



Published in final edited form as:

Technometrics. 2011 May 1; 53(2): . doi:10.1198/TECH.2011.10041.

Bayesian Local Contamination Models for Multivariate Outliers

Garritt L. Page and David B. Dunson

Department of Statistical Science, Duke University, Durham, NC 27704

Garritt L. Page: page@stat.duke.edu; David B. Dunson: dunson@stat.duke.edu

Abstract

In studies where data are generated from multiple locations or sources it is common for there to exist observations that are quite unlike the majority. Motivated by the application of establishing a reference value in an inter-laboratory setting when outlying labs are present, we propose a local contamination model that is able to accommodate unusual multivariate realizations in a flexible way. The proposed method models the process level of a hierarchical model using a mixture with a parametric component and a possibly nonparametric contamination. Much of the flexibility in the methodology is achieved by allowing varying random subsets of the elements in the lab-specific mean vectors to be allocated to the contamination component. Computational methods are developed and the methodology is compared to three other possible approaches using a simulation study. We apply the proposed method to a NIST/NOAA sponsored inter-laboratory study which motivated the methodological development.

Keywords

Bayesian robustness; Component-wise classification; Inter-laboratory studies; Mixtures

1. INTRODUCTION

Random effects or multilevel models are commonly used to model data that have a hierarchical or nested structure. These types of models are appealing because of their wide applicability and the availability of estimation and inference for subject-specific and global parameters. As with many statistical procedures, estimation and inference available from these models are sensitive to the presence of observations that are unlike the majority. If the data have a multivariate structure, then handling outliers can be even more complicated as observation vectors can be composed of a combination of elements, some of which are similar in magnitude to the majority of other observations and some of which are outlying.

As an example, consider a NIST/NOAA sponsored inter-laboratory study conducted to assess the quality of trace element measurements in the marine mammal population. The 33 participating laboratories were instructed to produce five replicate measurements of the concentration of 15 trace elements in a sample of marine mammal tissue. Figure 1 provides boxplots of measured concentrations of arsenic (As) and selenium (Se) where each boxplot corresponds to a lab. Notice that Lab 18 recorded concentrations for both trace elements that are much larger than the other labs. Conversely, Lab 17 reported concentrations for Se that are larger than the majority of the labs, but the concentrations measured for As seem quite reasonable relative to the majority of the other labs. In light of this, there is a need to develop methods that are able to handle multivariate outliers in a flexible way.

There is a large literature devoted to the development of statistical methodologies that are robust to the presence of outliers. Rocke (1983), Mandel (1995), and Bednarski and Zontek (1996) handled outliers by proposing robust estimators, such as M-estimators or estimators

resulting from Fréchet differentiable functionals. Muller and Uhlig (2001) and Lischer (1996) proposed estimators based on the differences between observations. Another approach is to discard all measurements that are classified as outliers via a detection method such as those proposed by Peña and Prieto (2001) and Penny and Jolliffe (2001). However, the uncertainty associated with outlier detection is not considered. Song, Zhang, and Qu (2007) provided an example of accommodating outliers by using heavy-tailed distributions. Although the fatter tails of the t -distribution more readily accommodate outliers compared to a Gaussian distribution, it still is fairly restrictive in that its shape is symmetric and unimodal.

Introducing some distributional perturbation by way of a finite mixture is another approach. Box and Tiao (1968) first proposed this method and used a two-component mixture to accommodate outliers. Approaches that use mixtures of some kind can be quite flexible but typically classify an entire multivariate object as being an outlier or not.

