Bayesian modeling of flash floods using generalized extreme value distribution with prior elicitation

Elijah Gaioni[1,*], Dipak Dey[2], Fabrizio Ruggeri[3]

[1]Watson Research Center, IBM, New York, USA,
[2]University of Connecticut, Statistics Department, Connecticut, USA,
[3]CNR IMATI, Milano, Italy

Abstract

Flash floods present a recurring problem in many parts of the Southwest United States, and the impact of such floods is felt on a social as well as an economic scale. The extent and severity of the damage resulting from such floods has been measured in many ways, including lives lost and dollar amounts of insurance claims. We focus on one historic US river, the Sabine, which has caused extensive flood damage in the past. Gauge height measurements along segments of the river permit an investigation into the distribution of the height of water at that location in the river over time. The height of water in a river is a function of not only current rainfall and snow melt, but also the geometry of the river itself and numerous characteristics of its surrounding areas, such as permeability of the surrounding soil and extent of human development. Quantifying some of these characteristics for direct incorporation into a model may be challenging in some instances, and the data itself may simply be unavailable in others. Consequently, as an alternative, an expert familiar with river flow may be able to indirectly impart some of this information to the model through quantiles of the quantity of interest, in this case, gauge height. Proper prior elicitation is a key element in Bayesian inference and the assessment of any prior distribution from experts' opinions is a critical aspect of this inference, both in getting the information and in transforming it into a functional form for the prior distribution. Many methods have been proposed to tackle the problem; most of them are based on the assessment of some features (e.g., quantiles, mean) of the parameter of interest, whereas very few look at features of the model itself, i.e., the observable quantities whose distribution is specified as a function of the parameter. We propose a novel approach which starts from quantiles of the parametric model, translates them into values of the parameters of interest, and uses them to specify a prior distribution. In conjunction with the likelihood, the prior is then used to develop the predictive distribution, which provides the basis for future expectations regarding the behavior of the river. The generalized extreme value distribution will be shown to model the height of water in the Sabine River quite well and we will discuss practical issues concerning the implementation of the approach, from graphical tools helpful in assessing the plausibility of the specified quantiles to adequate parameter transformations and sensitivity analysis.

Keywords: Bayesian robustness · Generalized extreme value distribution · Prior elicitation · Quantile.

Mathematics Subject Classification: Primary 62F15.

---

*Corresponding author. Email: egaioni@us.ibm.com

1. INTRODUCTION

The particular aspect of the distribution of gauge height that is crucially important when investigating floods is the upper tail, as opposed to the mean value or other measures of typically occurring water levels. The generalized extreme value (GEV) distribution arises naturally when modeling the maxima over a sequence of observations, and it frequently provides an empirically good fit to the extreme values in a data set even when the motivation through maxima doesn't apply. We propose this distribution as the likelihood function and compare its performance with the lighter-tailed normal distribution serving as the likelihood.

Our data consists of all available daily gauge height measurements at a point along the Sabine River near Ruliff, Texas, USA. The data was collected from September 27, 1997, through March 15, 2009, and is available on the United States geological survey website `http://waterdata.usgs.gov/usa/nwis/uv?site\_\no=08020900`. Since our interest is in the upper tail of the distribution of gauge heights, we began by discarding the smaller half of the data points, i.e., everything below the median gauge height. The retained data points are shown in Figure 1, along with a kernel-smoothed estimate of their density in Figure 2, where the flood stage is indicated by a line at the 25 foot mark in both figures.
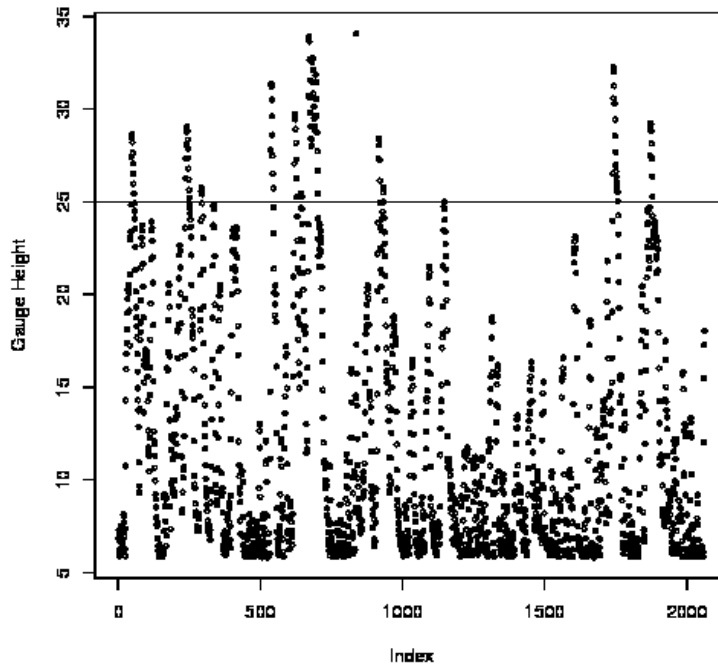

Figure 1. Time series of Sabine River sauge height.

The key inference pertaining to this river will concern the probability of the gauge height exceeding the 25 ft. flood stage threshold, resulting in flooding of the surrounding developed area. The probability of flooding will be shown to differ substantially under the different models. For purposes of model selection we hold out a portion of observations at the end of the data set for validation. There we compare the predicted probabilities of flooding under the GEV and normal models.

Bayesian analysis is markedly recognized by the subjective probability belief, or quantitative a priori description of unknown parameter θ. Without external support statisticians
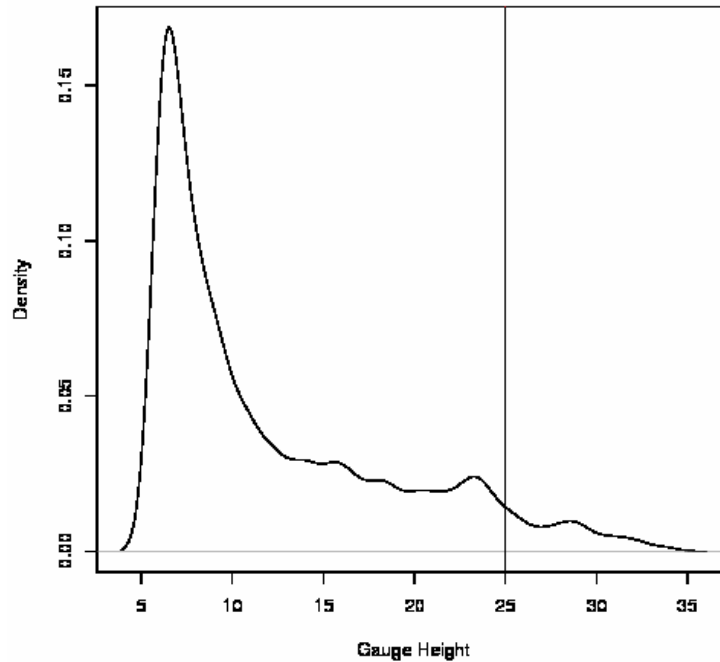
Figure 2. Kernel-Smoothed density estimate of Sabine River dauge height.

