER can be used to obtain confidence intervals for the case of zero-valued observations.

3.3 Discussion of Results from Simulation Study

Table 3

Comparison of the MOVER and GCI against Independence Jeffreys' Prior for Zero Values Included for Constructing Two-sided 95% Confidence Intervals for $ln(1-\delta)+\left(\mu+\frac{1}{2}\sigma^2\right)$

			MOVER		GCI		Equal-Tail		HPD Intervals	
δ	n	σ^2	Cover %	Width	Cover %	Width	Cover %	Width	Cover %	Width
0.1	15	1	95.03 (3.60, 1.37)	1.65	95.53 (2.34, 2.13)	1.72	95.03 (2.59, 2.38)	1.66	95.36 (1.26, 3.38)	1.59
		2	95.22 (3.13, 1.65)	2.78	95.50 (2.26, 2.24)	2.85	95.14 (2.62, 2.24)	2.79	95.56 (0.91, 3.53)	2.62
		3	94.87 (2.90, 2.23)	3.88	94.94 (2.35, 2.71)	3.94	94.86 (2.73, 2.41)	3.89	95.42 (0.80, 3.78)	3.59
	50	1	95.10 (3.02, 1.88)	0.78	95.79 (2.26, 1.95)	0.80	94.93 (2.47, 2.60)	0.78	94.98 (1.78, 3.24)	0.77
		2	95.16 (2.87, 1.97)	1.26	95.41 (2.37, 2.22)	1.29	95.28 (2.47, 2.25)	1.27	95.56 (1.47, 2.97)	1.24
		3	94.86 (2.69, 2.45)	1.73	94.87 (2.43, 2.70)	1.76	94.93 (2.68, 2.39)	1.73	95.13 (1.53, 3.34)	1.69
0.2	15	1	95.17 (3.30, 1.53)	1.87	95.99 (2.14, 1.87)	1.98	95.48 (2.44, 2.08)	1.89	96.18 (1.08, 2.74)	1.80
		2	95.41 (2.78, 1.81)	3.13	95.70 (2.02, 2.28)	3.23	94.98 (2.76, 2.26)	3.13	95.78 (0.88, 3.34)	2.91
		3	94.97 (3.11, 1.92)	4.38	94.93 (2.54, 2.53)	4.47	94.60 (2.83, 2.57)	4.37	95.10 (0.81, 4.09)	3.99
	50	1	95.00 (2.99, 2.01)	0.85	95.56 (2.26, 2.18)	0.88	94.90 (2.62, 2.48)	0.86	95.20 (1.84, 2.96)	0.85
		2	94.95 (2.92, 2.13)	1.37	95.28 (2.36, 2.36)	1.39	95.09 (2.74, 2.17)	1.37	95.20 (1.72, 3.08)	1.34
		3	95.30 (2.58, 2.12)	1.87	95.39 (2.24, 2.37)	1.90	95.13 (2.77, 2.10)	1.87	95.46 (1.45, 3.09)	1.83

From the above it is evident that the interval lengths and coverage of the equal-tail intervals are very similar to those of the MOVER, with the lengths being almost identical. It is interesting to note however, that the proportion of intervals above and below the true value differ substantially. The Bayesian HPD intervals are therefore a large improvement on the MOVER and generalized confidence intervals.

In Table 1 the equal tail Bayesian intervals using the Jeffreys' prior $p(\mu,\sigma^2)\propto\sigma^{-2}$ are identical to the generalized confidence intervals. This will not be the case in Table 3. The reason for this is the simulation of δ. In the Bayesian case (using the Independence Jeffreys prior) the posterior distribution of δ is the Beta, $B\left(n_0+\frac{1}{2},n_1+\frac{1}{2}\right)$, distribution while Zou et al (2009) (see also Tian (2005)) used two pivotal quantities, $B(n_0+1,n_1)$ and $B(n_0,n_1+1)$ for δ, which are combined with the pivotal quantity of the lognormal mean to simulate $ln\widetilde{M}=ln(1-\delta)+\left(\mu+\frac{1}{2}\sigma^2\right)$. From Table 3 it is also clear that the Bayesian equal-tail intervals are shorter than those of the generalized confidence interval procedure with just as good or better coverage probabilities.

Tian and Wu (2006) also considered an approach based on the adjusted log-likelihood ratio statistics for constructing a confidence interval for the mean of lognormally distributed data with excess zeros. Because of different parameter values only a few results could be compared. It does seem, however, that the procedures described in Table 3 result in better results than the adjusted log likelihood method.

Once again, other prior distributions were also evaluated, but only for the case of $\delta=0.1$. In addition, the proportion of intervals above and below the true parameter values was also not recorded.

Table 4
Comparison of the MOVER and Other Prior Distributions for Constructing Two-sided 95% Confidence Intervals for $ln(1-\delta)+\left(\mu+\frac{1}{2}\sigma^2\right)$

			Equal-Tail / MOVER		HPD Intervals	
n	σ^2	Prior / Method	Cover %	Width	Cover %	Width
15	1	MOVER	95.03	1.65		
		$p(\tilde{\boldsymbol{\theta}}) \propto \sigma^{-3}\delta^{-1/2}(1-\delta)^{1/2}$	94.35	1.53	94.10	1.48
		Reference Prior	95.45	1.75	96.35	1.67
		Probability-Matching Prior	96.25	1.67	94.41	1.59
	2	MOVER	95.22	2.78		
		$p(\tilde{\boldsymbol{\theta}}) \propto \sigma^{-3}\delta^{-1/2}(1-\delta)^{1/2}$	94.06	2.52	93.93	2.39
		Reference Prior	95.01	3.02	96.10	2.81
		Probability-Matching Prior	94.76	2.68	94.79	1.59
	3	MOVER	94.87	3.88		
		$p(\tilde{\boldsymbol{\theta}}) \propto \sigma^{-3}\delta^{-1/2}(1-\delta)^{1/2}$	94.47	3.45	94.02	3.23
		Reference Prior	95.26	4.27	96.48	3.90
		Probability-Matching Prior	94.22	3.76	94.45	3.49
50	1	MOVER	95.10	0.78		
		$p(\tilde{\boldsymbol{\theta}}) \propto \sigma^{-3}\delta^{-1/2}(1-\delta)^{1/2}$	94.79	0.76	94.72	0.76
		Reference Prior	94.93	0.79	94.95	0.79
		Probability-Matching Prior	94.80	0.77	94.85	0.76
	2	MOVER	95.16	1.26		
		$p(\tilde{\boldsymbol{\theta}}) \propto \sigma^{-3}\delta^{-1/2}(1-\delta)^{1/2}$	94.37	1.23	94.25	1.21
		Reference Prior	95.34	1.28	95.78	1.26
		Probability-Matching Prior	94.79	1.25	95.00	1.23
	3	MOVER	94.86	1.73		
		$p(\tilde{\boldsymbol{\theta}}) \propto \sigma^{-3}\delta^{-1/2}(1-\delta)^{1/2}$	95.03	1.68	94.51	1.65
		Reference Prior	95.11	1.77	95.53	1.72
		Probability-Matching Prior	94.65	1.72	94.66	1.68