In this article we develop a Bayesian methodology that allows the different elements of a multivariate vector to vary in their outlier status, while accommodating uncertainty in outlier classification in estimation and inferences. The general idea is to introduce some distributional contamination (Bayarri and Morales 2003) to accommodate multivariate outliers. This is done by constructing a hierarchical model where the process is modeled with a mixture of mixtures. A majority (non-outlying) component is modeled by a multivariate normal, while the other component corresponds to unusual observations and is modeled with a finite mixture. Much of the flexibility in the methodology is achieved by allowing allocation to the non-Gaussian component to vary for the different vector elements. One might think of this as a type of local clustering of the location/source mean vectors. The notion of local clustering or clustering multivariate objects on a subset of attributes has been introduced in the literature. Specifically, Dunson (2009) used the local clustering idea to choose a prior for an unknown random effects distribution within a hierarchical model and Hoff (2006) developed a model-based clustering approach that clusters a multivariate vector of attributes using a subset of the attributes.

The methods developed here are applicable to any study in which multivariate measurements arising from different studies are to be combined. This could include, for example, meta-analysis and inter-laboratory studies. For clarity and ease of exposition we motivate ideas from an inter-laboratory perspective. Because of this we provide a very brief introduction to inter-laboratory studies here. These studies are conducted to ensure measurement capability for commerce, evaluate national and international equivalence of measure, and validate measurement devices and methods or standard materials. Typically, the overarching goal in the analysis of data produced by inter-laboratory studies is to establish a reference value, which can sometimes be thought of as an estimate of a measurand (quantity intended to be measured), and estimate its uncertainty. To determine a “degree of equivalence” each laboratory’s measurements are compared to the reference value. This analysis is usually carried out under the guidelines set forth in *The Guide to the Expression of Uncertainty in Measurement* (GUM) created by the International Organization of Standardization (ISO). Since its inception criticisms and alternative approaches to estimating a reference value have been proposed in the literature (Gleser 1998; Rukhin and Vangel 1998; Rukhin 2007; and Toman 2007).

In what follows, Section 2 provides a detailed description of the model and Section 3 develops computational methods. Section 4 describes a simulation study conducted to compare the performance of the proposed method to three reasonable alternatives. Section 5 provides an example using a NIST/NOAA marine mammal inter-laboratory dataset. In Section 6 we make conclusions.

2. DESCRIPTION OF THE LOCAL CONTAMINATION MODEL

Though we focus on an inter-laboratory application for clarity, the methodology can be implemented in a more general setting. The proposed model will be referred to as a local contamination (LC) model. We assume that the replicates for each laboratory vary according to a normal distribution and that if a laboratory makes an unusual measurement for a particular element (compared to the rest of the labs) it will continue to do so for that element. However, within-laboratory outliers can be handled in a fairly straightforward manner through the use of a heavy-tailed residual density such as a multivariate t -distribution (e.g., Section 5.4). For sake of simplicity and ease of exposition we describe the LC model with a Gaussian residual distribution.

Let $\mathbf{y}_{ij} = (y_{ij1}, y_{ij2}, \dots, y_{ijp})'$ be the j th vector of p measurements taken by the i th laboratory with $j = 1, \dots, n_i$ and $i = 1, \dots, m$. Assume that

$$\mathbf{y}_{ij} \stackrel{\text{ind}}{\sim} N_p(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i) \quad (1)$$

with $\boldsymbol{\mu}_i = (\mu_{i1}, \mu_{i2}, \dots, \mu_{ip})'$ a lab-specific mean and $\boldsymbol{\Sigma}_i$ a lab-specific covariance which we model with an inverse-Wishart [i.e., $\boldsymbol{\Sigma}_i \sim \text{IW}(v, \mathbf{A})$]. For studies where measurements are physically constrained to be nonnegative one can simply log-transform the measurements prior to analysis.

A novel contribution of the LC model is how the $\boldsymbol{\mu}_i$'s are modeled. Ultimately, we model the $\boldsymbol{\mu}_i$ with a multivariate normal whose mean and variance allow each $\boldsymbol{\mu}_i$ to potentially be composed of elements that are members of a majority (non-outlying) component and others that are members of a non-majority (outlying) component which we refer to as a contamination. This latent allocation structure is modeled by introducing classification variables that classify the p elements of each lab mean vector $\boldsymbol{\mu}_i$ as being part of the majority or not. Let $\boldsymbol{\gamma}_i = (\gamma_{i1}, \gamma_{i2}, \dots, \gamma_{ip})'$ be a p -dimensional vector of 0's and 1's such that

