

Bayesian Estimation of Log-Normal Means with Finite Quadratic Expected Loss

Enrico Fabrizi* and Carlo Trivisano†

Abstract. The log-normal distribution is a popular model in biostatistics and other fields of statistics. Bayesian inference on the mean and median of the distribution is problematic because, for many popular choices of the prior for the variance (on the log-scale) parameter, the posterior distribution has no finite moments, leading to Bayes estimators with infinite expected loss for the most common choices of the loss function. We propose a generalized inverse Gaussian prior for the variance parameter, that leads to a log-generalized hyperbolic posterior, for which it is easy to calculate quantiles and moments, provided that they exist. We derive the constraints on the prior parameters that yield finite posterior moments of order r . We investigate the choice of prior parameters leading to Bayes estimators with optimal frequentist mean square error. For the estimation of the lognormal mean we show, using simulation, that the Bayes estimator under quadratic loss compares favorably in terms of frequentist mean square error to known estimators.

Keywords: Bayes estimators, generalized hyperbolic distribution, generalized inverse gamma distribution, Bessel functions.

1 Introduction

Suppose that a random variable X with mean ξ and variance σ^2 is normally distributed, such that $\exp(X) \sim \text{LogN}(\xi, \sigma^2)$. In this paper we consider the problem of Bayesian inference about functionals of (ξ, σ^2) of the form $\theta_{a,b} = \exp(a\xi + b\sigma^2)$ with $a, b \in \mathfrak{R}$ based on a random sample (X_1, \dots, X_n) . We may obtain the mean, the median, the mode and various non-central moments of the log-normal distribution for different choices of a, b . More specifically $a = 1$ and $b = 0$ yields the median $(\theta_{1,0})$, $a = 1$ and $b = -1$ yields the mode $(\theta_{1,-1})$ and $a = 1$ and $b = .5$ yields the mean $(\theta_{1,0.5})$. This estimation problem is of practical relevance because, when analyzing positively skewed data, it is common practice to take the log transformation and assume the normality of the transformed data. Krishnamoorthy et al. (2011), Zou et al. (2009), Gill (2004), Dagne (2001), Limpert et al. (2001) are just a few examples of research about inference on the log-normal mean and quantiles with applications in different fields of enquiry.

The problem has a long tradition in both the frequentist and the Bayesian literature. With reference to the latter, an important starting point is Zellner (1971). He considers diffuse priors of the type $p(\xi, \sigma) \propto \sigma^{-1}$ and he obtains the following results for the log-normal median: *i*) $p(\theta_{1,0}|\sigma, data)$ is a log-normal distribution; and *ii*) $p(\theta_{1,0}|data)$ is a log-t distribution.

*DISES, Università Cattolica, Piacenza, Italy, enrico.fabrizi@unicatt.it

†Dipartimento di Scienze Statistiche ‘P. Fortunati’, Università di Bologna, Italy, carlo.trivisano@unibo.it

Summarizing the log-t distribution is challenging using popular loss functions, such as the quadratic, because moments of all orders do not exist. For the log-normal mean Zellner (1971) shows that $p(\theta_{1,0.5}|\sigma, data)$ is a log-normal distribution. To obtain $p(\theta_{1,0.5}|data)$ he suggests studying the joint posterior $p(\log(\theta_{1,0.5}), \sigma|data)$, integrating out σ and then considering its exponential transformation. He notes that the integral over σ can be ‘expressed in terms of modified Bessel functions’ but that ‘it is the case that the posterior mean of $\theta_{1,0.5}$ does not exist’, which creates similar problems in summarizing the posterior distribution, (that is in obtaining ‘Bayes estimators’ of the parameter in question and easily interpretable measures of the information loss).

Most of Zellner’s paper focuses on the inference conditional on σ . He notes that, within the class of estimators of the form $k \exp(\bar{X})$ with $\bar{X} = n^{-1} \sum_{i=1}^n X_i$ and where k is a constant, the estimator for $\theta_{1,0.5}$ with minimum mean square error (MSE) is given by $\check{\theta}_{1,0.5} = \exp(\bar{X} + \sigma^2/2 - 3\sigma^2/2n)$. From a Bayesian point of view, this estimator may be justified as the minimizer of the posterior expected loss, provided that the relative quadratic loss function $L_{RQ} = [(\theta - \hat{\theta})/\theta]^2$ is adopted.

Another important reference is Rukhin (1986). Rukhin proposes the following ‘generalized’ prior:

$$p(\xi, \sigma) = p(\sigma) \propto \sigma^{-2\nu+n-2} \exp(-\sigma^2[\gamma^2/2 - 2(b - a^2/n)]), \quad (1)$$

with $\nu \in \Re$, $\gamma^2 > 4(b - a^2/n)$. Assuming the relative quadratic loss function L_{RQ} , he obtains an estimator for $\theta_{a,b}$ of the form $\hat{\theta}_{a,b}^{Ru} = \exp(a\bar{X})g(Y)$ that is given by

$$\hat{\theta}_{a,b}^{Ru} = \exp(a\bar{X}) \left(\frac{\beta}{\gamma}\right)^\nu \frac{K_\nu(\beta Y)}{K_\nu(\gamma Y)}, \quad (2)$$

$\beta = \gamma^2 - 2c$, $c = b - 3a^2/(2n)$ and $K_\nu(\cdot)$ is the modified Bessel function of the third kind (the Bessel-K function from now on). For a general introduction to Bessel functions, see Abramowitz and Stegun (1968), chapters 9 and 10. To obtain the values for the hyperparameters ν, γ , Rukhin (1986) chooses to minimize the frequentist MSE of $\hat{\theta}_{a,b}^{Ru}$. As the $K_\nu(\cdot)$ are quite difficult to handle, Rukhin uses a ‘small arguments’ approximation to $\hat{\theta}_{a,b}^{Ru}$ to propose a value for ν and a ‘large arguments’ approximation to propose a value for γ . Rukhin does not recognize that, with a simple change of variable, the prior he proposes may be seen as the product of a flat prior over the real line for ξ and the following prior on σ^2 :

$$p(\sigma^2) \propto (\sigma^2)^{-\nu+n/2-3/2} \exp(-\sigma^2[\psi^2/2 - 2(b - a^2/n)]) \quad (3)$$

which is the limit of a generalized inverse gamma distribution, $GIG(\lambda, \delta, \gamma)$ as $\delta \rightarrow 0$. The other parameters are given by $\lambda = -\nu + n/2 - 1/2$ and $\gamma^2 = \psi^2/2 - 2(b - a^2/n)$ (see section 2 for more details and notation). He does not provide the posterior distribution, so his proposal is inadequate for many inferential purposes (i.e., calculating of posterior variances or posterior probability intervals).

In this paper, we derive the posterior distribution of $\theta_{a,b}$ assuming a proper generalized inverse gamma prior on σ^2 (and a flat prior over the real line for ξ). We show

that this posterior is a log-generalized hyperbolic distribution and state the conditions on the hyperparameters that guarantee the existence of posterior moments of a given order. Once these conditions are met for the first two non-central moments, we discuss the Bayes estimators with the ordinary quadratic loss function $L_Q = (\theta - \hat{\theta})^2$.

The main results of the paper may be summarized as follows: *i*) we show that, given our choice of the prior distributions, Bayes estimators associated with the relative quadratic loss function L_{RQ} can be reconducted to posterior expectations provided that b is properly modified; *ii*) adopting a ‘small arguments’ approximation to the Bessel-K functions and a choice of hyperparameters aimed at minimizing the MSE, we show using simulation that our Bayes estimator of the mean, i.e. $\theta_{1,0.5}$ is substantially equivalent to the estimator proposed in Shen et al. (2006), which has been proven to be superior to many of the alternatives previously proposed in the literature.

The paper is organized as follows. In section 2 we briefly present the generalized inverse Gaussian and generalized hyperbolic distributions. In section 3, posterior distributions for σ^2 and $\theta_{a,b}$ are derived, and Bayes estimators under quadratic and relative quadratic losses are introduced. Section 4 is devoted to the choice of values to be assigned to the hyperparameters in order to obtain Bayes estimators with the minimum frequentist MSE. In section 5 we introduce a simulation exercise and discuss the results, while section 6 describes an application to a real data set. The concluding section 7 offers some conclusions and ideas for future research.

2 The generalized inverse Gaussian and generalized hyperbolic distributions

In this section we briefly introduce the generalized inverse Gaussian (GIG) and generalized hyperbolic (GH) distributions, establish the notation and mention some key properties that will be used later. For more details on these distributions, see Bibby and M. (2003) and Eberlein and von Hammerstein (2004) among others.