In this instance it appears as though the probability matching prior, $p(\tilde{\boldsymbol{\theta}}) \propto \delta^{-1/2}(1-\delta)^{-1/2}\sigma^{-2}\left(1+\frac{2}{\sigma^2}\right)^{1/2}$, and the Jeffreys' Rule prior ($p(\tilde{\boldsymbol{\theta}}) \propto \sigma^{-3}\delta^{-1/2}(1-\delta)^{1/2}$) tend to give coverage probabilities slightly less than 0.95. The reference prior, $p(\tilde{\boldsymbol{\theta}}) \propto \delta^{-1/2}(1-\delta)^{-1/2}\sigma^{-1}\left(1+\frac{2}{\sigma^2}\right)^{1/2}$ on the other hand gives the correct coverage probabilities, but the intervals are wider than that of the MOVER.

Example 2: *Two-sided 95% Confidence Intervals for* $(1-\delta)exp\left(\mu+\frac{1}{2}\sigma^2\right)$

On page 3761 of Zou et al (2009) the following example was given (as referenced in Zhou and Tu (2000)): diagnostic test charges on 40 patients were investigated. Among them, 10 patients had no diagnostic test charges and the charges for the remaining patients were approximated using a lognormal distribution. On the log scale the following values were observed: $\bar{y}=6.8535$ and $s^2=1.8696$.

For the MOVER the following interval was obtained and this is compared to results obtained from the Bayes and GCI methods:

	MOVER	GCI	Jeffreys-Rule	Independence Jeffreys	Reference Prior	Probability Matching Prior
Lower Limit (equal-tail)	955.50	970.81	1002.18	975.03	996.10	982.41
Upper Limit (equal-tail)	4491.55	4687.37	4690.34	4310.66	4802.44	4519.89
Lower Limit (HPD)			958.05	906.87	931.60	926.10
Upper Limit (HPD)			4345.73	3932.04	4356.93	4111.21

From the table it is clear that the intervals do not differ much. The intervals for the MOVER and the probability-matching prior are for all practical purposes the same. The shortest intervals are obtained from the Independence Jeffreys' and Probability-Matching prior. As mentioned before, the equal-tail Independence Jeffreys interval and the GCI interval will not be the same because of the difference in the simulation of δ.

3.4 Bayesian Confidence Intervals for $\tilde{\sigma}^2 = (1-\delta)exp(2\mu + \sigma^2)[exp\sigma^2 - (1-\delta)]$, the Variance of Lognormally distributed data with Zero-valued Observations

For the assessment of the extent of variability among health care costs or among exposure measurements, confidence intervals or tests concerning the variance $\tilde{\sigma}^2$ of lognormally distributed data with zero-valued observations becomes necessary. Krishnamoorthy, Mathew and Ramachandran (2006) made inference about the lognormal variance, while Bebu and Mathew (2008) obtained confidence intervals for the ratio of variances in the case of the bivariate lognormal distribution. However, as far as we know no procedures are known for computing confidence intervals for $\tilde{\sigma}^2$.

The following table presents a simulation study using Jeffreys' Independence prior

$$p(\tilde{\theta}) \propto \sigma^{-2}\delta^{-1/2}(1-\delta)^{-1/2}$$

From the table it is clear that the Bayesian procedure gives adequate coverage. The credibility intervals in the case of small n and large σ^2 are however quite wide.

With respect to the interval lengths, the HPD intervals are a great improvement on the equal-tail intervals, particularly if n is small.

Similar results were obtained from other simulation studies where $\delta = 0.3$ and $n = 15$ and 20. The Jeffreys' Rule prior on the other hand gives coverage probabilities that are somewhat too small.

Table 5
Simulation study for constructing two-sided 95% Bayesian confidence intervals for $\tilde{\sigma}^2 = (1-\delta)exp(2\mu+\sigma^2)[exp\sigma^2 - (1-\delta)]$ using Jeffreys' Independence prior

$\boldsymbol{\delta = 0.1}$			**Equal-tail Interval**			**HPD Interval**		
$\boldsymbol{n}$	$\boldsymbol{\mu}$		$\boldsymbol{\sigma^2 = 0.25}$	$\boldsymbol{\sigma^2 = 1}$	$\boldsymbol{\sigma^2 = 2.25}$	$\boldsymbol{\sigma^2 = 0.25}$	$\boldsymbol{\sigma^2 = 1}$	$\boldsymbol{\sigma^2 = 2.25}$
10	**0**	**Coverage**	95.07	95.05	95.34	95.43	95.05	95.45
		Length	82.874	1.4e16	3.3e35	7.0863	3.3e10	2.4e24
	1	**Coverage**	95.41	95.11	94.95	95.66	94.94	94.99
		Length	4033.4	1.7e14	5.5e53	411.00	3.18e9	4.5e34
	2	**Coverage**	95.06	94.86	94.93	95.47	94.65	95.29
		Length	1371.9	2.7e16	2e180	303.44	1.9e11	1.6e88
15	**0**	**Coverage**	95.46	94.90	95.20	95.41	94.81	95.45
		Length	2.7576	5.79e4	8.8e13	1.8573	3317.7	9.9e10
	1	**Coverage**	95.14	95.04	94.68	95.44	95.04	95.16
		Length	20.724	4.53e5	2.8e14	13.870	2.47e4	7.7e11
	2	**Coverage**	95.00	94.83	95.08	95.27	94.38	94.86
		Length	155.78	2.90e8	4.4e17	104.12	3.71e6	1.5e13
20	**0**	**Coverage**	95.12	94.53	95.05	95.59	95.12	95.50
		Length	1.6494	1465.4	1.64e9	1.2664	330.75	4.26e7
	1	**Coverage**	95.21	94.92	95.32	95.80	95.16	95.08
		Length	12.069	6765.9	2.1e10	9.2813	1897.6	4.18e8
	2	**Coverage**	95.30	95.23	95.52	95.80	94.95	95.63
		Length	90.337	2.02e5	2.5e11	69.396	2.70e4	2.79e9
50	**0**	**Coverage**	94.99	95.00	95.33	95.37	94.92	95.45
		Length	0.6237	27.362	6647.4	0.5691	20.024	2964.5
	1	**Coverage**	95.28	95.07	94.95	95.47	95.32	94.88
		Length	4.6485	203.27	4.98e4	4.2399	148.77	2.21e4
	2	**Coverage**	95.03	95.05	94.51	95.49	94.85	94.71
		Length	34.284	1460.6	4.16e5	31.267	1070.3	1.82e5