$$\gamma_{ik} = \begin{cases} 1 & \text{if } \mu_{ik} \in \mathcal{M}_i \\ 0 & \text{if } \mu_{ik} \notin \mathcal{M}_i, \end{cases}$$

where \mathcal{M}_i is a collection of elements that make up the majority component for the i th lab, $d_i = \sum_{k=1}^p \gamma_{ik}$ is the number of elements in the majority component for the i th lab, and $\gamma_{ik} \stackrel{\text{ind}}{\sim} \text{Ber}(\pi_i)$ with $\pi_i \stackrel{\text{iid}}{\sim} \text{Beta}(a_\pi, b_\pi)$ (a_π and b_π are user-supplied). π_i is the prior probability that lab i falls within the majority component for a randomly selected element. By specifying a hyperprior for π_i , we allow the data to inform about the proportion of outlying elements. Typically, one would elicit a_π and b_π based on prior knowledge of the proportion of outlying elements with $a_\pi \gg b_\pi$ so that $E(\pi_i) \gg 0.5$, with ideally $E(\pi_i) \approx 1 - \varepsilon$ for small ε . This leads to a local ε -contaminated model.

We now introduce atoms $\{(\boldsymbol{\mu}_h^*, \boldsymbol{\Sigma}_h^*); h = 0, 1, \dots, L\}$ that are used to construct the mean and covariance matrix corresponding to $\boldsymbol{\mu}_i$. The pair $(\boldsymbol{\mu}_0^*, \boldsymbol{\Sigma}_0^*)$ are the mean vector and covariance matrix that correspond to the majority component (which makes $\boldsymbol{\mu}_0^*$ of principal interest in an inter-laboratory study analysis), while $(\boldsymbol{\mu}_h^*, \boldsymbol{\Sigma}_h^*)$ for $h = 1, \dots, L$ correspond to the L clusters that make up the contamination. $\boldsymbol{\mu}_0^*$ is sampled a priori from $N_p(\mathbf{m}, \mathbf{S})$ while the

μ_h^{**} 's for $h = 1, \dots, L$ are sampled independently from $t_p(\mathbf{m}, \mathbf{S}, \nu)$ a priori. [Here $t_p(\mathbf{m}, \mathbf{S}, \nu)$ denotes a scaled (\mathbf{S}) and shifted (\mathbf{m}) p -dimensional t -distribution with ν degrees of freedom.] Multivariate t -distributions are used to accommodate the possibility of contamination cluster locations being highly variable. For computational purposes (detailed in Section 3) the \sum_h^* 's are diagonal matrices with diagonal entries σ_{hk}^{*2} , $k = 1, \dots, p$. To preserve conjugacy we use an inverse-Gamma prior $\sigma_{hk}^{*2} \stackrel{\text{ind}}{\sim} \text{IG}(a_\sigma, b_\sigma)$ for $k = 1, \dots, p$ (values for a_σ, b_σ are supplied by the user). Thus, depending on γ_i , the mean and covariance matrix corresponding to μ_i is made up of elements from μ_0^* and \sum_0^* and/or elements from μ_ℓ^* and \sum_ℓ^* for some $\ell = 1, \dots, L$.

The value of L actually represents an upper bound on the number of clusters that make up the contamination. The form of the contamination is motivated by a finite Dirichlet approximation to the Dirichlet process as proposed by Ishwaran and Zarepour (2002). As this upper bound increases, there is convergence to a nonparameteric limit, but the finite approximation is somewhat simpler to implement. Allowing the contamination to potentially consist of more than one cluster provides more flexibility in handling outliers compared to a one-cluster contamination. The necessity of this flexibility can be seen in Figure 1. It is fairly obvious that the As and Se entries of Lab 18's mean vector will occupy a cluster of the contamination. Also, it is possible that the As entry of Lab 26's mean vector will be allocated to the contamination. If this turns out to be the case, then at least two contamination clusters will be necessary to accommodate the two element means.