can only implement the whole computational process at his/her own will, say conjugate priors, Jeffreys' priors and other objective priors. On the other hand, expert opinion may be helpful when we investigate new, rare, complex or poorly understood phenomena.

Multivariate-normal related prior elicitation on the predictive prior space by requesting response summaries from experts was developed by Kadane et al. (1980), Garthwaite and Dickey (1988), Al-Awadhi and Garthwaite (1998) among others. However, those algorithms were limited to simple normal linear or AR(1) time series models. On the other hand, direct non-informative prior elicitation was discussed through piecewise conjugate priors (Meeden, 1992), entropy based priors (Jaynes, 1968, 1983), mixture of natural conjugate priors (Dalal and Hall, 1983) and others. Quantile based univariate prior elicitation for simple cases, say symmetric ones, was studied by Peterson and Miller (1964), Garthwaite and Dickey (1985), O'Hagan (1998) among others. A recent comprehensive review on probability elicitation was written by Garthwaite et al. (2005) and Dey and Liu (2007). (Berger, 1985, Chapter 3) also discussed subjective prior determination on the direct prior space by showing that a lack of sufficient tail information in continuous parameter space causes much difficulty for most of the approaches in practice, including the "histogram" approach, the "relative likelihood" approach, the entropy based method, and even the most used "matching a given functional form" method which often needs prior moments and others. Berger (1985, Chapter 3) envisioned that a quantile based approach poses as a better method since estimation of probabilities of regions are more attractive than working on moments. However, two situations deserve caution in application of quantile approaches: disagreement among multiple quantiles and incidental matching by multiple functional forms (Berger, 1985, Chapter 3). The key point to ease these concerns is to efficiently and precisely recover flexible parametric priors in a quantitative way other than those weak symmetric ones in order to implement the "sketching" principle (Berger, 1985, Chapter 3) for the downstream graphical verification.

In this paper, we propose a model based approach for direct model quantile elicitation,

which translates into prior distribution assessment. Although the proposed approach is quite general, we deem it particularly useful in cases such as this river data, in which assessments directly on the prior distribution are extremely difficult. In fact, this is typical of problems which can be modeled by a generalized extreme value distribution (e.g., extreme rainfalls and wind speeds).

The paper is organized as follows. Section 2 motivates the approach by discussing its application to the Sabine river data, which lends itself naturally to use of the generalized extreme value distribution. Section 3 formally presents the techniques used to fit the model using the quantile-based Bayesian prior elicitation technique. Section 4 tailors the approach to the GEV distribution. The practical implementation of this procedure is discussed in Section 5, and the full Sabine River data analysis is presented in Section 6. Section 7 briefly discusses robustness considerations, and Section 8 concludes the paper with a discussion of the key findings.

2. Flash Flooding Along the Sabine River

The Sabine River has the largest volume of water discharged at its mouth among all rivers in Texas. It is 555 miles long, with a drainage basin of nearly 10,000 square miles, and lies in a region with heavy rainfall. Furthermore, its basin is characterized by flat slopes and a wide floodplain.

Rivers in this region are known for dangerous floods and lie in the so-called "Flash Flood Alley" region of Texas. The quick onset of flash floods in this region stems from the proximity of its flat plains to the Texas Hill region, which channels large volumes of rainfall into shallow river beds. With average annual precipitation between thirty-seven and fifty inches, the region in which the Sabine River lies experiences frequent flooding, with large floods occurring every five years on average. Perhaps the greatest potential for disastrous consequences arises when this particular type of flooding, known as flash flooding, occurs.

Flooding leads to damage of the area surrounding the river when the water level rises above a height known as the "flood stage", the value of which depends on the proximity of the river to developed land. By estimating the distribution of gauge heights, we can calculate the return levels, which provide important information about the likelihood and severity of floods. A T-year return level refers to a flood of a magnitude that is exceeded, on average, only once every T-years. Thus, the higher return levels are very sensitive to the upper tails of the distribution of gauge height. Consequently, we focus on addressing the difficulties involved in extracting information pertaining to these upper tails through the methodology introduced below.

3. Statement of the Problem

Following the Bayesian paradigm, we consider a random vector X with distribution function $F(x|\theta)$ and density function $f(x|\theta)$ and we are interested in specifying the prior distribution $\pi(\theta)$, where θ can be scalar or vector valued.

Experts are asked to specify distinct quantiles $G_1 < \cdots < G_k$, with k fixed, from the marginal distribution of X, corresponding to the probabilities $q_1 < \cdots < q_k$. Each quantile specification G, corresponding to a probability q, leads to the equation in θ given by

$$q = F(G|\theta). \tag{1}$$

We suppose that the parameter θ is obtained as solution of Equation (1), i.e.,

$$\theta = h(q, G). \tag{2}$$

As an example, consider an exponential model $X \sim \mathcal{E}(\theta)$, then Equations (1) and (2) become, respectively $q = 1 - \exp(-\theta G)$ and $\theta = [-\log(1-q)]/G$.

In general, m quantiles are needed to obtain a solution

$$\theta = h(q_1, \ldots, q_m, G_1, \ldots, G_m) \tag{3}$$

when considering a multivariate parameter θ of size m.

Solutions to Equations (2) or (3) lead to expert's assessment on the value of the the parameter θ. Multiple quantile specifications lead to multiple assessment on θ. Specification of n quantiles when the parameter θ has size $m \leq n$ lead to $\binom{n}{m}$ values of θ: the number follows from all possible combinations of m elements chosen among n. In particular, n values of θ are obtained for univariate parameters.