4 Comparison of the Means of Two Lognormal Populations containing Zero Values

The situation that will be investigated in this section is similar to that of section 3 only with the added complexity of analyzing the results when not only a single sample is taken, but when the means of two populations containing zero values is considered.

Therefore, assume that the probability of obtaining zero-valued observations from the $j-th$ population ($j = 1,2$) is δ_j where $0 < \delta_j < 1$. Denote $Y_{ij} = lnX_{ij} \sim N(\mu_j, {\sigma_j}^2)$ where X_{ij} $(i = 1,2,\dots,n_{j1})$ is a random sample from the $j-th$ population. In addition $X_{ij} = 0$ for $i = n_{j1}+1,\dots,n_j$ and $n_{j0} = n_j - n_{j1}$.

To compare the two population means confidence intervals will be constructed for $ln\widetilde{M}_1 - ln\widetilde{M}_2$, where $\widetilde{M}_j = (1-\delta_j)exp\left(\mu_j + \frac{1}{2}{\sigma_j}^2\right)$, $(j = 1,2)$, is the mean of the $j-th$ population.

Let $\boldsymbol{\theta} = [\delta_1 \quad \mu_1 \quad {\sigma_1}^2 \quad \delta_2 \quad \mu_2 \quad {\sigma_2}^2]'$ then the likelihood function is given by

$$L(\boldsymbol{\theta} \mid data) \propto \prod_{j=1}^{2} \{\delta_j^{n_{j0}} (1-\delta_j)^{n_{j1}} \prod_{i=1}^{n_{j1}} (\frac{1}{\sigma_j^2})^{\frac{1}{2}} \exp[-\frac{(y_{ij}-\mu_j)^2}{2\sigma_j^2}]\} \qquad (17)$$

and the Fisher Information matrix is

$$I(\boldsymbol{\theta}) = diag\left[\frac{n_1}{\delta_1(1-\delta_1)} \quad \frac{n_1(1-\delta_1)}{\sigma_1^2} \quad \frac{n_1(1-\delta_1)}{2\sigma_1^4} \quad \frac{n_2}{\delta_2(1-\delta_2)} \quad \frac{n_2(1-\delta_2)}{\sigma_2^2} \quad \frac{n_2(1-\delta_2)}{2\sigma_2^4}\right] \qquad (18)$$

The Independence Jeffreys' prior is then given by:

$$p(\boldsymbol{\theta}) \propto \prod_{j=1}^{2} \sigma_j^{-2} \delta_j^{-1/2} (1-\delta_j)^{-1/2} \qquad (19)$$

and from $|I(\boldsymbol{\theta})|^2$ the Jeffreys' Rule prior then becomes

$$p(\boldsymbol{\theta}) \propto \prod_{j=1}^{2} \sigma_j^{-3} \delta_j^{-1/2} (1-\delta_j)^{1/2} \qquad (20)$$

As was done in Zhou and Tu (2000) the following different sample sizes will be considered:

Table 6
Sample Sizes Analyzed by Monte Carlo Simulation Techniques

n_1	n_2
10	10
25	25
50	50
100	100
10	25
25	10
25	50

Zero proportions with different skewness coefficients are also considered. Based on these generated samples the credibility intervals are constructed.

The results from the simulation study performed by Zhou and Tu (2000) for the Maximum Likelihood and Bootstrap methods have been supplied as well for the purposes of comparison. In addition, the MOVER and Bayesian frameworks for the two different Jeffreys' prior distributions are compared (due to space only average results for designs 3 and 5 are given). The overall average is also given. The coverage probabilities and interval lengths are given in Table 8.

In the following table the parameter settings used in the simulation study are presented (the skewness coefficients for samples 1 and 2 are reported under headings γ_1 and γ_2):

Table 7
Parameter Settings used in the Simulation Study

Design	σ_1^2	σ_2^2	δ_1	δ_2	γ_1	γ_2
1	3.0	1.0	0.0	0.0	96.4851	6.1849
2	4.0	4.0	0.0	0.0	414.3593	414.3593
3	3.0	1.0	0.1	0.1	100.9809	6.1763
4	2.0	0.5	0.0	0.1	23.7323	2.6848
5	2.0	0.5	0.1	0.2	24.5572	2.5806

Table 8
Comparison of Results for the Ratio of Two Populations – Summary Results for Simulation Studies – 95% Two-sided Confidence Intervals for $ln\widetilde{M}_1 - ln\widetilde{M}_2$

		Equal – Tail*		HPD Intervals	
Design	Method	Coverage	Width	Coverage	Width
3	ML	92.37 (6.75, 0.88)	2.51		
	Bootstrap	92.98 (3.93, 3.08)	2.70		
	Independence Jeffreys'	95.28 (2.32, 2.40)	3.59	96.02 (1.22, 2.76)	3.40
	Jeffreys' Rule	94.68 (2.24, 3.08)	3.15	95.15 (1.20, 3.66)	3.02
	MOVER	95.62 (2.11, 2.27)	3.62		
5	ML	92.74 (6.19, 1.08)	1.87		
	Bootstrap	93.46 (3.49, 3.05)	1.98		
	Independence Jeffreys'	95.41 (2.35, 2.24)	2.58	95.95 (1.33, 2.72)	2.46
	Jeffreys' Rule	94.98 (2.12, 2.90)	2.28	95.28 (1.23, 3.49)	2.20
	MOVER	95.68 (2.06, 2.26)	2.57		
Overall	ML	93.09 (5.66, 1.25)	2.46		
	Bootstrap	93.15 (3.68, 3.17)	2.62		
	Independence Jeffreys'	95.32 (2.36, 2.32)	3.44	96.09 (1.33, 2.57)	3.30
	Jeffreys' Rule	94.79 (2.25, 2.95)	3.06	95.31 (1.35, 3,33)	2.96
	MOVER	95.64 (2.07, 2.29)	3.46		

*Refers to Equal-tail Bayesian intervals, Maximum Likelihood Methods, Bootstrap and MOVER estimates.