The density of the GIG distribution may be written as follows:

$$p(x) = \left(\frac{\gamma}{\delta}\right)^\lambda \frac{1}{2K_\lambda(\delta\gamma)} x^{\lambda-1} \exp\left\{-\frac{1}{2}(\delta^2 x^{-1} + \gamma^2 x)\right\}, \quad (4)$$

$x > 0$. If $\delta > 0$ the permissible values for the other parameters are $\gamma \geq 0$ if $\lambda < 0$ and $\gamma > 0$ if $\lambda = 0$. If $\delta \geq 0$ then γ, λ should be strictly positive.

The moments of the GIG can be expressed as functions of the Bessel-K functions by

$$E(X^j) = \left(\frac{\delta}{\gamma}\right)^j \frac{K_{\lambda+j}(\delta\gamma)}{K_\lambda(\delta\gamma)}. \quad (5)$$

The mode of the GIG is a simple function of the distribution’s parameters. More specifically, $Mo(X) = (\gamma^{-2})(\lambda - 1 + \sqrt{(\lambda - 1)^2 + \delta^2 \gamma^2})$ for strictly positive γ , and $Mo(X) = \delta^2/[2(1 - \lambda)]$ for $\gamma = 0$. We note that using theorem 1.2 from Laforgia and

Natalini (2010), and the positive skewness of the GIG distribution (Nguyen et al. 2003), it may easily be shown that for $\gamma > 0$

$$\frac{(\lambda - 1) + \sqrt{(\lambda - 1)^2 + \delta^2 \gamma^2}}{\gamma^2} \leq E(X) \leq \frac{(\lambda + 1) + \sqrt{(\lambda + 1)^2 + \delta^2 \gamma^2}}{\gamma^2}. \tag{6}$$

Many important distributions may be obtained as special cases of the GIG. For $\lambda > 0$ and $\gamma > 0$, the gamma distribution emerges as the limit when $\delta \rightarrow 0$. The inverse-gamma is obtained when $\lambda < 0$, $\delta > 0$ and $\gamma \rightarrow 0$ and an inverse Gaussian distribution is obtained when $\lambda = -\frac{1}{2}$.

Barndorff-Nielsen (1977) introduces the generalized hyperbolic (GH) distribution as a normal variance-mean mixture where the mixing distribution is GIG. That is, if $(X|W = w) \sim N(\mu + \beta w, w)$ and $W \sim GIG(\lambda, \delta, \gamma)$ then the marginal distribution of X will be GH (i.e., $X \sim GH(\lambda, \alpha, \beta, \delta, \mu)$, where $\alpha^2 = \beta^2 + \gamma^2$). The probability density function of the GH is given by

$$f(x) = \frac{\left(\frac{\gamma}{\delta}\right)^\lambda}{\sqrt{2\pi}K_\lambda(\delta\gamma)} \frac{K_{\lambda-1/2}(\alpha\sqrt{\delta^2 + (x-\mu)^2})}{(\sqrt{\delta^2 + (x-\mu)^2}/\alpha)^{1/2-\lambda}} \exp(\beta(x-\mu)), \tag{7}$$

$x \in \Re$, and where $\gamma^2 = \alpha^2 - \beta^2$. The parameter domain is defined by the following conditions: *i*) $\delta \geq 0, \alpha > 0, \alpha^2 > \beta^2$ if $\lambda > 0$; *ii*) $\delta > 0, \alpha > 0, \alpha^2 > \beta^2$ if $\lambda = 0$; *iii*) $\delta > 0, \alpha \geq 0, \alpha^2 \geq \beta^2$ if $\lambda < 0$. The parameter α determines the shape, β determines the skewness (the sign of the skewness is consistent with that of β), μ is a location parameter, δ serves for scaling and λ influences the size of mass contained in the tails. The class of GH distributions is closed under affine transformations i.e if $X \sim GH(\lambda, \alpha, \beta, \delta, \mu)$ and $Z = b_0X + b_1$ then $Z \sim GH(\lambda, \alpha/|b_0|, \beta/|b_0|, |b_0|\delta, b_0\mu + b_1)$. An essential tool in what follows is the moment generating function of the GH distribution:

$$M_{GH}(t) = \exp(\mu t) \left(\frac{\gamma^2}{\alpha^2 - (\beta + t)^2} \right)^{\lambda/2} \frac{K_\lambda(\delta\sqrt{\alpha^2 - (\beta + t)^2})}{K_\lambda(\delta\gamma)} \tag{8}$$

which exists provided that $|\beta + t| < \alpha$.

3 Bayes estimators of $\theta_{a,b}$

3.1 Derivation of the posterior distribution for σ^2 and $\theta_{a,b}$

The representation of the GH distribution as a normal mean-variance mixture with the GIG as mixing distribution introduced in the previous section provides the basis for obtaining the posterior distribution of $\eta_{a,b} = \log(\theta_{a,b})$ when assuming a GIG prior for σ^2 . More specifically we can prove the following result.

Theorem 3.1. Assume the following: *i*) $p(\eta_{a,b}|\sigma^2, X) \sim N(\bar{\eta}_{a,b}, a^2\sigma^2/n)$ with $\bar{\eta}_{a,b} = a\bar{X} + b\sigma^2$ and *ii*) $p(\xi, \sigma^2) = p(\xi)p(\sigma^2)$, with $p(\sigma^2) \sim GIG(\lambda, \delta, \gamma)$ and $p(\xi)$ an improper

distribution uniform over the real line.
 It follows that

$$p(\sigma^2|data) \sim GIG(\bar{\lambda}, \bar{\delta}^*, \gamma), \tag{9}$$

$$p(\eta_{a,b}|data) \sim GH(\bar{\lambda}, \bar{\alpha}, \bar{\beta}, \bar{\delta}, \bar{\mu}) \tag{10}$$

where $\bar{\delta}^* = \sqrt{Y^2 + \delta^2}$, $Y^2 = \sum_{i=1}^n (X_i - \bar{X})^2$, $\bar{\lambda} = \lambda - \frac{n-1}{2}$, $\bar{\alpha} = \sqrt{\frac{n}{a^2}(\gamma^2 + \frac{nb^2}{a^2})}$ and $\bar{\beta} = n\frac{b}{a^2}$. Let $\bar{\gamma}^2 = \bar{\alpha}^2 - \bar{\beta}^2$. As a consequence $\bar{\gamma}^2 = \frac{n}{a^2}\gamma^2$, $\bar{\delta} = \sqrt{\frac{a^2}{n}(Y^2 + \delta^2)}$ and $\bar{\mu} = a\bar{X}$.

Proof. To prove (9) simply note that:

$$\begin{aligned} p(\sigma^2|data) &\propto \int \sigma^{-n} \exp \left\{ -\frac{1}{2\sigma^2}(Y^2 + n(\xi - \bar{X})^2) \right\} \\ &\times \left(\frac{\gamma}{\delta}\right)^\lambda \frac{1}{2K_\lambda(\delta\gamma)} (\sigma^2)^{\lambda-1} \exp \left\{ -\frac{1}{2}(\delta^2\sigma^{-2} + \gamma^2\sigma^2) \right\} d\xi \\ &\propto \int (\sigma^2)^{-n/2+\lambda-1} \exp \left\{ -\frac{1}{2\sigma^2}(Y^2 + \delta^2) - \frac{\gamma^2\sigma^2}{2} \right\} \\ &\times \exp \left\{ -\frac{1}{2\sigma^2}n(\xi - \bar{X})^2 \right\} d\xi \\ &= (\sigma^2)^{\lambda-\frac{n+1}{2}} \exp \left\{ -\frac{1}{2} \left(\frac{Y^2 + \delta^2}{\sigma^2} - \frac{\gamma^2\sigma^2}{2} \right) \right\}. \end{aligned}$$

The second statement is a special case of the [Barndorff-Nielsen \(1977\)](#) result. □

We are not primarily interested in $p(\eta_{a,b}|data)$, but rather in $\theta_{a,b} = \exp(\eta_{a,b})$ which is distributed as a log-GH, a distribution that has not, to our knowledge, received any attention in the literature. In any case, we can calculate the moments of $p(\theta_{a,b}|data)$ that we need for summarizing the posterior distribution with a quadratic loss function by using the moment-generating function of the GH distribution ($M_{GH}(t)$) and more specifically the fact that $E(\theta|data) = M_{GH}(1)$ and $Var(\theta|data) = M_{GH}(2) - [M_{GH}(1)]^2$.

If we are able to generate samples from the GH distribution, moreover, we may obtain a sample from its exponential transformation. The quantiles and probability intervals may then be calculated using Monte Carlo (MC) techniques. From among the variety of software available for generating random GH numbers, we mention the **ghyp** package running under R ([Breymann and Lüthi 2010](#)).