We treat the $\binom{n}{m}$ values of θ as a sample from the expert's prior distribution and we use it to specify one by, for instance, using sample moments to specify the hyperparameters of the prior distribution.

Back to the exponential model $X \sim \mathcal{E}(\theta)$, the statistician could deem a gamma prior $\mathcal{G}(\alpha, \beta)$ as a suitable model for the prior because of its conjugacy property. The parameters α and β are such that $\mathrm{E}(\theta) = \frac{\alpha}{\beta}$ and $\mathrm{Var}(\theta) = \frac{\alpha}{\beta^2}$, so that

$$\alpha = \frac{[\mathrm{E}(\theta)]^2}{\mathrm{Var}(\theta)} \text{ and } \beta = \frac{\mathrm{E}(\theta)}{\mathrm{Var}(\theta)}. \tag{4}$$

Given the sample $\theta_1, \ldots, \theta_n$, the sample mean and sample variance are computed and substituted in Equation (4) and the values of α and β are therefore determined.

4. Generalized Extreme Value Distribution

We discuss how the proposed approach applies to data modeled with a generalized extreme value (GEV) distribution. The interest for the GEV model is twofold, both mathematical and practical. The GEV model depends upon three parameters (location, scale and shape) which are related to its quantiles in a nontrivial way and they require a transformation before the proposed approach can be applied satisfactorily (as discussed in Section 5).

The practical interest in applying the model to the GEV model resides in the difficulty of assessing prior distributions directly on the parameters, e.g., through some of their quantiles. Typical data modeled by a GEV distribution are extreme rainfalls (e.g., Coles and Tawn, 1996) and extreme wind speeds (e.g., Coles and Powell, 1996). Experts are expected to have opinions on the probability of rainfall (or wind speed) exceeding some thresholds rather than on quantiles of the distribution of the GEV location, scale and shape parameters.

Based on three distinct quantiles, we present the equations whose solution (in general, numerical) links parameters to quantiles. Therefore, quantile assessments are transformed into parameter ones which are used to choose hyperparameters of the prior distribution on the parameter.

We suppose three different probabilities $q_1 < q_2 < q_3$ are chosen and the expert assigns distinct values $G_1 < G_2 < G_3$ to the corresponding quantiles.

The cumulative density function of the GEV distribution is given by

$$F(x) = \exp\left\{-\left[1 + \lambda\left(\frac{x-\mu}{\sigma}\right)\right]_+^{-1/\lambda}\right\}$$

and its density function is

$$f(x) = \frac{1}{\sigma}\left[1 + \lambda\left(\frac{x-\mu}{\sigma}\right)\right]_+^{-1/\lambda-1}\exp\left\{-\left[1 + \lambda\left(\frac{x-\mu}{\sigma}\right)\right]_+^{-1/\lambda}\right\},$$

with $\mu \in \mathbb{R}$, $\lambda \in \mathbb{R}$ and $\sigma \in \mathbb{R}^+$.

The following holds for G_q, quantile of order q:

$$q = \exp\left\{-\left[1 + \lambda\left(\frac{G_q-\mu}{\sigma}\right)\right]_+^{-1/\lambda}\right\}. \tag{5}$$

From now on, we suppose $1 + \lambda\left(\frac{G_q-\mu}{\sigma}\right) > 0$.

Starting from Equation (5) and considering the pair (q_1, G_1), simple computations lead to

$$\sigma = \frac{\lambda(G_1 - \mu)}{[-\log(q_1)]^{-\lambda} - 1}, \tag{6}$$

and, after introducing the pair (q_2, G_2),

$$\mu = \frac{G_1 K_2(\lambda) - G_2 K_1(\lambda)}{K_2(\lambda) - K_1(\lambda)}, \tag{7}$$

with $K_i(\lambda) = \frac{\lambda}{\sigma}(G_{q_i} - \mu)$, $i = 1, 2$. It is worth noticing that Equation (7) holds for $\lambda \neq 0$.

Finally, when considering the other pair (q_3, G_3), it follows that λ is the solution to

$$\Delta_1 \exp(-\alpha_1\lambda) + \Delta_2 \exp(-\alpha_2\lambda) + \Delta_3 \exp(-\alpha_3\lambda) = 0, \tag{8}$$

with $\Delta_1 = G_2 - G_3$, $\Delta_2 = G_3 - G_1$, $\Delta_3 = G_1 - G_2$ and $\alpha_i = \log(-\log(q_i))$, $i = 1, 2, 3$. Note that $\exp(-\alpha_1\lambda) < \exp(-\alpha_2\lambda) < \exp(-\alpha_3\lambda)$ and $\Delta_1 < 0$, $\Delta_2 > 0$ and $\Delta_3 < 0$.

Once Equation (8) is solved and λ found, then μ and σ are obtained, in sequence, from Equation (7) and (6). Although in most cases Equation (8) will lead to numerical solutions, a closed form solution can be obtained in the very simple (but useful) case in which $\alpha_3 = 3\alpha_1$ and $\alpha_2 = 2\alpha_1$. In this case, Equation (8) becomes

$$\Delta_3 x^3 + \Delta_2 x^2 + \Delta_1 x = 0,$$

taking $x = \exp(-\alpha_1\lambda)$. We consider the solution $\tilde{x} = \frac{G_2 - G_3}{G_1 - G_2}$, different from $x = 0, 1$, which correspond to $\lambda = \infty$ and $\lambda = 0$, respectively. In this case Equations (8), (7) and (6)

become

$$\tilde{\lambda} = \frac{\log([G_1 - G_2]/[G_2 - G_3])]}{\log(-\log(q_1))},$$

$$\tilde{\mu} = \frac{G_1^2 + G_2^2 - G_1 G_2 - G_1 G_3}{G_2 - G_3} \quad \text{and}$$

$$\tilde{\sigma} = \frac{[G_2 - G_1]^3}{[G_2 - G_3][2G_2 - G_3 - G_1]} \times \frac{\log([G_1 - G_2]/[G_2 - G_3])}{\log(-\log(q_1))}.$$

A very interesting case, for practical purposes, is given by $q_1 = 0.80$, $q_2 = 0.9514$ and $q_3 = 0.9884$, where the last two can be approximated by 0.95 and 0.99, respectively.

We prefer to consider all the three parameters defined on the real line, as motivated later in Section 5. Therefore we consider $\log(\sigma)$ instead of σ.