In this case it is once again evident that the Bayesian methods are substantially better than both the ML and bootstrap methods. Both the ML and Bootstrap methods results in substantial undercoverage. The MOVER is a vast improvement on these methods and results in adequate coverage. However, the Bayesian methods result in adequate coverage and particularly the HPD intervals result in more efficient (narrower) intervals than even the MOVER. According to Table 8, the prior that gives the shortest HPD intervals with correct coverage is the Jeffreys' Rule prior. The Independence Jeffreys' prior on the other hand gives the best coverage for equal-tail intervals.

4.1 Example: Rainfall Data

For the purposes of comparison of the different methods, an example was chosen using raw data obtained from the South African Weather Service. The data consisted of the monthly rainfall totals for the cities of Bloemfontein and Kimberley, two South African cities, over a period of 69 to 70 years of measurement. However, these two cities are both located in relatively arid regions and are characterized by mainly summer rainfall. For that reason, the winter months do contain some rainfall data, but also contain many years where the total monthly rainfall data was zero. Probability plots as well as the Shapiro-Wilks (1965) test indicate that the lognormal distribution is a better fit than the normal distribution.

The aim is to produce two-sided 95% confidence intervals for the ratio of the two means, $\frac{\tilde{M}_1}{\tilde{M}_2} = \frac{(1-\delta_1)exp\left(\mu_1+\frac{1}{2}\sigma_1^2\right)}{(1-\delta_2)exp\left(\mu_2+\frac{1}{2}\sigma_2^2\right)}$. If show that the confidence interval includes 1 then the conclusion is that there is no difference between the mean rainfalls.

The data can be summarized as follows:

Table 9
Summary of the Rainfall Data

City	Parameter	Value
Bloemfontein	Number of Years of Available Data	70
	Number of Zero Valued Observations	18
	Mean of Log-Transformed Data	1.9578
	Variance of Log-Transformed Data	2.1265
Kimberley	Number of Years of Available Data	69
	Number of Zero Valued Observations	10
	Mean of Log-Transformed Data	1.0526
	Variance of Log-Transformed Data	3.1589

In order to compare the results, both the MOVER was applied to the data as well as the Bayesian methods described in the preceding text for the following priors: Independence Jeffery's Prior (Prior 1 in the table), the Jeffery's Rule Prior (Prior 2), the Reference Prior and the Probability Matching Prior. The maximum likelihood and bootstrap procedures derived by Zhou and Tu (2000) are also given. The results are presented in the following table:

Table 10
Summary of Results for the Rainfall Data – Two-sided 95% Confidence Intervals for $\frac{\tilde{M}_1}{\tilde{M}_2} = \frac{(1-\delta_1)exp\left(\mu_1+\frac{1}{2}\sigma_1^2\right)}{(1-\delta_2)exp\left(\mu_2+\frac{1}{2}\sigma_2^2\right)}$

	Maximum Likelihood	Bootstrap	MOVER	Prior 1	Prior 2	Reference Prior	Probability Matching Prior
Lower Limit	0.5009	0.5271	0.4486	0.4485	0.4520	0.4410	0.4587
Upper Limit	3.2805	3.4710	3.3428	3.3245	3.3048	3.3548	3.3141

From Table 10 it is clear that the Bayesian and MOVER intervals are for practical purposes the same. The maximum likelihood and bootstrap intervals on the other hand differ somewhat from these intervals. It can furthermore be seen that 1 is included in all the 95% confidence intervals and thus there is no difference between the mean rainfalls for the two cities.

The following graphs also illustrate these results:

Prior 1 – The Independence Jeffreys Prior

Histogram of Simulated Log Transformed Ratio of Means

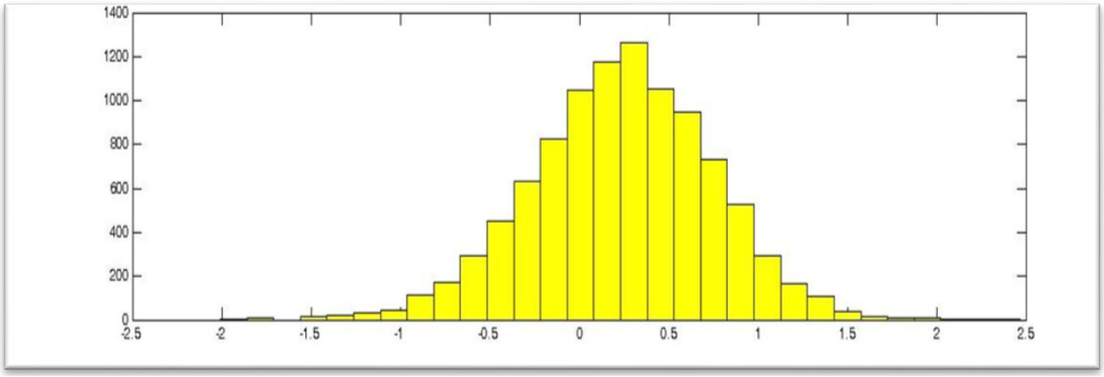

Histogram of Simulated Lognormal Ratio of Means

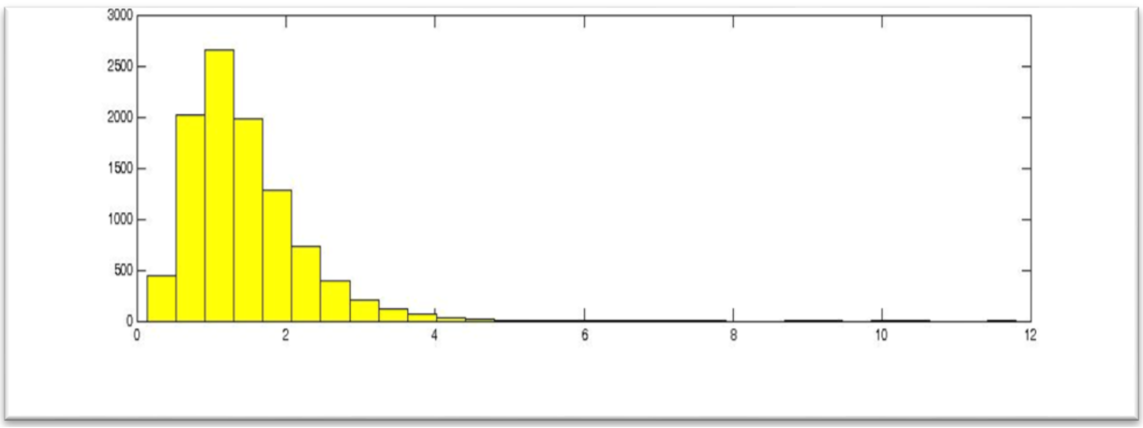

5 The Bivariate Lognormal Distribution

The purpose of this section is to develop Bayesian procedures for computing credibility intervals for the ratio of means and variances of the bivariate lognormal distribution. The Bayesian procedures using different prior distributions will be compared with the generalized confidence interval (GCI) approach used by Bebu and Mathew (2008) and the MOVER described by Zhou et al (2009).

As mentioned by Bebu and Mathew (2008) the bivariate (and multivariate) lognormal distributions are particularly suited for the study of size distributions. Asbestos fiber sizes for example are often bivariately lognormally distributed.

5.1 Notation and Description of the Setting

Let $[X_1 \quad X_2]'$ follow a bivariate lognormal distribution so that $[Y_1 \quad Y_2]' = [lnX_1 \quad lnX_2]'$ follows a bivariate normal distribution with mean parameters $\boldsymbol{\mu} = [\mu_1 \quad \mu_2]'$ and covariance matrix $\Sigma = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}$ where ρ is the correlation coefficient between Y_1 and Y_2. Thus

$$E(X_i) = exp\left(\mu_i + \frac{1}{2}\sigma_i^2\right), \qquad i = 1,2$$

Bayesian confidence intervals for the parameters

$$\tilde{\theta} = (\mu_1 - \mu_2) + \frac{1}{2}(\sigma_1^2 - \sigma_2^2) \qquad (21)$$

will be constructed in order to compare the ratio of the lognormal means.

For the ratio of the variances Bayesian confidence intervals for

$$\tilde{\delta} = \frac{Var(Y_1)}{Var(Y_2)} = exp(2\tilde{\theta})\left\{\frac{exp(\sigma_1^2)-1}{exp(\sigma_2^2)-1}\right\} \qquad (22)$$

will be obtained.

Let $[X_{1k}, X_{2k}]', k = 1,2, \dots, n$ denote a random sample from the bivariate lognormal distribution and let $\boldsymbol{Y}_k = [Y_{1k}, Y_{2k}]' = [lnX_{1k}, lnX_{2k}]'$. The sufficient statistics (for $n \geq 3$) are $\bar{\boldsymbol{Y}} = [\bar{Y}_1, \bar{Y}_2]'$ and $S = \sum_{k=1}^{n}(\boldsymbol{Y}_k - \bar{\boldsymbol{Y}})(\boldsymbol{Y}_k - \bar{\boldsymbol{Y}})' = \begin{bmatrix} s_{11} & r\sqrt{s_{11}s_{22}} \\ r\sqrt{s_{11}s_{22}} & s_{22} \end{bmatrix}$ where $\bar{Y}_i = \frac{1}{n}\sum_{k=1}^{n} Y_{ik}$, $s_{ij} = \sum_{k=1}^{n}(Y_{ik} - \bar{Y}_i)(Y_{jk} - \bar{Y}_j)'$ and $r = \frac{s_{12}}{\sqrt{s_{11}}\sqrt{s_{22}}}$. It is well known that $S \sim W_2(\Sigma, \mathrm{n}-1)$ i.e. the bivariate Wishart distribution with scale matrix Σ and degrees of freedom $\mathrm{n}-1$.

Bebu and Mathew (2008) (see also Berger and Sun (2008)) used the following properties of the Wishart distribution to construct a generalized pivot statistic:

$$T_1 = \frac{S_{11}}{\sigma_1^2} \sim \chi_{n-1}^2 \qquad (23)$$

$$T_2 = \frac{S_{22}(1-r^2)}{\sigma_2^2(1-\rho^2)} \sim \chi_{n-2}^2 \qquad (24)$$

$$T_3 = \left[\frac{S_{11}}{\sigma_2^2(1-\rho^2)}\right]^{1/2}\left[\frac{r\sqrt{S_{22}}}{\sqrt{S_{11}}} - \frac{\rho\sigma_2}{\sigma_1}\right] = Z \qquad (25)$$

Here Z is a standard normal random variable and χ_{n-1}^2 and χ_{n-2}^2 are chi-squared random variables with the indicated degrees of freedom.

From (22) to (24) it follows that simulated values for the unknown parameters σ_1^2, σ_2^2 and ρ can be obtained from the following distributions:

$$\frac{S_{11}}{\chi_{n-1}^{2*}} = \sigma_1^{2*} \qquad (26)$$

$$\sigma_2^{2*} = S_{22}(1 - r^2)\left\{\frac{1}{\chi_{n-2}^{2*}} + \frac{1}{\chi_{n-1}^{2*}}\left(\frac{Z^*}{\sqrt{\chi_{n-2}^{2*}}} - \frac{r}{\sqrt{1 - r^2}}\right)^2\right\} \qquad (27)$$

$$\rho^* = \varphi(Y^*), \quad \varphi(Y^*) = \frac{Y^*}{\sqrt{1+Y^{*2}}}, \quad Y^* = \frac{-Z^*}{\sqrt{\chi_{n-1}^{2*}}} + \frac{\sqrt{\chi_{n-2}^{2*}}}{\sqrt{\chi_{n-1}^{2*}}}\frac{r}{\sqrt{1-r^2}} \qquad (28)$$

The asterisk is used to represent a random realized observation from the implied distribution. Equations (26) – (28) are (for $a = 1$ and $b = 2$) similar to R_{22}, R_{12} and R_{11} given in Bebu and Mathew (2008).