To identify the ℓ th contamination cluster used to construct the mean vector and covariance matrix associated with μ_i we introduce an S_i for each of the $i = 1, \dots, m$ labs. The S_i 's take on values $\ell = 1, \dots, L$ with $\Pr(S_i = \ell) = v_\ell$. For simplicity, all elements of the mean and covariance matrix of μ_i that are allocated to the contamination come from the same contamination cluster. The vector of probability weights on the L clusters that make up the contamination is modeled as

$$\nu = (v_1, v_2, \dots, v_L)' \sim \text{Dir}(\alpha/L, \dots, \alpha/L)$$

with $\alpha = 1$ being fixed in the applications that we consider to favor high weights on few clusters.

Finally, we model the μ_i 's independently for $i = 1, \dots, m$ with

$$\mu_i \stackrel{\text{ind}}{\sim} N_p(\delta_{\gamma_i}, \mathbf{T}_{\gamma_i}).$$

The entries of the vector δ_{γ_i} depend on γ_{ik} in the following way. If $\gamma_{ik} = 1$, then $\delta_{\gamma_{ik}} = \mu_{0k}^*$, otherwise, $\delta_{\gamma_{ik}} = \mu_{S_i k}^*$ for all $k = 1, \dots, p$. This structure can be succinctly written as

$$\delta_{\gamma_i} = \gamma_i \otimes \mu_0^* + (1 - \gamma_i) \otimes \mu_{S_i}^*,$$

where \otimes denotes the Hadamard or element-wise multiplication.

\mathbf{T}_{γ_i} is constructed in a similar fashion. Since \sum_h^* for $h = 0, \dots, L$, are diagonal matrices \mathbf{T}_{γ_i} is diagonal as well. The k th entry on the diagonal of \mathbf{T}_{γ_i} is the k th diagonal value of \sum_0^* (or σ_{0k}^{*2}) if $\gamma_{ik} = 1$; otherwise, it is the k th diagonal entry of $\sum_{S_i}^*$ (or $\sigma_{S_i k}^{*2}$). Thus we can construct \mathbf{T}_{γ_i} as

$$\mathbf{T}_{\gamma_i} = \gamma_i \gamma_i' \otimes \sum_0^* + (1 - \gamma_i)(1 - \gamma_i)' \otimes \sum_{S_i}^* .$$

Figure 2 provides a graphical representation of how the mean and covariance matrix of μ_i is constructed.

Values for all parameters denoted by Roman letters need to be specified by the user. To automate prior specification and avoid problems inherent to dealing with multivariate data on vastly different scales, we recommend standardizing the p variables prior to analysis. If this is done, the prior specification of $\mathbf{m} = \mathbf{0}$ and $\mathbf{S} = \mathbf{I}_p$ is natural. Because the p variables are standardized using all data (including any outliers), values for (a_σ, b_σ) should be chosen so the majority of prior distribution mass associated with the σ_{hk}^2 's is less than 1. Informal investigation indicated that inferences from the LC procedure are fairly insensitive to the prior specification for σ_{hk}^2 . In this article we use an inverse-Gamma distribution with mean 0.25 and standard deviation 0.5 giving $\Pr(\sigma_{hk}^2 > 1) \approx 0.025$. Following suggestions made by Verdine and Wasserman (1991), to make probability of being an outlier small (≈ 0.05), we set $a_\pi = 9.5$ and $b_\pi = 0.5$.