$M_{\eta|data}(t)$ exists only if $|\bar{\beta} + t| < \bar{\alpha}$, or equivalently if $\bar{\alpha}^2 - (\bar{\beta} + t)^2 > 0$ (i.e., $\bar{\gamma}^2 > t^2 + 2n\frac{b}{a^2}t$). This condition implies the following constraint on the prior parameter

γ :

$$\gamma^2 > \frac{a^2}{n}t^2 + 2bt. \quad (11)$$

The existence of posterior moments requires that γ is above a positive threshold when $a \neq 0, b > 0$ (as for the expected value). The threshold is asymptotically 0 for the median, (i.e., $\theta_{0,1}$), and it is negative for the mode (whenever $n > t/2$): so, it does not represent a restriction. With respect to the inference on the expected value ($\theta_{1,0.5}$) note that the popular inverse gamma prior on σ^2 , a special case of the GIG for $\lambda < 0, \delta > 0$ when $\gamma \rightarrow 0$, does not respect condition (11) thereby leading to a posterior distribution with non-existent moments. Note that this result is consistent with the following remark from Zellner (1971) concerning the inference about $\theta_{1,0}$: posterior moments exist only for the limit as $n \rightarrow \infty$ (that is, when the log-t posterior converges to the log-normal). Similarly, the uniform prior over the range $(0, A)$ for σ (Gelman 2006) implies that $p(\sigma^2) \propto \frac{1}{\sigma} \mathbf{1}_{(0,A)}$, which may be seen as an approximation to a $Gamma(\frac{1}{2}, \epsilon)$ (where $\epsilon = (4A^2)^{-1}$) truncated at A^2 . For $\lambda > 0, \gamma > 0$ and $\delta \rightarrow 0$, $GIG(\lambda, \delta, \gamma) \rightarrow Gamma(\lambda, \gamma^2/2)$. If we let $A \rightarrow \infty$, therefore, $p(\sigma) \propto 1$ is equivalent to a GIG prior with $\gamma \rightarrow 0$ and thus implies non-existent posterior moments.

Consistent with intuition, condition (11) implies that, in practice, to obtain a posterior distribution of $\theta_{a,b}$ with finite moments, a prior with short tails should be chosen.

3.2 Bayes estimators under quadratic and relative quadratic losses

If we summarize $p(\theta|data)$ using the ordinary quadratic loss function we obtain $\hat{\theta}_{a,b}^{QB} = E(\theta_{a,b}|data)$ or

$$\begin{aligned} \hat{\theta}_{a,b}^{QB} &= \exp(\bar{\mu}) \left(\frac{\bar{\gamma}^2}{\bar{\alpha}^2 - (\bar{\beta} + 1)^2} \right)^{\bar{\lambda}/2} \\ &\times \frac{K_{\bar{\lambda}}(\bar{\delta} \sqrt{\bar{\alpha}^2 - (\bar{\beta} + 1)^2})}{K_{\bar{\lambda}}(\bar{\delta} \bar{\gamma})} \end{aligned} \quad (12)$$

$$\begin{aligned} &= \exp(a\bar{X}) \left(\frac{\gamma^2}{\gamma^2 - (\frac{a^2}{n} + 2b)} \right)^{(\lambda - \frac{n-1}{2})/2} \\ &\times \frac{K_{\{\lambda - \frac{n-1}{2}\}}(\sqrt{(Y^2 + \delta^2)(\gamma^2 - (\frac{a^2}{n} + 2b))})}{K_{\{\lambda - \frac{n-1}{2}\}}(\sqrt{(Y^2 + \delta^2)\gamma^2})}. \end{aligned} \quad (13)$$

We provided two alternative expressions for $\hat{\theta}_{a,b}^{QB}$: (12) is indexed on the posterior parameters, and (13) highlights the role of the prior parameters, which will be useful for studying the choice of hyperparameters that is discussed in the next section.

Under a relative quadratic loss function, the Bayes estimator is defined as $\hat{\theta}_{a,b}^{RQB} = E(\theta_{a,b}^{-1})/E(\theta_{a,b}^{-2})$ (see Zellner (1971)). The following result shows that $\hat{\theta}_{a,b}^{RQB}$ may be

reconducted to a Bayes predictor under a quadratic loss function with a different choice of b and modified prior parameters.

Theorem 3.2. *For the Bayes estimator under relative quadratic loss function, we have that $\hat{\theta}_{a,b}^{RQB} = \hat{\theta}_{a,b^*}^{QB}$ with $b^* = b - 2a^2/n$ provided that the prior $p(\sigma^2) \sim GIG(\lambda, \delta, \gamma_*)$ with $\gamma_*^2 = \gamma^2 - 4a^2/n + 4b$ is assumed.*

Proof. Let $\tau_{a,b} = -\log(\theta_{a,b})$. From the stated properties of the GIG distribution, we have $\tau_{a,b}|data \sim GH(\bar{\lambda}, \bar{\alpha}, \bar{\beta}, \bar{\delta}, -\bar{\mu})$ and $2\tau_{a,b}|data \sim GH(\bar{\lambda}, \bar{\alpha}/2, \bar{\beta}/2, 2\bar{\delta}, -2\bar{\mu})$. Using (8) it may be shown that

$$\hat{\theta}^{RQB} = \exp(a\bar{X}) \left[\frac{\frac{n}{a^2}(\gamma^2 - \frac{4a^2}{n} + 4b)}{\frac{n}{a^2}(\gamma^2 - \frac{a^2}{n} + 2b)} \right]^{\{\lambda - \frac{n-1}{2}\}/2} \times \frac{K_{\{\lambda - \frac{n-1}{2}\}}\left(\sqrt{(Y^2 + \delta^2)(\gamma^2 - \frac{a^2}{n} + 2b)}\right)}{K_{\{\lambda - \frac{n-1}{2}\}}\left(\sqrt{(Y^2 + \delta^2)(\gamma^2 - \frac{4a^2}{n} + 4b)}\right)}. \tag{14}$$

If we set $b_* = b - 2a^2/n$ and $\gamma_*^2 = \gamma^2 - \frac{4a^2}{n} + 4b$ we obtain a formula that has exactly the structure of (13).

□

4 Choice of hyperparameters

We can easily see that $\hat{\theta}_{a,b}^{QB}$ is sensitive to the choice of the prior parameters; therefore a careful choice of λ, δ, γ is an essential part of the inferential procedure. Following Rukhin (1986), our aim is to choose the hyperparameters to minimize the frequentist MSE of the Bayes estimators. In practice this choice is a complicated task because expression (13) contains a ratio of Bessel-K functions that is quite intractable. Following Rukhin (1986) again, we will use a ‘small argument’ approximation to obtain the MSE-‘optimal’ values of the hyperparameters. Unfortunately this method is viable only for λ and δ because the small value approximation is free of γ . A more ‘heuristic’ argument will be proposed for the latter parameter. The simulation in section 5 shows that the parameters determined in this manner also lead to good estimator performance when the arguments of the Bessel-K functions are no longer small. According to (10), priors with light tails are required to guarantee a finite posterior expectation and expected loss when $b > 0$. For this reason, when a guess of σ^2 is available a priori, it may be used to improve the performances of the Bayes estimators. How this can be done, without breaking down the MSE even when the guess is grossly wrong is described in section 4.2.

4.1 Choice of hyperparameters using the small arguments approximation to the modified Bessel functions of the third kind

Consider first the following approximation of $\hat{\theta}_{a,b}^{QB}$ using the ‘small argument’ approximation of the Bessel-K functions.

Theorem 4.1. *Under the assumptions that $(Y^2 + \delta^2)\gamma^2 < 1$, $(Y^2 + \delta^2)(\gamma^2 - (\frac{a^2}{n} + 2b)) < 1$ and $\lambda < \frac{n-1}{2}$ we have that*

$$\hat{\theta}^{QB} \cong \exp(a\bar{X}) \exp \left\{ - \frac{(Y^2 + \delta^2)(a^2 + 2nb)}{4n(\lambda - \frac{n-3}{2})} \right\} = \hat{\theta}^{qb}. \tag{15}$$

Proof. Now consider the following power series representation of a Bessel function of the first kind:

$$I_\nu(z) = \left(\frac{1}{2}z\right)^\nu \sum_{k=0}^\infty \frac{\left(\frac{1}{4}z^2\right)^k}{k!\Gamma(\nu + k + 1)} \tag{16}$$

$\nu \in \Re$. Note that for $z < 1$, the addends in the sum part of (16) are decreasing, so for small values ($z \rightarrow 0$) the series may be approximated by its first terms ($k = 0, 1, 2, \dots$). Moreover,

$$K_\nu(z) = \frac{\pi}{2 \sin(\nu\pi)} [I_{-\nu}(z) - I_\nu(z)] \tag{17}$$

$\nu \in \Re - Z$. Note that both (16) and (17) are proved in Abramowitz and Stegun (1968), chapter 9. Combining the two and assuming $\nu > 0$ we may write

$$\begin{aligned} [I_{-\nu}(z) - I_\nu(z)] &= \left(\frac{1}{2}z\right)^{-\nu} \left[\frac{1}{\Gamma(1-\nu)} + \frac{\left(\frac{1}{4}z^2\right)}{\Gamma(2-\nu)} + \dots \right. \\ &\quad \left. - \left(\frac{1}{2}z\right)^{2\nu} \frac{1}{\Gamma(\nu+1)} - \left(\frac{1}{2}z\right)^{2\nu+2} \frac{1}{\Gamma(\nu+2)} - \dots \right]. \end{aligned}$$