Berger and Sun (2008) further proved that if the prior

$$\pi_{ab}(\mu_1, \mu_2, \sigma_1, \sigma_2, \rho) \propto \frac{1}{\sigma_1^{3-a}\sigma_2^{2-b}(1-\rho^2)^{2-b/2}} \qquad (29)$$

is used then for $a = 1$ and $b = 2$ the constructive posterior distributions of σ_1^2, σ_2^2 and ρ are given by equations (26) – (28). In other words, if the right-Haar prior ($a = 1$, $b = 2$) is used then the posterior distributions are also the fiducial distributions for the parameters as found by Fisher (1930) and used by Bebu and Mathew (2008).

Special cases of this class of priors are the Jeffreys-rule prior $\pi_{10} = \pi_J$, the right-Haar prior $\pi_H = \pi_{12}$, the Independence Jeffreys prior $\pi_{21} = \pi_{IJ}$ and $\pi_{11} = \pi_{RO}$.

In our simulation experiment of $\tilde{\theta}$ and $\tilde{\delta}$ the above mentioned priors will be used as well as five other priors: $\pi_{MS} \propto \sigma_1^{-1}\sigma_2^{-1}(1-\rho^2)^{-1/2}$, $\pi_{RP} \propto \sigma_1^{-1}\sigma_2^{-1}(1-\rho^2)^{-1}$, $\pi_S \propto \sigma_1^{-1}\sigma_2^{-1}$, $\pi_{R\sigma} \propto \sqrt{1+\rho^2}\sigma_1^{-1}\sigma_2^{-1}(1-\rho^2)^{-1}$, $\tilde{\pi}_{R\sigma} \propto \sigma_1^{-1}\sigma_2^{-1}(1-\rho^2)^{-1}(2-\rho^2)^{-1/2}$. For these five priors independent samples from their marginal posterior $\pi\ (\sigma_1, \sigma_2, \rho|data)$ can also easily be obtained by an acceptance-rejection algorithm (Berger and Sun (2008)). Once

$$\Sigma^* = \begin{bmatrix} \sigma_1^{2*} & \rho\sigma_1{}^*\sigma_2{}^* \\ \rho\sigma_1{}^*\sigma_2{}^* & \sigma_2^{2*} \end{bmatrix} \qquad (30)$$

is simulated it follows that

$$\theta|\Sigma^*, \text{data} \sim \text{N}\left\{(\bar{y}_1 - \bar{y}_2) + \frac{1}{2}(\sigma_1^{2*} - \sigma_2^{2*}), \frac{1}{\text{n}}(\sigma_1^{2*} + \sigma_2^{2*} - 2\rho\sigma_1{}^*\sigma_2{}^*)\right\} \qquad (31)$$

The $100(1-\alpha)\%$ confidence interval for $\tilde{\theta}$ in the case of the MOVER follows from section 2.7 (see also Zou et al (2009)) as

$$\tilde{L} = \hat{\theta}_1 - \hat{\theta}_2 - \sqrt{\left(\hat{\theta}_1 - l_1\right)^2 + \left(u_2 - \hat{\theta}_2\right)^2 - 2r\left(\hat{\theta}_1 - l_1\right)\left(u_2 - \hat{\theta}_2\right)}$$

$$\tilde{U} = \hat{\theta}_1 - \hat{\theta}_2 + \sqrt{\left(\hat{\theta}_2 - l_2\right)^2 + \left(u_1 - \hat{\theta}_1\right)^2 - 2r\left(\hat{\theta}_2 - l_2\right)\left(u_1 - \hat{\theta}_1\right)}$$

The ratio of the variances

$$\tilde{\delta} = exp\left(2\tilde{\theta}\right)\left\{\frac{exp(\sigma_1^2) - 1}{exp(\sigma_2^2) - 1}\right\}$$

can also easily be simulated using (30) and (31).

5.2 Numerical Results

For the nine different priors the coverage probabilities of the 95% Bayesian confidence intervals for $\tilde{\delta}$ and $\tilde{\theta}$ are given in Tables 11 and 12. These intervals are compared to the MOVER procedure.

The coverage probabilities and interval lengths of the MOVER are in some cases better and in other cases worse than those of the Bayesian procedures. In general one can say that the Bayesian methods result in narrower intervals. A surprise is the small interval lengths and good coverage probabilities of the priors π_S and π_{MS} .

Although this is a small simulation study it gives an idea of the performance of the various priors. For the the sake of brevity aggregated results (by different sample sizes and correlation coefficient values) are given below.

Table 11

Coverage Probabilities of the 95% Confidence Intervals for $\tilde{\theta}$ and $\tilde{\delta}$ in the case of the Bivariate Lognormal Distribution: a Comparison with the MOVER, by sample size