3. COMPUTATION

The joint posterior distribution of the parameters in the LC model is analytically intractable. We use a Gibbs sampler to obtain correlated draws from the posterior distribution. The full conditional distributions which can be used to construct a Markov chain whose stationary distribution is the joint posterior are described. In what follows we use $[\theta|-]$ to denote the distribution of θ conditioned on all other parameters. As an example, $[\mu_i|-]$ is shorthand for

$[\mu_i | \{\mu_j\}_j, \{\sum_i\}_{i=1}^m, \{\gamma_i\}_{i=1}^m, \{\mu_h^*\}_{h=0}^L, \{\sum_h^*\}_{h=0}^L, \{\pi_i\}_{i=1}^m, \{v\}_{\ell=1}^L, \{S_i\}_{i=1}^m, \mathbf{y}]$. The following full conditionals are fairly common and can be derived using routine algebra:

$$[\mu_i|-] \sim N_p \left([n_i \sum_i^{-1} + \mathbf{T}_{\gamma_i}^{-1}]^{-1} [n_i \sum_i^{-1} \bar{\mathbf{y}}_{ij} + \mathbf{T}_{\gamma_i}^{-1} \delta_{\gamma_i}] [n_i \sum_i^{-1} + \mathbf{T}_{\gamma_i}^{-1}]^{-1} \right), \quad (2)$$

$$[\sum_i|-] \sim \text{IW} \left(n_i + v, \sum_{j=1}^{n_i} (\mathbf{y}_{ij} - \mu_i)(\mathbf{y}_{ij} - \mu_i)' + \mathbf{A} \right), \quad (3)$$

$$[\gamma_{ik}|-] \sim \text{Ber}(p^*) \text{ with } p^* = \varphi(\mu_{ik}; \mu_{0k}^*, \sigma_{0k}^{*2}) \pi_i / \left(\varphi(\mu_{ik}; \mu_{0k}^*, \sigma_{0k}^{*2}) \pi_i + \varphi(\mu_{ik}; \mu_{S_i k}^*, \sigma_{S_i k}^{*2}) (1 - \pi_i) \right), \quad (4)$$

$$[\pi_i|-] \sim \text{Beta} \left(\sum_{k=1}^p \gamma_{ik} + a, p - \sum_{k=1}^p \gamma_{ik} + b \right), \quad (5)$$

$$[\nu|-] \sim \text{Dir}(\alpha_1^*, \dots, \alpha_L^*) \text{ with } \alpha_\ell^* = \sum_{i=1}^m 1[S_i = \ell] + \frac{\alpha}{L}, \ell = 1, \dots, L, \quad (6)$$

where $\phi(\cdot; \mu, \sigma^2)$ denotes a normal density with mean μ and variance σ^2 .

Full conditionals are also needed for the component-specific parameters. We first provide the full conditionals for μ_0^* and \sum_0^* . Here, we introduce some useful notation. Let $\mathbf{T}_{0:\gamma_i=0}$ denote the matrix that results from setting the entries of the rows and columns of the matrix \mathbf{T}_{γ_i} that are associated with $\gamma_{ik} = 0$ for $k = 1, \dots, p$ to zero. Similarly $\mathbf{T}_{0:\gamma_i=1}$ denotes the matrix whose row and column entries associated with $\gamma_{ik} = 1$ for $k = 1, \dots, p$ are set to zero. The full conditional for μ_0^* is

$$[\mu_0^*|-] \sim N_p \left(\left[\sum_{i=1}^m \mathbf{T}_{0:\gamma_i=0}^{-1} + \mathbf{S}_0^{-1} \right]^{-1} \times \left[\sum_{i=1}^m \mathbf{T}_{0:\gamma_i=0}^{-1} \mu_i + \mathbf{S}_0^{-1} \mathbf{m}_0 \right], \left[\sum_{i=1}^m \mathbf{T}_{0:\gamma_i=0}^{-1} + \mathbf{S}_0^{-1} \right]^{-1} \right). \quad (7)$$

This result is a consequence of the assumption that \sum_0^* is diagonal.

Since \sum_0^* is a diagonal we consider the σ_{0k}^{2*} 's individually for each k . The complete conditional for the k th diagonal element of \sum_0^* is

$$[\sigma_{0k}^{2*}|-] \sim \text{IG} \left(a_\sigma + \frac{1}{2} \sum_{i=1}^m \gamma_{ik}, \left[\frac{1}{b_\sigma} + \frac{1}{2} \sum_{i=1}^m \gamma_{ik} (\mu_{ik} - \mu_{0k}^*)^2 \right]^{-1} \right). \quad (8)$$

Derivations of (7) and (8) are provided in the Appendix.