Note that the terms in the second line are negligible for $z \rightarrow 0$. Assuming $\nu - 1 \cong \nu$ and $(\nu - 1)(\nu - 2) \cong (\nu - 1)^2$

$$K_\nu(z) \cong \frac{\pi}{2 \sin(\nu\pi)} \left(\frac{1}{2}z\right)^{-\nu} \frac{1}{\Gamma(1-\nu)} \exp\left(\frac{1}{4} \frac{z^2}{1-\nu}\right). \tag{18}$$

We may re-write $\hat{\theta}^{QB}$ as

$$\hat{\theta}^{QB} = \exp(a\bar{X}) \left(\frac{\gamma^2 - u}{\gamma^2}\right)^{\nu/2} \frac{K_\nu\left(\sqrt{(Y^2 + \delta^2)(\gamma^2 - u)}\right)}{K_\nu\left(\sqrt{(Y^2 + \delta^2)\gamma^2}\right)},$$

where $u = \frac{a^2 + 2nb}{n}$, $\nu = -\bar{\lambda}$. Note that, as we assume $2\lambda < n - 1$, the case $\nu < 0$ is not relevant. Using the fact that $K_a(\cdot) = K_{-a}(\cdot)$ and replacing z with the arguments of the Bessel-K that appear in $\hat{\theta}^{QB}$ completes the proof. \square

Rukhin (1986) proves that to minimize the frequentist MSE of estimators in the form $\exp(aX)g(Y)$ such as $\hat{\theta}^{qb}$, and $\hat{\theta}^{QB}$,

$$E \left[g(Y) - \exp(c\sigma^2) \right]^2, \tag{19}$$

where $c = b - 3a^2/(2n)$, should be minimized. Unfortunately, minimization of (19) with respect to (λ, δ) does not lead to a unique minimum. The optimum MSE is reached for a set of (λ, δ) pairs that are described by equation (20).

Theorem 4.2. *Under the assumptions of theorem 4.1, the value of λ in (15) that minimizes (19) is given by*

$$\lambda_{opt} = \frac{n-3}{2} - \frac{(n-1)(a^2+2nb)}{4nc} - \frac{(a^2+2nb)\delta^2}{4nc\sigma^2} \tag{20}$$

provided that $b \notin (-\frac{a^2}{2n}, \frac{3a^2}{2n})$. The result holds for any δ in \mathbb{R}^+ .

Proof. Following (19) we should minimize

$$\phi = E \left\{ \exp \left[\frac{a^2+2nb}{2n(\nu-1)}(Y^2+\delta^2) \right] + \exp(2c\sigma^2) - 2 \exp \left[\frac{a^2+2nb}{4n(\nu-1)}(Y^2+\delta^2) + c\sigma^2 \right] \right\}$$

where $\nu = -\bar{\lambda}$.

Because $\frac{Y^2}{\sigma^2} \sim \chi^2(n-1) = \text{Gamma}(\frac{n-1}{2}, \frac{1}{2})$ then $Y^2 \sim \text{Gamma}(\frac{n-1}{2}, \frac{1}{2\sigma^2})$. Using the standard formula for the moment generating function it follows that $E[\exp(rY^2)] = (1-2r\sigma^2)^{-\frac{n-1}{2}}$, provided that $r < \frac{1}{2\sigma^2}$. Thus

$$E \left\{ \exp \left[\frac{a^2+2nb}{2n(\nu-1)}Y^2 \right] \right\} = \left(1 - \frac{a^2+2nb}{n(\nu-1)}\sigma^2 \right)^{-\frac{n-1}{2}} \cong \exp \left(\frac{(n-1)(a^2+2nb)\sigma^2}{2n(\nu-1)} \right)$$

when $\sigma^2 < \frac{n(\nu-1)}{a^2+2nb}$. The approximated function to be minimized may be written as

$$\begin{aligned} \tilde{\phi} &= \exp \left\{ \frac{(a^2+2nb)}{2n(\nu-1)} [(n-1)\sigma^2 + \delta^2] \right\} \\ &- 2 \exp \left\{ \frac{(a^2+2nb)}{4n(\nu-1)} [(n-1)\sigma^2 + \delta^2] + c\sigma^2 \right\} + \exp(2c\sigma^2). \end{aligned}$$

Taking partial derivatives of $\tilde{\phi}$, and equating both partial derivatives to 0 leads to exactly the same equation:

$$\exp \left\{ \frac{(a^2+2nb)}{2n(\nu-1)} [(n-1)\sigma^2 + \delta^2] \right\} = \exp \left\{ \frac{(a^2+2nb)}{4n(\nu-1)} [(n-1)\sigma^2 + \delta^2] + c\sigma^2 \right\}.$$

Taking the log and solving for ν , we obtain

$$\nu = 1 + \frac{(n - 1)(a^2 + 2nb)}{4nc} + \frac{(a^2 + 2nb)}{4nc} \frac{\delta^2}{\sigma^2}. \tag{21}$$

So, we may obtain an ‘optimal’ value of ν for any choice of (positive) δ^2 . Formula (20) follows from noting that $\nu = \frac{n-1}{2} - \lambda$. \square

Note that if $b \in (-\frac{a^2}{2n}, \frac{3a^2}{2n})$ we would obtain $\lambda_{opt} > \frac{n-1}{2}$ for which the approximation (15) on which this choice of λ is based is not valid anymore. In practice, the median ($a = 1, b = 0$) is the only relevant case that falls outside the applicability of theorem 4.2. To choose λ when $b = 0$ we may observe that the estimators obtained under quadratic loss, (13), and relative quadratic loss, (14), are very close and asymptotically equivalent. Since $b^* \notin (-\frac{a^2}{2n}, \frac{3a^2}{2n})$ we may then apply (20) replacing $b^* = b - \frac{2a^2}{n}$ instead of b .

The λ_{opt} in (20) is a function not only of δ , as anticipated, but also of the unknown σ^2 ; therefore, an optimal choice of (δ, λ) should depend, at least in principle, on a prior guess for σ^2 . A method for circumventing the problem that is implicitly suggested by the generalized prior (3) is to let $\delta \rightarrow 0$. This condition may be approximated in practice by a δ that is much smaller than σ so as to make the third addend in (20) negligible. This approximation can be justified by noting (from (6)) that δ has the same order of magnitude of the expectation and the mode of σ^2 ; therefore, choices of the type $\delta = k\sigma_0^2$ for some constant k and prior guess of the variance σ_0^2 imply a negligible third addend in (20) in the ‘small value setting’ we assumed for the derivation of $\hat{\theta}^{qb}$.

The estimator $\hat{\theta}_{1,0.5}^{qb}$ is connected to popular estimators that have already been discussed in the literature. Note that if $\delta = 0, a = 1, b = 0.5$ and we replace λ_{opt} into (3.1), we obtain $\hat{\theta}_{1,0.5}^{qb} = \exp((\bar{X}) + \frac{S^2(n-3)}{2n})$, which is the MSE-optimal estimator (3.9) proposed by Zellner (1971) with the assumed known σ^2 replaced by $S^2 = Y^2/(n - 1)$. This estimator is also similar to the one proposed in Evans and Shaban (1976) with the function g truncated to the first term.

As far as γ is concerned, we propose choosing a value close to the minimum value to assure the existence of the first two posterior moments. Therefore, we specify the GIG with the heaviest possible tail among those yielding $p(\theta|data)$ with finite variance:

$$\gamma_0^2 = \max \left\{ 0, 4 \left(\frac{a^2}{n} + b \right) \right\} + \epsilon \tag{22}$$

where ϵ is a positive, typically small constant. Note that γ_0 depends on n . In any case, we found in our simulations that $\hat{\theta}_{1,0.5}^{QB}$ is not particularly sensitive to alternative choices of γ_0 that are close to (i.e., of the same order of magnitude of) the γ_0 we propose. Much larger values lead to inefficient $\hat{\theta}_{1,0.5}^{QB}$ with far larger frequentist MSEs. Moreover, the resulting GIG distribution will have a particularly light right tail for positive b . This result implies that the prior will be relatively peaked. In figure 1 we plot the log-density of the GIG prior for the discussed choice of the hyperparameters (and various n) and

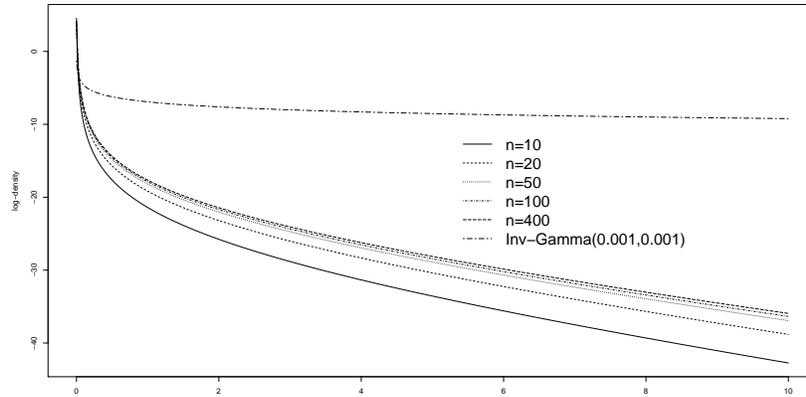


Figure 1: Log density of the GIG prior for the discussed choice of hyperparameters (and various n) and of a $Inv - Gamma(0.001, 0.001)$.

compare it to that of a $Inv - Gamma(0.001, 0.001)$, a commonly used prior for variance parameters.