We now suppose that the expert is able to provide $n > 3$ quantiles, so that $k = \binom{n}{3}$ values of $\theta = (\mu, \log(\sigma), \lambda)$ are obtained as solutions to Equations (6), (7) and (8).

From such sample $\theta_1, \ldots, \theta_k$, it is possible to obtain sample means, sample variances and sample covariances and they can be used to determine the hyperparameters of the prior distribution. For large k a trivariate Gaussian approximation is a possible choice, provided it gives negligible probability to the inadmissible values of the parameters. The choice of the Gaussian distribution will be motivated in Section 5.

5. PRACTICAL IMPLEMENTATION

So far we have assumed that experts can elicit quantiles which are used to find hyperparameters of priors with pre-specified functional form (gamma for the exponential model and Gaussian for the GEV model). Our goal is to provide a procedure which allows interactive elicitation, permitting checks and changes by the expert during the process and requiring as little intervention as possible by the statistician performing the modeling.

The former goal is achieved by presenting the expert the graphs of the predictive density functions, which are obtained as a result of his/her quantiles elicitation (see, e.g., Figure 10). Although the expert is asked only about few quantiles, his/her knowledge is such that he/she could assess the plausibility of such a density function over the entirety of its domain.

There are cases (e.g., the exponential model) in which a specific functional form of the prior distribution could be considered for many reasons (e.g., the gamma distribution as conjugate prior for the parameter of the exponential model). In our quest for a procedure which requires the least possible effort in modeling and allows for implementation of user-friendly software, we prefer to resort to a unique framework, namely Gaussian priors. First of all, we would transform parameters in such a way that their domain coincides with that of a (multidimensional) Gaussian distribution, or, as it is standard practice in statistics, such that the Gaussian distribution assigns negligible probability to the unfeasible values of the parameter. At this point, given a (possibly transformed) sample $(\theta_1, \ldots, \theta_k)$, a Q-Q plot is used to assess the normality assumption. If the Q-Q plot shows a lack of normality, then Box-Cox transformations are applied to the parameter θ so that the Gaussian assumption is justified for the new transformed parameter. Therefore, the sample $(\theta_1, \ldots, \theta_k)$ is used to estimate the means and variance-covariance matrix of the Gaussian prior distributions on the transformed parameter $\tilde{\theta}$. Finally, the posterior distribution on the actual parameter θ is obtained by transformation from the posterior distribution on $\tilde{\theta}$.

6.   Examples

In this Section we apply the proposed approach to the Sabine River data indicated in the Introduction. The Sabine data was split into two parts: the fitted portion and the validation portion. The time series and kernel smoothed density estimates of these two portions are provided below in Figures 3, 4, 5 and 6.
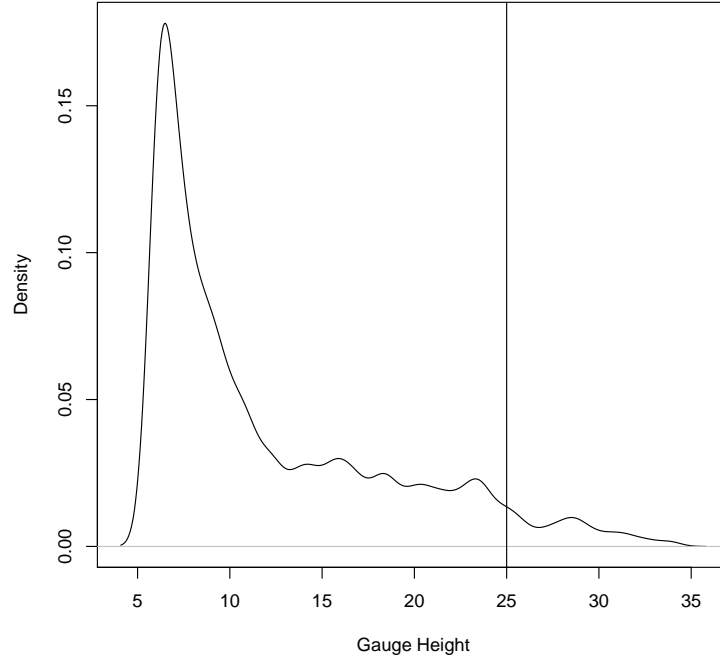

Figure 3. Kernel smoothed estimate of fitted portion of Sabine River gauge height density

We use empirical quantiles from the historical Sabine River data to serve as expert-provided prior information. Thus, the following quantiles are assumed:
Percentiles: $(0.01, 0.05, 0.10, 0.25, 0.50, 0.75, 0.90, 0.95, 0.99)$.
Quantiles: $(5.85, 5.95, 6.15, 6.84, 9.3, 16.02, 23.1, 25.76, 30.83)$
The $\binom{9}{3}$ combinations of percentile/quantile pairs for the GEV-MVN model result in the sample parameter points shown in Figure 7, and are used to develop the following method of moment type estimators for the MVN prior:
Sample mean vector: $(\mu, \sigma, \xi) = (7.08, 3.10, 2.11)$.
Sample covariance matrix: $\begin{pmatrix} 2.85 & 1.68 & 0.74 \\ 1.68 & 1.04 & 0.51 \\ 0.74 & 0.51 & 0.42 \end{pmatrix}$.

Similarly, the $\binom{9}{2}$ combinations of percentile/quantile pairs for the normal-MVN model result in the sample parameter points shown in Figure 8, and are used to develop the following method of moment type estimators for the MVN prior:
Sample mean vector: $(\mu, \sigma) = (11.56, 5.53)$.
Sample covariance matrix: $\begin{pmatrix} 11.20 & 3.55 \\ 3.55 & 10.28 \end{pmatrix}$.