		$\tilde{\theta}$				$\tilde{\delta}$			
Method	n	Cover	Length	Cover (HPD)	Length (HPD)	Cover	Length	Cover (HPD)	Length (HPD)
MOVER	5	94.0678	26.9783						
π_{21}		94.2578	55.4800	98.1156	41.1144	93.5489	1.25E+159	91.7333	3.33E+276
π_{10}		93.0000	25.2844	95.5978	20.2322	92.0822	1.11E+214	92.8200	1.11E+150
π_{12}		94.7411	32.5488	97.3300	27.2381	93.9511	1.50E+135	96.3733	1.29E+84
π_{RO}		93.7956	26.9144	96.1956	22.0011	93.0289	1.11E+277	94.6044	3.33E+178
π_{MS}		95.7544	26.5189	97.3089	22.2567	95.6589	1.25E+71	95.7978	3.33E+209
π_{RP}		95.3578	33.1967	97.8133	26.7178	95.0000	1.11E+252	94.4178	3.33E+158
π_S		95.9333	24.6289	96.7889	21.0100	96.1856	2.22E+231	96.9478	6.67E+155
$\pi_{R\sigma}$		95.3678	34.5844	98.1344	27.7144	94.8211	2.22E+291	94.0944	1.11E+192
$\tilde{\pi}_{R\sigma}$		95.3511	34.5833	98.0511	27.6844	94.8844	1.11E+298	94.0978	3.33E+182
MOVER	10	94.6078	10.6871						
π_{21}		94.8511	12.1234	97.1533	10.8180	94.5156	3.33E+70	93.1311	6.67E+54
π_{10}		94.4533	9.8200	95.7022	8.8960	94.1178	2.22E+38	94.4156	1.11E+31
π_{12}		94.7467	10.5210	96.0222	9.5938	94.4511	2.22E+64	95.6844	1.11E+49
π_{RO}		94.5911	10.0373	95.8822	9.1216	94.3400	4.44E+44	94.8844	1.11E+34
π_{MS}		95.2367	9.9548	96.3678	9.0937	95.2633	2.22E+45	95.3856	5.56E+34
π_{RP}		94.9733	10.7874	96.7533	9.7636	94.9533	4.44E+56	94.5178	1.11E+44
π_S		95.2800	9.3731	95.6433	8.6132	95.5100	2.22E+51	96.1800	8.89E+39
$\pi_{R\sigma}$		95.1467	10.9351	96.9400	9.8853	95.0200	6.67E+47	94.2600	4.44E+36
$\tilde{\pi}_{R\sigma}$		95.0800	10.9809	96.9956	9.9244	94.9689	1.11E+60	94.2022	2.22E+47
MOVER	20	95.4089	6.0731						
π_{21}		95.0000	6.0153	96.2000	5.7101	94.8444	1.11E+16	93.9444	1.11E+13
π_{10}		94.6311	5.5602	95.5222	5.2979	94.5956	8.89E+15	94.5911	3.33E+13
π_{12}		95.0422	5.7659	95.6556	5.4981	94.8800	1.11E+21	95.5933	3.33E+17
π_{RO}		94.8267	5.6294	95.5511	5.3669	94.8244	1.11E+17	95.0222	3.33E+14
π_{MS}		94.9844	5.5978	95.5600	5.3469	95.0044	2.22E+19	95.1378	7.78E+15
π_{RP}		95.2911	5.7830	96.1667	5.5156	95.1333	5.56E+14	94.7689	1.11E+12
π_S		94.9122	5.4377	95.2022	5.1983	95.0322	1.00E+15	95.5322	1.11E+12
$\pi_{R\sigma}$		95.0689	5.8067	96.1689	5.5379	95.0333	1.11E+18	94.4267	5.56E+14
$\tilde{\pi}_{R\sigma}$		95.1511	5.8232	96.0289	5.5526	95.0444	3.33E+17	94.4600	2.22E+14

The results indicate that in the small sample size situation the MOVER performs well when comparing the difference of the means. Efficiency in terms of average interval length does improve as the sample size increases. The advantage of the Bayesian framework is the construction of HPD intervals. For all choices of prior distribution both the coverage and the average interval length are improved with the HPD intervals.

With regards to coverage alone some choices of prior distributions performed well, given that the intention was not specifically for application to the ratio of means from two populations, but more towards the variance. It is clear that π_{RP} , $\pi_{R\sigma}$, $\tilde{\pi}_{R\sigma}$, π_S and π_{MS} all achieve better coverage than the independence Jeffreys' prior and in particular, π_S achieves this with an average interval length that is better than both the MOVER and all Jeffreys priors.

Table 12

Coverage Probabilities of the 95% Confidence Intervals for $\tilde{\theta}$ and $\tilde{\delta}$ in the case of the Bivariate Lognormal Distribution: a Comparison with the MOVER, by correlation coefficient

		$\tilde{\theta}$				$\tilde{\delta}$			
Method	ρ	Cover	Length	Cover (HPD)	Length (HPD)	Cover	Length	Cover (HPD)	Length (HPD)
MOVER	-0.9	97.754	16.748						
π_{21}		94.836	23.267	96.958	18.073	94.609	2.22E+148	92.820	1.11E+93
π_{10}		94.536	13.309	95.540	11.352	94.264	3.33E+74	94.171	5.56E+50
π_{12}		94.910	14.265	96.008	12.379	94.880	1.03E+17	95.516	1.25E+77
π_{RO}		94.736	13.589	95.713	11.649	94.633	1.11E+84	94.929	1.11E+53
π_{MS}		95.129	11.935	95.680	10.567	95.628	1.00E+67	96.000	2.22E+45
π_{RP}		95.169	14.650	96.444	12.446	95.371	1.00E+93	94.804	2.22E+61
π_S		94.871	10.931	94.789	9.816	95.772	1.11E+76	96.801	7.78E+38
$\pi_{R\sigma}$		95.273	14.947	96.689	12.670	95.282	2.22E+83	94.629	3.33E+52
$\tilde{\pi}_{R\sigma}$		95.364	15.022	96.733	12.720	95.307	4.44E+85	94.469	1.00E+55
MOVER	0.1	94.539	14.744						
π_{21}		94.784	28.826	97.333	23.352	94.164	3.75E+70	93.151	3.33E+276
π_{10}		93.607	15.810	95.767	13.622	92.907	1.11E+214	93.578	1.11E+150
π_{12}		94.424	21.120	96.618	18.529	93.616	7.50E+130	96.236	1.25E+72
π_{RO}		94.000	17.045	96.002	14.965	93.342	1.11E+277	94.896	3.33E+178
π_{MS}		94.914	19.590	97.152	17.051	94.370	2.50E+45	94.081	3.33E+209
π_{RP}		94.940	21.984	97.396	18.770	94.420	1.11E+252	93.907	3.33E+158
π_S		94.912	18.780	97.037	16.463	94.608	2.22E+231	94.326	6.67E+155
$\pi_{R\sigma}$		94.824	23.002	97.398	19.524	94.216	2.22E+291	93.584	1.11E+192
$\tilde{\pi}_{R\sigma}$		94.907	22.915	97.324	19.445	94.349	1.11E+298	93.720	3.33E+182
MOVER	0.9	91.791	12.246						
π_{21}		94.489	21.525	97.178	16.218	94.136	1.11E+159	92.838	7.78E+99
π_{10}		93.942	11.545	95.516	9.452	93.624	1.11E+98	94.078	6.67E+63
π_{12}		95.196	13.451	96.382	11.422	94.787	1.00E+135	95.900	1.00E+84
π_{RO}		94.478	11.947	95.913	9.875	94.218	7.78E+91	94.687	1.11E+63
π_{MS}		95.932	10.546	96.404	9.079	95.929	1.11E+71	96.240	4.44E+44
π_{RP}		95.513	13.133	96.893	10.781	95.296	4.44E+88	94.993	4.44E+57
π_S		96.342	9.729	95.809	8.543	96.348	1.11E+61	97.533	7.78E+39
$\pi_{R\sigma}$		95.486	13.378	97.157	10.944	95.377	1.11E+90	94.568	3.33E+59
$\tilde{\pi}_{R\sigma}$		95.311	13.451	97.018	10.997	95.242	1.11E+105	94.571	6.67E+67

From the above results it is apparent that in situations of both high negative and positive correlations the MOVER results in substantial overcoverage for the difference of two means, even though the interval width is markedly less than the Jeffreys rule prior. The independence Jeffreys prior and the right Haar prior are both improvements on both the MOVER and the Jeffreys rule prior with regards to the difference of the two means. The trend is again seen when comparing the HPD intervals. Of particular interest again is the performance of the π_S prior, which results in the best coverage as well as the most efficient intervals in terms of average length.