Because a choice of δ close to 0 implies a peak close to 0, in the next section we explore how a prior guess on the population variance may be used to improve the performances of the Bayes estimators.

4.2 Choice of hyperparameters based on prior guesses of σ^2

Assume that a prior guess σ_0^2 for σ^2 is available and this value is ‘far’ from 0. The subject of this section is how to specify priors that incorporate this prior guess and preserve the optimality properties we have described. We restrict our attention to positive b , because when $b \leq 0$, the existence of the first two posterior moments does not necessarily imply a prior with light tails, and we may choose priors with a more distributed probability mass.

We start by introducing the following approximation of $E(\sigma^2)$ that is based on the inequality (6) and that has proved in simulations to be effective for a wide range of choices for λ, δ, γ :

$$E(\sigma^2) \cong \frac{\lambda + \sqrt{\lambda^2 + \delta^2 \gamma^2}}{\gamma^2}. \tag{23}$$

In the following we denote the approximate expected value from (23) by σ_E^2 . We then introduce the following result:

Theorem 4.3. *Assuming that i) λ is a function of δ as expressed in (20) and that ii) $b > \frac{3a^2}{2n}$, we have $\sigma_E^2 = k\sigma^2$ for $0 < k < 1$.*

Proof. We may rewrite (20) as $\lambda = c_1 - c_2\delta^2$ with $c_1 = -1 - (n-1)\frac{2a^2}{2nb-3a^2}$ and $c_2 = \frac{a^2+2nb}{2(2nb-3a^2)\sigma^2}$. Substituting this identity into (23) and solving for δ^2 , we obtain:

$$\delta^2 = \frac{k\sigma^2(k\gamma^2\sigma^2 - 2c_1)}{1 - 2kc_2\sigma^2}. \quad (24)$$

Of course, δ^2 must be positive. The signs of both c_1 and c_2 depend on that of c (defined in (2)). Because $2nb - 3a^2 > 0$ (i.e., $b > \frac{3a^2}{2n}$) implies that $c_1 < 0$ and $c_2 > 0$, the numerator of δ^2 is always positive and we may focus on the denominator. It is positive whenever

$$k < \frac{2nb - 3a^2}{2nb + a^2} < 1, \quad (25)$$

a condition that reduces to $k < \frac{n-3}{n+1}$ for $a = 1, b = 0.5$. \square

The main implication of this result is that the relationship between λ and δ implied by (20) leads naturally to conservative priors whose expected values cannot exceed σ^2 . This finding is consistent with all the theory of MSE-optimal estimation of log-normal parameters, in which efficiency is improved at the price of some negative bias. As a byproduct, (24) suggests how to choose δ^2 as a function of γ^2 , a prior guess σ_0^2 and k .

To clarify the interpretation of the parameter k , we study its relation to the coefficient of variation of the prior distribution. Note first that if $W \sim GIG(\lambda, \delta, \gamma)$, with $\lambda < -1$, then

$$CV^2(W) \cong -\frac{1}{\lambda + 1}. \quad (26)$$

To see how this result follows, note that

$$CV^2(W) + 1 = \frac{K_{\lambda+2}(\delta\gamma)}{K_{\lambda+1}(\delta\gamma)} \frac{K_{\lambda}(\delta\gamma)}{K_{\lambda+1}(\delta\gamma)} = \frac{E(W_1)}{E(W_2)},$$

with $W_1 \sim GIG(\lambda + 1, \delta, \gamma)$ and $W_2 \sim GIG(\lambda, \delta, \gamma)$. Using the approximation (23) for the expected values, we find that

$$CV^2(W) + 1 = \frac{(\lambda + 1) + \sqrt{(\lambda + 1)^2 + \delta^2\gamma^2}}{\lambda + \sqrt{\lambda^2 + \delta^2\gamma^2}}.$$

Because $\sqrt{m_1^2 + m_2} \cong |m_1| + \frac{m_2}{2|m_1|}$ for a positive m_2 , $CV^2(W) + 1 \cong \lambda(\lambda + 1)^{-1}$ for $\lambda < -1$, from which (26) follows. We can now state the following result.

Theorem 4.4. *Given assumptions i) and ii) of theorem 4.3, assuming that δ is selected according to (24) and using the approximation in (26) for the squared coefficient of variation, it follows that $CV^2(\sigma^2)$ is a monotonically decreasing function of k .*

Proof. Because $b > \frac{3a^2}{2n}$ implies that $c_1 + 1 < 0$ and $c_2 > 0$, $\lambda = c_1 - c_2\delta^2$ implies that $\lambda + 1 < 0$. It follows from (26) that the squared coefficient of variation will be

monotonically decreasing in k whenever $-(\lambda + 1)$ is monotonically increasing (that is, when its derivative is positive) $\forall k \in (0, \frac{2nb-3a^2}{2nb+a^2})$.

To determine the sign of $\frac{d}{dk}\{c_2\delta^2 - (c_1 + 1)\}$, we focus on the sign of the numerator in $\frac{d}{dk} \frac{c_2k\sigma^2(k\gamma^2\sigma^2-2c_1)}{1-2kc_2\sigma^2}$.

$$(2c_2k\gamma^2\sigma^4 - 2c_1c_2\sigma^2)(1 - 2kc_2\sigma^2) + (c_2k^2\gamma^2\sigma^4 - 2c_1c_2k\sigma^2)2c_2\sigma^2 > 0$$

may be simplified to

$$2c_2\sigma^2(k\gamma^2\sigma^2 - c_2k^2\gamma^2\sigma^4 - c_1) > 0.$$

This inequality holds whenever $k\gamma^2\sigma^2(1 - c_2k\sigma^2) > c_1$. Given that $c_1 < 0$, this condition holds $\forall k \in (0, \frac{2nb-3a^2}{2nb+a^2})$ because for k within this range, we know that $1 - 2c_2k\sigma^2 > 0$ implies that $1 - c_2k\sigma^2 > 0$. □

When selecting a value of k within its permissible range, we should consider that the larger the value of k , the closer the expected value of the prior is to the guess σ_0^2 , but also the smaller the prior’s coefficient of variation (that is the more informative is the distribution). Likewise, a larger a priori coefficient of variation yields an expected value farther to the left of σ_0^2 .

The parameter k may be interpreted as a measure of the amount of a priori information available. If we are quite uncertain about our prior guess for σ^2 it makes sense to have not only a prior with a large coefficient of variation but also a conservatively small expected value. In fact, it can be easily shown through simulation that prior specification that puts a sizeable part of the prior mass beyond σ^2 yields Bayes estimators with huge frequentist MSEs. Therefore, the relationship (through k) between a prior expected value and the coefficient of variation is consistent with the optimization of frequentist MSE of the Bayes estimators.

We also note that the relationship between k and the squared CV is nonlinear. The speed of the squared CV reduction increases with k ; therefore, choices of k close to its maximum imply very peaked prior distributions. On the other extreme, a small k leads to a prior that is practically the same as those considered in section 4.1.

5 Simulation

In this section, we compare, using their frequentist MSEs, the Bayes estimators introduced in the previous sections with other estimators, both frequentist and Bayesian. We focus on estimating the mean of the log-normal (i.e., $\theta_{1,0.5}$), a popular problem in the literature and one for which several competing estimators exist. Specifically, we consider the unbiased estimator of Evans and Shaban (1976) and the estimator proposed by Shen et al. (2006) which the authors proved to be more efficient than the various alternatives that had previously been proposed in the literature. We also consider the Bayes estimator of Rukhin (1986). As a general benchmark, we also consider the MSE-optimal estimator of Zellner (1971) which is based on a known σ^2 .