The 95% highest posterior density credible region for the normal-MVN model parameters is shown in Figure 9, and the posterior densities for the GEV-MVN and normal-MVN models are illustrated in Figures 11 and 10.
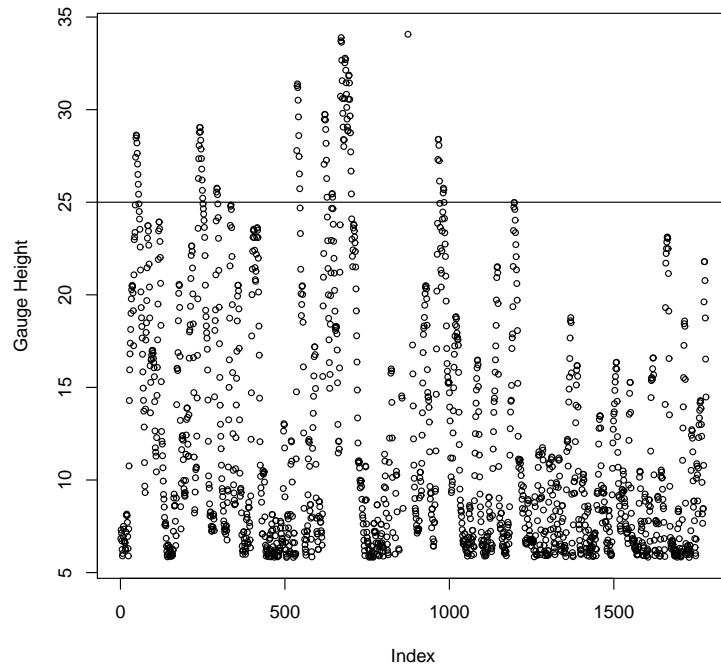
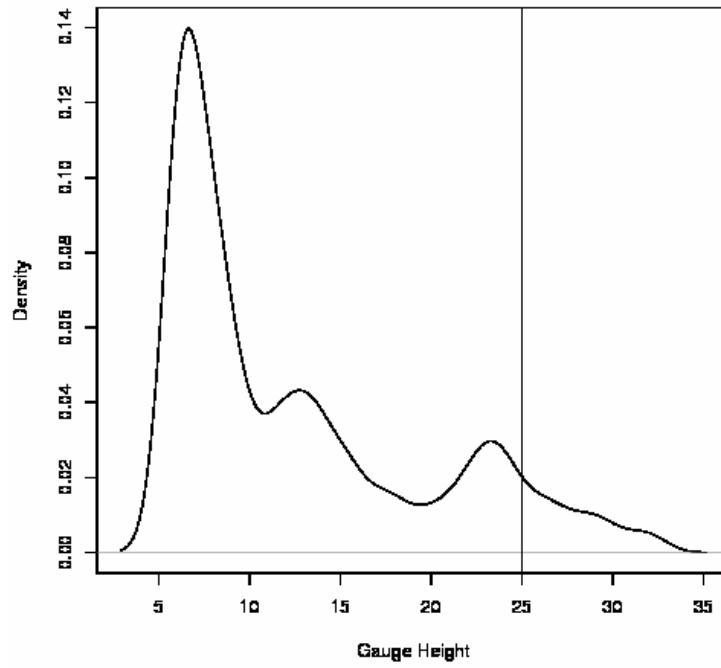Figure 4. Time series of fitted portion of Sabine River gauge height data.

Figure 5. Kernel smoothed estimate of validation portion of Sabine River gauge height density.
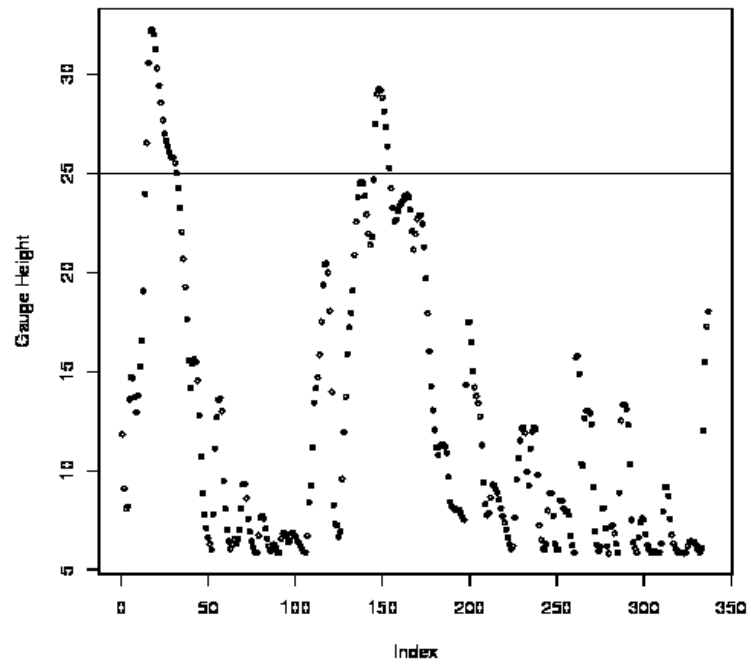
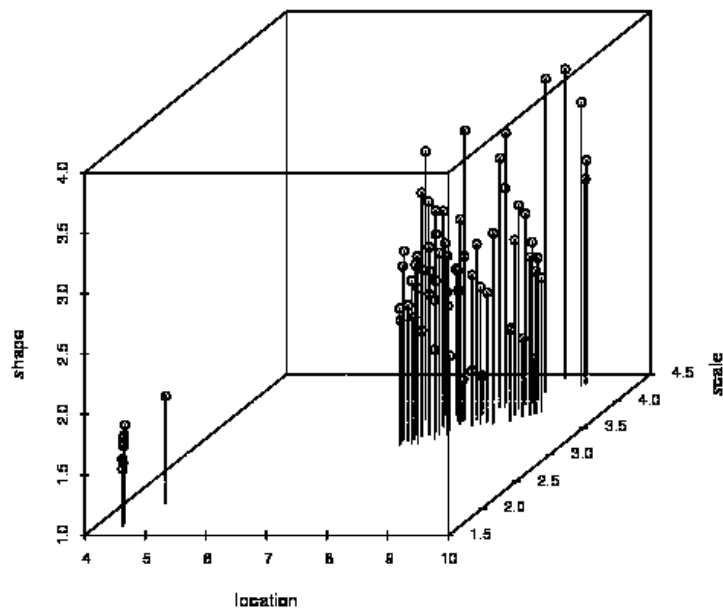Figure 6. Time series of validation portion of Sabine River gauge height data.