Furthermore, in situations of low correlation between the two means the MOVER appears to perform rather well. The only improvement on the MOVER is the HPD interval for the independence Jeffreys prior. Thus, both the MOVER and the independence Jeffreys prior are particularly well suited to situations of small correlation.

When comparing the two variances we once again see that HPD intervals are an improvement on the standard equal-tailed intervals with regards to both coverage and interval length. When there is not a high degree of correlation it seems that the right Haar prior is the most likely choice.

6 Concluding Remarks

The lognormal distribution is currently used extensively to describe the distribution of positive random variables. One particular application of the data is statistical inference with regards to the mean of the data. In this article we compare the performance of the MOVER-based confidence interval estimated and the generalized confidence interval procedure to coverage of credibility intervals obtained using Bayesian methodology and a variety of different prior distributions to estimate the appropriateness of each. For the Bayesian approach both the equal-tail and highest posterior density (HPD) credibility intervals are presented. The simulation studies show that the constructed Bayesian confidence intervals have satisfying coverage probabilities and in some cases outperform the MOVER and generalized confidence interval results. The Bayesian inference procedures (hypothesis tests and confidence intervals) are also extended to the difference between two lognormal means as well as to the case of zero-valued observations. The Bayesian methods considered all provide better coverage than the maximum likelihood and bootstrap methods. The HPD intervals are an improvement on both the equal-tail and MOVER intervals. A surprise is the small interval lengths and good coverage of the Jeffreys' Rule HPD intervals.

For the assessment of the extent of variability among health care costs or among exposure measurements, confidence intervals or tests concerning the variance $\tilde{\sigma}^2$ of lognormally distributed data with zero-valued observations becomes necessary. However, as far as we know no procedures are known for computing confidence intervals for $\tilde{\sigma}^2$. From our simulation study using Jeffreys' independence prior it becomes clear that the Bayesian procedure gives adequate coverage. The credibility intervals in the case of small sample sizes and large σ^2 are however quite wide. With respect to interval lengths the HPD intervals are a great improvement on the equal-tail intervals, particularly when n is small.

In the last section of this paper the bivariate lognormal distribution is considered and Bayesian confidence intervals (using nine different priors) are simulated for the difference between two correlated lognormal means as well as for the ratio of two correlated lognormal variances. It is also mentioned that if the right-Haar prior is used then the posterior distributions are also the fiducial distributions for the parameters as found by Fisher (1930) and used by Bebu and Mathew (2008). Although this is only a short simulation study it can be concluded that some of the Bayesian intervals are an improvement over the MOVER with respect to coverage probabilities and interval lengths.

8 BIBLIOGRAPHY

Bebu, I. and Mathew, T. (2008). "Comparing the means and variances of a bivariate lognormal distribution," *Statistics in Medicine*, 27, pp 2684 – 2696.

Berger, J.O., Bernardo, J.M. (1992). "On the development of the reference prior Method," *Bayesian Statistics 4: Proceedings of the Fourth Valencia Meeting.* Eds. Bernardo, J.M., Berger, J.O., Dawid, A.P., Smith, A.F.M. Oxford, U.K., Clarendon Press, pp. 35 – 60.

Berger, J.O., Sun, D. (2008). "Objective priors for the bivariate normal model," *The Annals of Statistics*, volume 36, no 1, pp 963-988

Box, G.E.P., Tiao, G.C. (1973). "Bayesian inference in statistical analysis," *Reading, MA,* Addison-Wesley.

Chen, Y.H., Zhou, X.H. (2006). "Interval estimates for the ratio and difference of two lognormal means," *Statistics in Medicine*, 25, pp 4099 – 4113.

Datta, G.S., Ghosh, J.K. (1995). "On priors providing frequentist validity of Bayesian inference," *Biometrika*, 82, pp. 37 – 45.

Harvey, J., Groenewald, P.C.N. and ven der Merwe, A.J. (2010). "Bayesian confidence intervals for the mean of a lognormal distribution: a comparison with the MOVER and generalized confidence interval procedure," Technical report no. 408, Department of Statistics and Actuarial Science, University of the Free State.

Krishnamoorthy, K., Mathew, T. (2003). "Inferences on the means of lognormal distributions using generalized p-values and generalized confidence intervals." *Journal of Statistical Planning and Inference,* 115, pp. 103 - 121.

Krishnamoorthy, K., Mathew, T., Ramachandran, G. (2006). "Generalized p-values and confidence intervals: A novel approach for analyzing lognormally distributed exposure data," *Journal of Occupational and Environmental Hygiene,* 3 pp 642 – 650.

Jeffreys, H. (1967). "Theory of probability," *Oxford University Press,* London.

Shapiro, S.S., Wilk, M.B. (1965). "An analysis of variance test for normality (Complete Samples)." *Biometrika,* 52: pp. 591 - 611.

Tian, L., Wu, J. (2007a). "Inferences on the common mean of several lognormal populations: The generalized variable approach," *Biometrical Journal*, 49, pp 944 – 951.

Tian, L., Wu, J. (2007b). "Inferences on the mean response in a log-regression model: The generalized variable approach," *Statistics in Medicine,* 25, pp 5180 - 5188.

Tian, L., (2005). "Inferences on the mean of zero-inflated lognormal data: the generalized variable approach." *Statistics in Medicine*, pp. 3223 - 3232.

Tian, L., Wu, J., (2006). "Confidence intervals for the mean of lognormal data with excess zeros." *Biometrical Journal*, 48 no. 1: pp. 149 - 156.

Zhou, X.H., Tu, W. (2000). "Interval estimation for the ratio in means of log-normally distributed medical costs with zero values," *Computational Statistics and Data Analysis*, 35, pp. 201 – 210.

Zou, Y.Z., Taleban, J., Huo, C.Y. (2009). "Confidence interval estimation for lognormal data with application to health economics," *Computational Statistics and Data Analysis*, 53, pp. 3755 – 3764.