Table 1: Comparison of alternative estimators in terms of MSE: prior for σ^2 not incorporating guesses.

n	Est./ σ^2	MSE					
		0.1	0.5	1	2	5	20
n=11	$\hat{\theta}^{BQ}$	0.0095	0.0514	0.109	0.228	0.531	0.982
	$\hat{\theta}^{BR}$	0.0096	0.0540	0.120	0.268	0.661	0.998
	$\hat{\theta}^{Ru}$	0.0095	0.0513	0.110	0.244	0.680	1.291
	$\hat{\theta}^{SBZ}$	0.0095	0.0505	0.105	0.217	0.544	0.998
	$\hat{\theta}^{ES}$	0.0098	0.0594	0.147	0.422	2.533	65.77
	$\hat{\theta}^{Zel}$	0.0093	0.0458	0.090	0.173	0.384	0.944
n=101	$\hat{\theta}^{BQ}$	0.0010	0.0060	0.014	0.036	0.133	0.729
	$\hat{\theta}^{BR}$	0.0010	0.0060	0.014	0.038	0.170	0.929
	$\hat{\theta}^{Ru}$	0.0011	0.0067	0.017	0.052	0.308	18.970
	$\hat{\theta}^{SBZ}$	0.0010	0.0060	0.014	0.036	0.134	0.805
	$\hat{\theta}^{ES}$	0.0010	0.0061	0.015	0.039	0.183	4.562
	$\hat{\theta}^{Zel}$	0.0010	0.0049	0.010	0.020	0.486	0.182
n=401	$\hat{\theta}^{BQ}$	0.0003	0.0016	0.004	0.010	0.041	0.375
	$\hat{\theta}^{BR}$	0.0003	0.0016	0.004	0.010	0.045	0.566
	$\hat{\theta}^{Ru}$					0.162	9.755
	$\hat{\theta}^{SBZ}$	0.0003	0.0016	0.004	0.010	0.041	0.376
	$\hat{\theta}^{ES}$	0.0003	0.0016	0.004	0.010	0.045	0.780
	$\hat{\theta}^{Zel}$	0.0003	0.0013	0.003	0.005	0.012	0.049

The simulation set-up is essentially the same as that introduced in [Zhou and Gao \(1997\)](#) and also used by [Shen et al. \(2006\)](#). We assume that $X \sim N(-\sigma^2/2, \sigma^2)$; then $\exp(X)$ is log-normally distributed, and our estimand $E[\exp(X)] = 1$. We consider 6 distinct values for σ^2 ($\sigma^2 = 0.1, 0.5, 1, 2, 5, 20$) and three sample sizes ($n = 11, 101, 400$). The results we present used 100,000 MC samples, and were obtained using R.

In [tables 1 and 2](#), we present the results obtained under the priors specified according to the suggestions of [Section 4.1](#); specifically we set $\delta = 0.01$. In presenting the results, we denote the estimator of formula (3.1) in [Rukhin \(1986\)](#) by $\hat{\theta}^{Ru}$, with the prior parameters chosen using formulas (3.6) and (3.7) of the same paper. Let $\hat{\theta}^{ES}$ denote the unbiased estimator of [Evans and Shaban \(1976\)](#) and let $\hat{\theta}^{SBZ}$ denote the estimator from [Shen et al. \(2006\)](#). We also denote the optimal estimator assuming a known σ^2 that was discussed above by $\hat{\theta}^{Zel}$. Because the simulations only address estimating the mean, we omit the indexes related to the choice of a, b .

Despite the prior parameters having been chosen using a ‘small arguments’ approximation to the Bessel-K function, results show that $\hat{\theta}^{QB}$ performs well regardless of the size of σ^2 . Specifically $\hat{\theta}^{QB}$ is close to $\hat{\theta}^{SBZ}$ for all values of σ^2, n . We emphasize this as $\hat{\theta}^{SBZ}$ is the reference frequentist estimator in the recent literature.

Both estimators are negatively biased (see table 2) and the bias rapidly increases with σ^2 , especially when $n = 11$, even though $\hat{\theta}^{QB}$ is less biased and more variable with respect to $\hat{\theta}^{SBZ}$. In any case, note that $\hat{\theta}^{Zel}$, which is MSE-optimal and assumes known σ^2 , is also negatively biased. From Theorem 1 in Shen et al. (2006), $\hat{\theta}^{SBZ}$ is obtained from

$$\hat{\theta}^{SBZ}(\sigma^2) = \exp \left\{ \bar{X} + \frac{Y^2(n-1)}{2(n+4)(n-1) + 3(n-1)\sigma^2} \right\}$$

in which σ^2 in the denominator is replaced by its Maximum Likelihood (ML) estimator $Y^2/(n-1)$. The approximation of $\hat{\theta}^{QB}$ for small values of σ^2 , obtained from (15) by replacing λ with (20) and setting $\delta = \sigma^2$ (and assuming σ^4 negligible) may be written as

$$\hat{\theta}^{qb}(\sigma^2) = \exp \left\{ \bar{X} + \frac{Y^2(n-3)}{2n(n-4) + 2n\sigma^2} \right\}$$

which is close to $\hat{\theta}^{SBZ}$.

$\hat{\theta}^{QB}$ is also more efficient than $\hat{\theta}^{ES}$. For small σ^2 , $\hat{\theta}^{QB}$ is only moderately biased and shows a similar MSE, while for large population variances, the unbiasedness of $\hat{\theta}^{ES}$ is quite costly in terms of variance.

When comparing $\hat{\theta}^{QB}$ to $\hat{\theta}^{RQB}$, we find that the two perform similarly for small σ^2 , but $\hat{\theta}^{QB}$ becomes clearly superior to the predictor based on the relative quadratic loss function as σ^2 increases. $\hat{\theta}^{Ru}$ exhibits an MSE close to that of $\hat{\theta}^{QB}$ for small σ^2 , but its properties deteriorate dramatically for large population variances. We have already noted that the prior chosen for σ^2 in Rukhin (1986) is the limit of a GIG for $\delta \rightarrow 0$; nonetheless the proposed choice for γ is inconsistent with the existence of the first two moments of the posterior distribution for the mean of the log-normal. The results for $\hat{\theta}^{Ru}$ when $n = 400$ and $\sigma^2 \leq 2$ are missing because of numerical problems. These estimators involve the calculation of Bessel-K functions with very large orders (the order increases linearly with n) and arguments very close to 0, which results in huge values. These values lead to numerical instability and the generation of errors with the software we used. Because they are not essential to our purposes, we do not investigate the problem further.

In table 3, we report selected results for the estimators of the log-normal mean in the case where the priors for σ^2 incorporate prior guesses using the methodology described in Section 4.2. These priors are relevant only for large σ^2 ; therefore, we report the results only for $\sigma^2 = 1, 5, 20$. They involve a parameter k that controls both the closeness of the prior expectation to σ_0^2 (the prior guess) and the peakedness of the prior. Let M denote the maximum admissible value for k according to (25); we consider the cases where $k = 0.5M$ and $k = M$. The purpose of this second set of comparisons involving $\hat{\theta}^{QB}$ is mainly to determine how much efficiency improves when a guess of the population variance is available and used in the prior specification and how sensitive these gains are to incorrect guesses for σ^2 . For this purpose, we consider the following values for σ_0^2 : $0.4\sigma^2, 0.8\sigma^2, \sigma^2, 1.2\sigma^2, 1.6\sigma^2, 2\sigma^2$. These values correspond to errors in guessing σ^2 that

Table 2: Comparison of alternative estimators in terms of Bias: prior for σ^2 not incorporating guesses.

n	$Est./\sigma^2$	Bias					
		0.1	0.5	1	2	5	20
n=11	$\hat{\theta}^{BQ}$	-0.0075	-0.0470	-0.103	-0.225	-0.555	-0.990
	$\hat{\theta}^{BR}$	-0.0254	-0.1278	-0.246	-0.445	-0.793	-0.999
	$\hat{\theta}^{Ru}$	-0.0102	-0.0540	-0.107	-0.207	-0.448	-0.921
	$\hat{\theta}^{SBZ}$	-0.0108	-0.0621	-0.134	-0.288	-0.680	-0.999
	$\hat{\theta}^{ES}$	0.0010	0.0015	0.002	0.004	0.004	0.007
	$\hat{\theta}^{Zel}$	-0.0075	-0.0408	-0.082	-0.159	-0.354	-0.827
n=101	$\hat{\theta}^{BQ}$	-0.0014	-0.0067	-0.015	-0.039	-0.154	-0.812
	$\hat{\theta}^{BR}$	-0.0034	-0.0186	-0.042	-0.104	-0.341	-0.963
	$\hat{\theta}^{Ru}$	0.0029	0.0160	0.033	0.067	0.176	0.818
	$\hat{\theta}^{SBZ}$	-0.0023	-0.0111	-0.024	-0.055	-0.191	-0.888
	$\hat{\theta}^{ES}$	-0.0004	-0.0005	0.000	0.000	0.001	-0.014
	$\hat{\theta}^{Zel}$	-0.0015	-0.0061	-0.011	-0.022	-0.052	-0.185
n=401	$\hat{\theta}^{BQ}$	0.000	-0.002	-0.004	-0.010	-0.044	-0.415
	$\hat{\theta}^{BR}$	-0.001	-0.005	-0.011	-0.029	-0.118	-0.732
	$\hat{\theta}^{Ru}$					0.278	1.661
	$\hat{\theta}^{SBZ}$	-0.0006	-0.0029	-0.006	-0.015	-0.054	-0.447
	$\hat{\theta}^{ES}$	-0.0001	-0.0002	0.000	0.000	0.000	0.002
	$\hat{\theta}^{Zel}$	-0.0004	-0.0015	-0.003	-0.006	-0.013	-0.050

range from large and negative to large and positive. The results are reported only for $n = 11, 101$ because the impact of the prior specification is less interesting for $n = 400$.