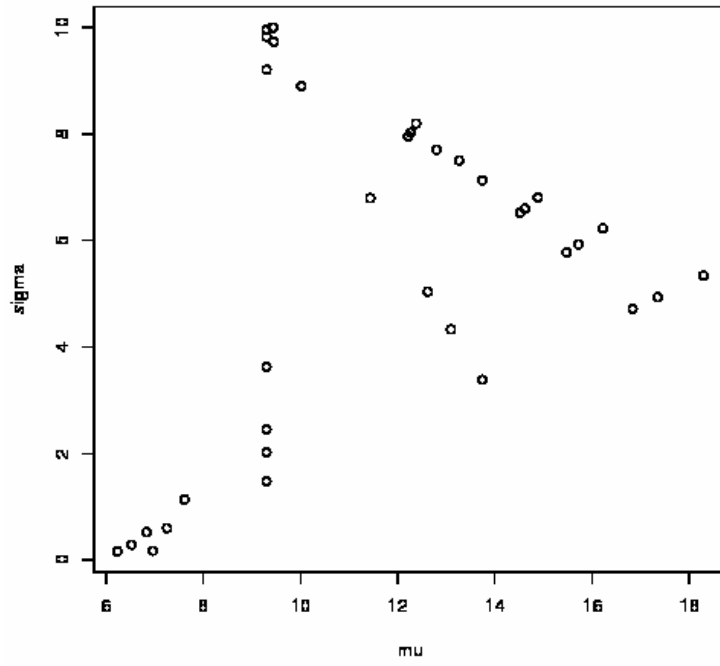
Figure 7. Sample parameter points for GEV-MVN model.

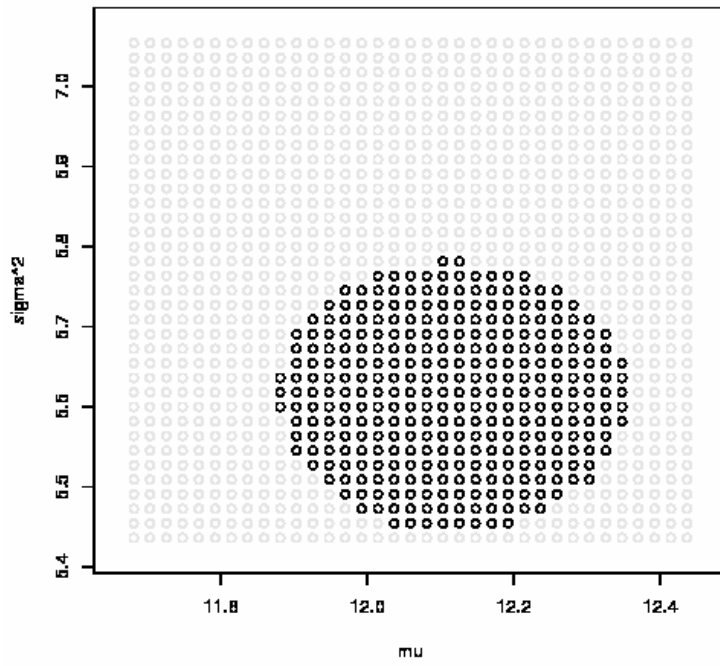Figure 8. Sample parameter points for normal-MVN model.

Figure 9. 95% highest posterior density credible region for normal model parameters.
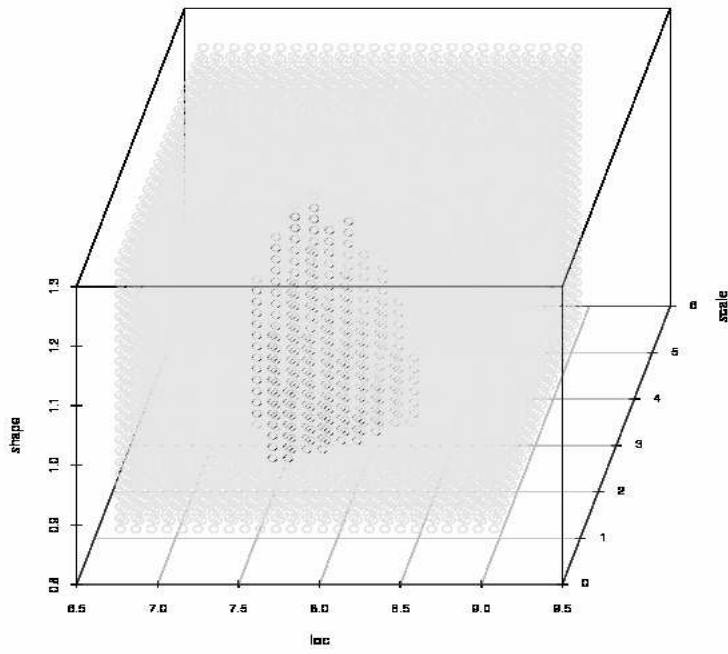
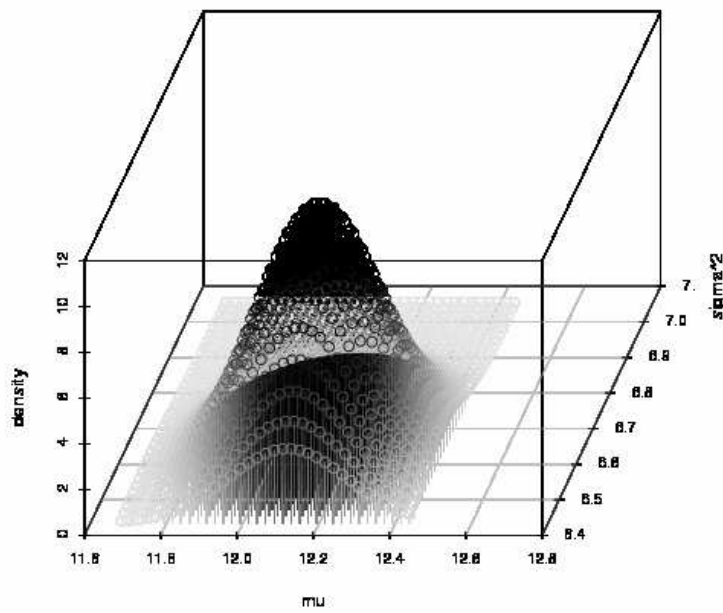Figure 10. Posterior density for GEV-MVN model.

Figure 11. Posterior Density for GEV-MVN Model

The posterior mean and covariance for the GEV model are:

Posterior mean vector: $(\mu, \sigma, \xi) = (7.65, 2.52, 0.96)$,

posterior covariance matrix: $\begin{pmatrix} 0.027 & 0.032 & -0.010 \\ 0.032 & 0.038 & -0.012 \\ -0.010 & -0.012 & 0.004 \end{pmatrix}$.