With respect to table 3, we note that, as expected, when the guess σ_0^2 is exactly equal to σ^2 there are large gains in efficiency relative to the parallel results in table 1. The gains are larger for $k = M$ than for $k = 0.5M$. To better appreciate the size of these gains, the MSEs and biases should also be compared with those of $\hat{\theta}^{Zel}$ from table 1, whose MSEs represent a lower bound.

When $k = 0.5M$, there is much less sensitivity to wrong guesses for σ^2 than when $k = M$. As anticipated in Section 4.2 this finding is explained by the prior becoming peaked and the prior coefficient of variation decreasing dramatically when k gets close to its upper bound. We observe this decreased sensitivity for both under- and over-statement of σ^2 . Specifically, when $\sigma_0^2 = 0.4\sigma^2$, the prior implied by $k = 0.5M$ is more diffuse, which compensates for its expectation being farther to the right than in the case of $k = M$. For values of the guess close to σ^2 , a more peaked distribution leads to estimators that behave better. However, as soon as the prior expected value exceeds the underlying true value (when $\sigma_0^2 = 1.2\sigma^2$, for example), we start to observe an increase in the MSE and a switch in the sign of the bias (from negative to positive) for the estimators associated with $k = M$. For larger positive errors in guessing σ^2 , the properties of these estimators deteriorate fast and dramatically.

For $k = 0.5M$, we observe a smoother behavior. The estimators with the best MSE are not those where $\sigma_0^2 = \sigma^2$. This result is due to $E(\sigma^2) = k\sigma_0^2$ which implies that a value of σ^2 moderately greater than σ^2 yields a prior with an expected value closer to σ^2 . It is also unwise in this situation to intentionally overstate σ^2 . From the case of $\sigma_0^2 = 2\sigma^2$, we may observe that the properties of $\hat{\theta}^{QB}$ when $k = 0.5M$ also deteriorate when there is gross positive error in guessing the population variance.

We conclude that when a reasonable prior guess for σ^2 is available, the efficiency of the $\hat{\theta}^{QB}$ Bayes estimators may be improved. Except for the case of k close to the upper limit of its admissible range, these improvements are substantial and can be obtained with a reasonable level of robustness with respect to guessing errors. Concerning the choice of k , we presented empirical results for $k = 0.5M$ that represent the best performance for the simulation setting considered here and also for others we do not report for the sake of brevity. A smaller k tends to reproduce the results we have seen for the choices of prior parameters discussed in Section 4.1 while a larger k shares, although to a lesser degree, the problems illustrated for $k = M$.

6 An application to a real environmental problem

The statistical assessment of data collected from contaminated sites is a major task for environmental agencies. Contaminant concentration data quite often appear to follow a positively skewed distribution and the log-normal is frequently used as a reference model. In the US, the Environmental Protection Agency (EPA) recommends basing remediation decisions on the estimation of the upper confidence limit (UCL) for the

Table 3: Comparison of alternative estimators: prior incorporating guesses of σ^2 .

		<i>MSE</i>			<i>Bias</i>		
		$\sigma^2 = 1$	$\sigma^2 = 5$	$\sigma^2 = 20$	$\sigma^2 = 1$	$\sigma^2 = 5$	$\sigma^2 = 20$
		$\sigma_0^2 = 0.4\sigma^2$					
$k = .5M$	11	0.101	0.521	0.988	-0.157	-0.638	-0.994
$k = M$	11	0.110	0.593	0.995	-0.227	-0.743	-0.997
$k = .5M$	101	0.014	0.137	0.792	-0.031	-0.226	-0.873
$k = M$	101	0.023	0.309	0.973	-0.117	-0.540	-0.986
		$\sigma_0^2 = 0.8\sigma^2$					
$k = .5M$	11	0.097	0.474	0.977	-0.121	-0.564	-0.987
$k = M$	11	0.091	0.432	0.952	-0.140	-0.564	-0.974
$k = .5M$	101	0.014	0.127	0.736	-0.021	-0.181	-0.830
$k = M$	101	0.013	0.131	0.772	-0.052	-0.310	-0.876
		$\sigma_0^2 = \sigma^2$					
$k = .5M$	11	0.096	0.453	0.963	-0.101	-0.515	-0.979
$k = M$	11	0.088	0.377	0.875	-0.092	-0.421	-0.913
$k = .5M$	101	0.014	0.123	0.680	-0.015	-0.149	-0.779
$k = M$	101	0.011	0.067	0.324	-0.013	-0.106	-0.498
		$\sigma_0^2 = 1.2\sigma^2$					
$k = .5M$	11	0.096	0.438	0.940	-0.080	-0.455	-0.964
$k = M$	11	0.091	0.406	0.986	-0.040	-0.227	-0.701
$k = .5M$	101	0.014	0.120	0.609	-0.009	-0.109	-0.693
$k = M$	101	0.013	0.131	2.934	0.029	0.189	1.238
		$\sigma_0^2 = 1.6\sigma^2$					
$k = .5M$	11	0.099	0.457	0.876	-0.036	-0.296	-0.879
$k = M$	11	0.117	1.316	80.534	0.074	0.401	2.648
$k = .5M$	101	0.014	0.132	0.683	0.005	-0.004	-0.281
$k = M$	101	0.030	1.845	3483.366	0.124	1.236	52.430
		$\sigma_0^2 = 2\sigma^2$					
$k = .5M$	11	0.108	0.627	1.666	0.012	-0.069	-0.539
$k = M$	11	0.181	6.355	13668.7	0.204	1.573	44.9
$k = .5M$	101	0.015	0.188	5.421	0.019	0.141	1.067
$k = M$	101	0.074	13.207	2684534	0.238	3.463	1468.1

mean of a log-normal population. The reference method for this estimation is based on Land’s H statistic (EPA 2002), although it has been criticized as it is liable to lead to overestimation of the actual UCL. The issue is discussed in Singh et al. (2002), that also provide real data sets that we use to apply our method.

In this section, for two data sets considered in Singh et al. (2002), we calculate estimates of the log-normal mean using all estimators considered in section 5 along with the maximum likelihood estimator and $\hat{\theta}^{qb}$ introduced in (15); $\hat{\theta}^{Zel}$ is not considered as σ^2 is unknown. Moreover we consider the problem of estimating the UCL at the 0.95 confidence level. We propose to use the 95-th percentile of the posterior $p(\theta_{1,0.5}|data)$ as we found by simulation that the posterior tracks the sampling distribution of $\hat{\theta}_{1,0.5}^{QB}$ very well. Unfortunately, for the choice of hyperparameters discussed in section 4.1, the frequentist coverage is less than 0.95; the reason is that the MSE-optimal $\hat{\theta}_{1,0.5}^{QB}$ is a negatively biased estimator. So, for the purpose of the UCL estimator we introduce an alternative choice of the prior parameter λ that leads to a UCL that guarantees a frequentist coverage close to 0.95.

For the new choice of λ we proceed as follows. As we base the interval estimator on the percentile of the posterior distribution, we can consider $p(\eta_{1,0.5}|data)$ instead of $p(\theta_{1,0.5}|data)$; we choose hyperparameters in order to have $E_s(\hat{\eta}_{1,0.5}^{QB}) \cong \xi + \frac{1}{2}\sigma^2$ (with E_s denoting the expectation with respect to the sampling distribution and $\hat{\eta}^{QB} = E(\eta_{1,0.5}|data)$); $\exp\{q_{0.95}(\eta_{1,0.5}|data)\}$ will then be our estimate of the UCL at the 0.95 level. Using theorem 3.1 and approximation (23) we have that

$$\hat{\eta}_{1,0.5}^{QB} = E(\eta_{1,0.5}|data) \cong \bar{\mu} + \bar{\beta} \left\{ \frac{\bar{\lambda} + \sqrt{\bar{\lambda}^2 + \bar{\delta}^2 \bar{\gamma}^2}}{\bar{\gamma}^2} \right\}.$$

Because $\sqrt{m_1^2 + m_2} \cong |m_1| + \frac{m_2}{2|m_1|}$ for a positive m_2 , assuming $\bar{\lambda} < 0$ we obtain:

$$\hat{\eta}_{1,0.5}^{QB} \cong \bar{X} + \frac{Y^2 + \delta^2}{2(n - 1 - 2\lambda)}.$$

If we assume δ^2 positive but very small as suggested in section 4.1 we have that for $\lambda = 0$, $E_s(\hat{\eta}_{1,0.5}^{QB}) \cong \eta_{1,0.5}^{QB}$. It may be shown using the same simulation of section 5 that for this new choice of λ , our estimator of the UCL yields a frequentist coverage of the associated interval very close to 0.95 whenever the coefficient of variation of the underlying log-normal population is lower or equal than 2.5, as will be the case for the data sets considered in this section.