The posterior mean and covariance for the normal model are:

Posterior mean vector: $(\mu, \sigma) = (12.14, 6.61)$.

Posterior covariance matrix: $\begin{pmatrix} 0.037 & 0.004 \\ 0.004 & 0.017 \end{pmatrix}$.

These posterior densities are then combined with the likelihood specifications to provide the predictive density for a new data set, z, given the original data, x, defined as

$$f(z|x) = \int f(z|\theta)\pi(\theta|x)d\theta.$$

The predictive density is sketched in Figure 12 for the normal-MVN model, and in Figure 13 for the GEV-MVN model.
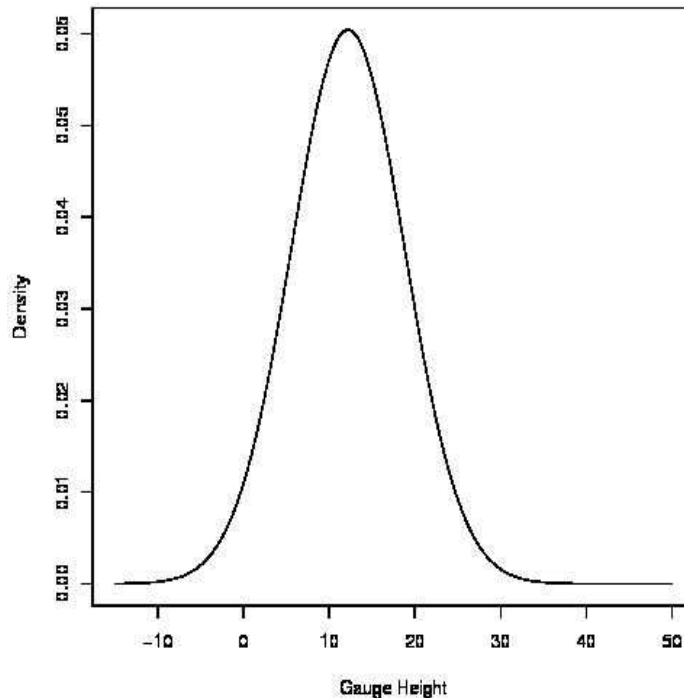

Figure 12. Predictive density for the normal-MVN model.

We then use the validation data set to compare the performance of these two predictive models in anticipating the relative frequency of flood occurrence. In the validation data set the gauge heights exceed the flood stage 8.01% of the time. The normal-MVN model predicts 2.61%, which is too low by 5.40%, and the GEV-MVN model predicts 10.85%, which is too high by 2.84%.

This predictive criterion for selecting the better model based on flood prediction chooses the GEV-MVN model. The Sabine River gauge height appears to require a heavy upper-tailed distribution. Recalling that we initially discarded the data points below the median, we see that the normal-MVN model considers exceedance of the flood stage a 76 day return level event, whereas the GEV-MVN model considers it an 18 day return level event. This illustrates the pronounced differences in inference based on light-tailed and heavy-tailed distributions when modeling extreme events.
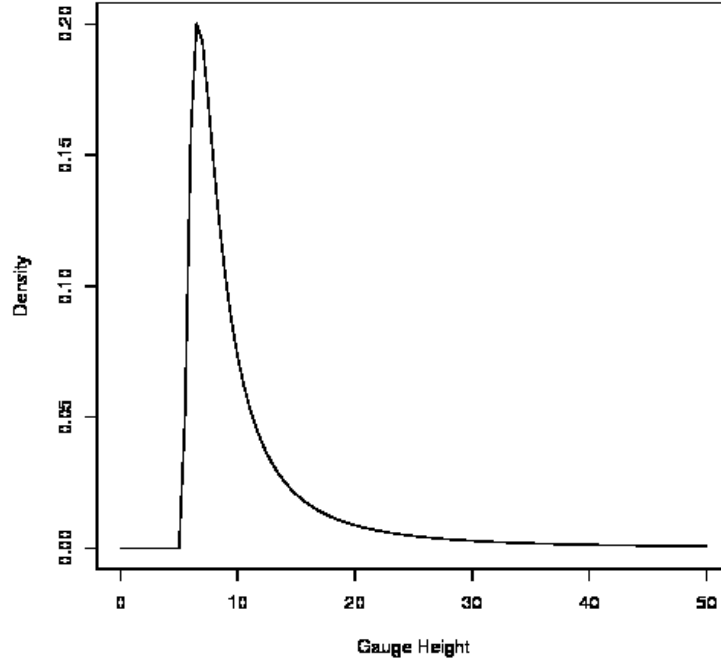
Figure 13. Predictive density for the GEV-MVN model.

7. Bayesian Robustness

The paper uses experts' quantile specification to determine a unique prior distribution, in general a Gaussian one as motivated in the previous sections. Although the choice of a unique prior is convenient for inference purposes, the selection of its functional form and parameters is rather arbitrary. This aspect, considered as a critical issue by both Bayesians and non-Bayesians, can be addressed by taking classes of priors and performing a sensitivity analysis, as described in Rios Insua and Ruggeri (2000) and references therein.

The case of quantile specification on the model $X \sim f(x|\theta)$, with X defined on $\Re$ and $\theta \in \Theta$, has been addressed by Betrò et al. (1994) as a particular generalized moments constrained class of priors, defined as

$$\Gamma = \left\{ \pi : \int_\Theta h_i(\theta)\pi(d\theta) = q_i, i = 1, \ldots, k \right\}, \tag{9}$$

where $G_1 < \cdots < G_k$ are the quantiles corresponding to $q_1 < \cdots < q_k$, and

$$h_i(\theta) = \int_{-\infty}^{G_i} f(x|\theta)dx, i = 1, \ldots, k. \tag{10}$$

When interested in inferences about posterior quantities like $E^{\pi(\theta|x)}g(\theta)$, the range

$$\rho = \sup_{\pi \in \Gamma} E^{\pi(\theta|x)}g(\theta) - \inf_{\pi \in \Gamma} E^{\pi(\theta|x)}g(\theta) \tag{11}$$

is a typical sensitivity measure to study the effects of priors varying in a class Γ. As shown in Betrò et al. (1994), the range in Equation (11) is obtained for discrete distributions concentrated in at most $k+1$ points when considering the class Γ defined in Equation (9). Efficient algorithms have been developed by Betro and coauthors; for references, see the most recent paper by Betrò (2009).

8. Conclusions

This approach, whereby the families of likelihood and prior distributions are specified by the experts, along with the quantiles, circumvents the need to elicit prior information on the unobservable prior space. We emphasize that the methodology illustrated above enabled indirect specification of the tail behavior through expert elicited quantiles of the observable quantities. This tail behavior was important in correctly determining the return level for the extreme events leading to potentially severe flooding.

An important goal was to minimize the number of arbitrary choices that the statistician must make. Indeed, this methodology enables one to work exclusively with expert-provided quantiles, inducing a sample of points in the prior space, which are then used to estimate the hyperparameters. The prior distribution both provides extra flexibility in modeling the response, and serves to reconcile the $n > m$ quantiles, which would overspecify a single parametric family of likelihoods alone, with the expert-provided information. The family of likelihoods, the specification of which is guided by the expert, is modified by the quantiles in this fashion. In this case we compared the performance of GEV and normal likelihoods, finding the former to be superior in predicting the tail behavior of the flood process.

This approach of eliciting information on the observable space, not the parameter space, as is more common, resembles an empirical Bayes approach. However, instead of assuming the availability of actual data, we can work exclusively with expert-provided quantiles. Thus, the basis for our modeling process is quite different than the traditional setup. To the best of our knowledge, there are no standard techniques that handle this data framework. However, the nearest to what could be considered a standard approach to fitting a model based on a set of quantiles would perhaps involve working with the cdfs. Here one could minimize a measure of the distance between the cdf points provided by the expert-provided quantiles on the one hand, and a cdf belonging to the family of parametric likelihoods provided by the experts on the other. This would be related to the Kolmogorov-Smirnov test statistic if the measure used is the maximum absolute deviation between the two cdfs. The important difference is that the cdf implied by the expert-provided quantiles in this case, is not actually an empirical cdf, since it is not derived from observed data.

For comparative purposes, we note that by proceeding in this more standard manner, minimizing the distance between a cdf in the $\text{GEV}(\mu, \sigma, \xi)$ family of distributions and the expert-provided quantiles, results in the following:

If one uses the maximum absolute deviation as a measure of the inter-cdf distance, one obtains $\mu = 8.51, \sigma = 3.31, \xi = 0.44$, which results in a probability of 6.9% of exceeding the flood stage. This understates the flood frequencies observed in the validation data set, but not by very much.

We further note that the solution offered in this paper was illustrated with a parametric model, using quantile assessment on it to specify a parametric prior distribution, with known functional form. We note that we could also have used a similar approach when dealing with a nonparametric model. As an example, consider a Dirichlet process with a parameter given, apart from a constant, by a GEV distribution. Quantiles could be used either to specify the GEV distribution directly (looking for the parameter values giving the best fit to the quantiles) or a second-level prior distribution on the GEV parameters, as done here. The former approach would be fully nonparametric, whereas the other is a mixture of Dirichlet processes (Antoniak, 1974).

A semi-parametric approach could consider a GEV distribution whose parameters are chosen by a Dirichlet process with a Gaussian parameter obtained as in the current paper; this approach leads to a so-called Dirichlet process mixture model (Lo, 1984).

REFERENCES

Al-Awadhi, S.A., Garthwaite, P.H., 1998. An elicitation method for multivariate normal distributions. Communications in Statistics – Theory and Methods, 27, 112342.

Antoniak, C.E., 1974. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. Annals of Statistics, 2, 1152-1174.

Berger, J., 1985. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. Annals of Statistics, 2, 1152-1174.

Betrò, B., 2009. Numerical treatment of Bayesian robustness problems. International Journal of Approximate Reasoning, 50, 279-288.

Betrò, B., Męczarski, M., Ruggeri, F., 1994. Robust Bayesian analysis under generalized moments conditions. Journal of Statistical Planning and Inference, 41, 257-266.

Coles, S.G., Powell, E.A., 1996. Bayesian methods in extreme value modelling: a review and new developments. International Statistical Review, 64, 119-136.

Coles, S.G., Tawn, J.A., 1996. A Bayesian analysis of extreme rainfall data. Applied of Statistics, 45, 463-478.

Dalal, S.R., Hall, W.J., 1983. Approximating priors by mixtures of natural conjugate priors. Journal of The Royal Statistical Society Series B – Statistical Methodology, 45, 272-286.

Dey, D.K., Liu, J., 2007. A quantitative study of quantile based direct prior elicitation from expert opinion. Bayesian Analysis, 2, 137-166.

Garthwaite, P.H., Dickey, J.M., 1985. Double- and single-bisection methods for subjective probability assessment in a location-scale family. Journal of Econometrics, 29, 149-163.

Garthwaite, P.H., Dickey, J.M., 1988. Quantifying expert opinion in linear regression problems. Journal of The Royal Statistical Society Series B – Statistical Methodology, 50, 462-474.

Garthwaite, P.H., Kadane, J.B., O'Hagan, A., 2005. Statistical methods for eliciting probability distributions. Journal of the American Statistical Association, 100, 680-701.

Jaynes, E.T., 1968. Prior probabilities. IEEE Transactions on Systems Science and Cybernetics, SSC-4, 227.

Jaynes, E.T., 1983. Papers On Probability. Statistics and Statistical Physics, R.D. Rosenkrantz. D. Reidel publishing Co., Dordrecth.

Kadane, J.B., 1980. Predictive and structural methods for eliciting priot distributions. A. Zellner (Ed.), Studies in Bayesian Econometrics and Statistics in Honor of Harold Jeffreys, North Holland Publishing Company, 89-93.

Lo, A.Y., 1984. On a class of Bayesian nonparametric estimates: I density estimates. Annals of Statistics, 12, 351-357.

Meeden, G., 1992. An elicitation procedure using piecewise conjugate priors. Goel, P.K. and Iyengar, N.S. (Eds.), Bayesian Analysis in Statistics and Econometrics, Springer-Verlag, 423-432.

O'Hagan, A., 1998. Eliciting expert beliefs in substantial practical applications. The Statistician, 47, 21-35.

Peterson, C.R., Miller, A., 1964. Mode, median and mean as optimal strategies. Journal of Experimental Psychology, 68, 363-367.

Ríos-Insua, D., Ruggeri, F. (Eds.), 2000. Robust Bayesian Analysis. Springer, New York.