We compare our estimate of the UCL to the method based on the H statistic, three others recommended by Singh et al. (2002) (jackknife, standard bootstrap, Chebychev inequality) and the upper limit of the bootstrap based confidence interval associated with $\hat{\theta}^{SBZ}$ (see Shen et al. (2006)).

The data set we consider is the concentration of aluminium and manganese at the Naval Construction Battalion Center (example 4.6). For both contaminants we have two samples of size $n = 17$, with log-normality not rejected using a Shapiro-Wilks test

Table 4: Estimates for the log-normal means: Naval Construction Battalion Center data.

	$\hat{\theta}^{ML}$	$\hat{\theta}^{ES}$	$\hat{\theta}^{SBZ}$	$\hat{\theta}^{Ru}$	$\hat{\theta}^{RQB}$	$\hat{\theta}^{qb}$	$\hat{\theta}^{QB}$
a1	2002.7	1704.8	1214.6	1449.9	1476.9	1570.8	1339.7
mn	1262.6	1100.9	819.4	951.6	1004.1	1016.5	892.1

Table 5: Estimates for the UCL (95% confidence): Naval Construction Battalion Center data.

	jackknife	bootstrap	Chebychev	H -Stat	SBZ	QB
a1	3283.34	3663.20	5314.99	9102.73	6286.47	3757.57
mn	1889.52	1821.55	3291.95	5176.16	3671.14	2375.48

at the 0.1 confidence level. The following data are observed: $\mathbf{a1}=(71, 71, 107, 113, 125, 163, 164, 264, 290, 527, 586, 979, 2640, 2660, 3560, 5920, 13200)$ and $\mathbf{mn}=(15.8, 28.2, 85.6, 90.6, 150.0, 199.0, 259.0, 281.0, 390.0, 777.0, 824.0, 838.0, 1010.0, 1350.0, 1490.0, 3250.0, 4300.0)$. Note that the aluminium data are more skewed than the manganese data ($CV(\mathbf{a1}) = 1.81$, $CV(\mathbf{mn}) = 1.32$).

Point estimates of the log-normal means are reported in table 4. The various estimators applied to the same data lead to different estimated values, with differences increasing with σ^2 and thus with the skewness of the distribution. Note that $\hat{\theta}^{ML}$ returns the largest estimate in both cases, while $\hat{\theta}^{SBZ}$ the smallest. The fact that $\hat{\theta}^{SBZ} < \hat{\theta}^{QB}$ is consistent with the larger negative bias of the first we noted in the previous section. Less efficient estimators tend to produce larger estimates.

Estimates of the UCL (95% confidence level) are reported in table 5. The estimates of the 95-th percentile of $p(\theta_{1,0.5}|data)$ are based on the `ghyp` package of R (Breyman and Lüthi 2010). As expected, we note that the estimates based on the H statistic and the bootstrap interval for $\hat{\theta}^{SBZ}$ described in Shen et al. (2006) produce very large estimates: between the second largest and the largest observed value in the case of $\mathbf{a1}$ and larger than the observed maximum for the case of \mathbf{mn} . Estimates based on our method are smaller, close to those based on resampling methods for the $\mathbf{a1}$ sample, between those and the method based on Chebychev inequality for the \mathbf{mn} sample.

7 Conclusions and future work

In this paper, we considered the popular log-normal model and the specific problems associated with estimating many of its parameters (including mean, median and mode). These problems are caused by the fact that log-t and other distributions that are typically associated with the analysis of the log-normal model have no finite moments. Our approach is Bayesian but parallel problems arise from a frequentist perspective. Specifically we wanted to continue using the popular quadratic loss function to summarize the posterior distribution. We found that a generalized inverse gamma prior for

the population variance allows formally stating the conditions on the prior parameters that lead to posterior distributions with finite moments; moreover, the Bayes estimators of log-normal parameters associated with the quadratic loss function have desirable frequentist properties.

Further developments of this research are possible in many directions. We are interested in applying generalized inverse gamma priors to the variance components in normal mixed models specified for the log of Poisson means that are commonly used (for example) in epidemiology. An extension of this methodology may also be applied to prediction problems in finite population modeling where normal models on the log-scale are popular.

Acknowledgments

The authors would like to thank Professor Partha Lahiri of the University of Maryland for the informal discussion we had that was to lead to this present study. The work of Carlo Trivisano was partly funded by the Grant no. 2008CEFF37001 (PRIN,MIUR)

References

- Abramowitz, A. and Stegun, I. (1968). *Handbook of mathematical functions*. Dover, NY: Addison–Wesley, 5th edition edition. 976, 982
- Barndorff-Nielsen, O. (1977). “Exponentially decreasing distributions for the logarithm of particle size.” *Proceedings of the Royal Statistical Society*, A353: 401–419. 978, 979
- Bibby, B. and M., S. (2003). “Generalized hyperbolic and inverse Gaussian distributions: limiting cases and approximation of processes.” In Rachev, S. (ed.), *Handbook of Heavy Tailed Distributions in Finance*, 211–248. New York, NY: Elsevier Science B.V. 977
- Breymann, W. and Lüthi, D. (2010). *ghyp: a package on Generalized Hyperbolic distributions*. available at cran.r-project.org/web/packages/ghyp/index.html. 979, 994
- Dagne, G. (2001). “Bayesian transformed models for small area estimation.” *TEST*, 10: 375–391. 975
- Eberlein, E. and von Hammerstein, E. (2004). “Hyperbolic processes in finance.” In Dalang, R., Dozzi, M., and F., R. (eds.), *Seminar on Stochastic Analysis, Random Fields and Applications IV, Progress in Probability*, 221–264. Berlin, DE: Birkhäuser Verlag. 977
- EPA (2002). *Limits for exposure point concentrations at hazardous waste sites*. Washington, DC: OSWER 9285.6-10. 993
- Evans, I. and Shaban, S. (1976). “New estimators (of smaller MSE) for parameters of a lognormal distribution.” *Biometrische Zeitschrift*, 18: 453–466. 984, 987, 988

- Gelman, A. (2006). “Prior distributions for variance parameters in hierarchical models.” *Bayesian Analysis*, 1: 515–533. 980
- Gill, P. (2004). “Small-sample inference for the comparison of means of log-normal distributions.” *Biometrics*, 60: 525–527. 975
- Krishnamoorthy, K., Mallick, A., and Mathew, T. (2011). “Inference for the Lognormal mean and quantiles based on samples with left and right type I censoring.” *Technometrics*, 51: 72–83. 975
- Laforgia, A. and Natalini, P. (2010). “Some inequalities for modified Bessel functions.” *Journal of Inequalities and Applications*, Article ID 253035: 10 pages. 977
- Limpert, E., Stahel, W., and Abbt, M. (2001). “Log-normal distributions across the sciences: keys and clues.” *BioScience*, 51: 341–352. 975
- Nguyen, T., Chen, J., Gupta, A., and Dinh, K. (2003). “A proof of the conjecture on positive skewness of Generalized Inverse Gaussian distributions.” *Biometrika*, 90: 245–250. 978
- Rukhin, A. (1986). “Improved estimation in lognormal models.” *Journal of the American Statistical Association*, 81: 1046–1049. 976, 981, 983, 987, 988, 989
- Shen, H., Brown, L., and Zhi, H. (2006). “Efficient estimation of log-normal means with application to pharmacokinetic data.” *Statistics in Medicine*, 25: 3023–3038. 977, 987, 988, 989, 993, 994
- Singh, A., Singh, A., and Engelhardt, M. (2002). *The lognormal distribution in environmental applications*. Washington, DC: EPA/600/R-97/006. 993
- Zellner, A. (1971). “Bayesian and non-Bayesian analysis of the log-normal distribution and log-normal regression.” *Journal of the American Statistical Association*, 66: 327–330. 975, 976, 980, 984, 987
- Zhou, X. and Gao, S. (1997). “Confidence intervals of the log-normal mean.” *Statistics in Medicine*, 16: 783–790. 988
- Zou, G., Taleban, J., and Huo, C. (2009). “Confidence interval estimation for lognormal data with application to health economics.” *Computational Statistics and Data Analysis*, 53: 3755–3764. 975