To facilitate computation we use a scale mixture of normal representation of the t -distributions that correspond to the contamination cluster locations. This requires introducing an auxiliary variable $[\omega_\ell \sim \text{IG}(w/2, 2/w)]$ for each of the L contamination clusters where w is the degrees of freedom (which we set to 4). The complete conditionals for

$\{(\mu_h^*, \sum_h^*); h = 1, \dots, L\}$ are then

$$[\mu_h^*|-] \sim N_p \left(\left[\sum_{i:S_i=h} (\omega_h \mathbf{T})_{0:\gamma_i=1}^{-1} + \mathbf{S}_h^{-1} \right]^{-1} \times \left[\sum_{i:S_i=h} (\omega_h \mathbf{T})_{0:\gamma_i=1}^{-1} \mu_i + \mathbf{S}_h^{-1} \mathbf{m}_h \right], \left[\sum_{i:S_i=h} (\omega_h \mathbf{T})_{0:\gamma_i=1}^{-1} + \mathbf{S}_h^{-1} \right]^{-1} \right), \quad (9)$$

$$[\sigma_{hk}^{2*}|-] \sim \text{IG} \left(a_\sigma + \frac{1}{2} \sum_{i:S_i=h} (1 - \gamma_{ik}), \left[\frac{1}{b_\sigma} + \frac{1}{2} \sum_{i:S_i=h} (1 - \gamma_{ik}) (\mu_{ik} - \mu_{hk}^*)^2 \right]^{-1} \right). \quad (10)$$

The full conditional of the auxiliary variable ω is $[\omega_h|-] \sim \text{IG}(0.5(w+p))$,

$$[\omega_h|-] \sim \text{IG}(0.5(w+p), (0.5[(\mu_h - \mathbf{m})' \mathbf{S}^{-1} (\mu_h - \mathbf{m}) + w])^{-1}). \quad (11)$$

The full conditional of S_i is discrete with probability mass function

$$\Pr[S_i=\ell|-]=\frac{\varphi(\mu_i;\delta_{\ell\gamma_i},\mathbf{T}_{\ell\gamma_i})v_\ell}{\sum_{h=1}^L\varphi(\mu_i;\delta_{h\gamma_i},\mathbf{T}_{h\gamma_i})v_h}\text{ for } \ell=1,\dots,L, \quad (12)$$

where $\varphi(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes a multivariate normal density, $\delta_{\ell\gamma_i}=\gamma_i \otimes \mu_0^*+(1-\gamma_i) \otimes \mu_{S_i=\ell}^*$, and $\mathbf{T}_{\ell\gamma_i}=\gamma_i\gamma_i' \otimes \sum_0^*+(1-\gamma_i)(1-\gamma_i)' \otimes \sum_{S_i=\ell}^*$.

A Markov chain associated with the joint distribution of interest can be had by iteratively cycling through the complete conditional distributions on an individual basis.

When using auxiliary variables for classifying in a mixture setting it is possible for chains coming from an MCMC algorithm to mix poorly. To improve mixing, we use an adaptive type of Gibbs sampler (Roberts and Rosenthal 2007). This was done by replacing p^* in (4) with

$$p^*+\exp\{-0.01(t-1)\}(0.5-p^*),$$

where $t = 1, \dots, M$ denotes the t th MCMC iterate. This initially pulls p^* to 0.5 but then converges to (4) exponentially fast as $t \rightarrow \infty$. Since the adaptation vanishes at an exponential rate, the necessary regularity conditions are satisfied and the algorithm converges to the correct distribution.

4. SIMULATION STUDY

We compare the performance of the LC model in estimating a reference vector to three other reasonable alternatives by way of a simulation study. The simulation study consists of generating several multivariate datasets (with and without outlying elements) and for each, estimating a reference vector using three competing methods and the LC model. The four procedures were compared using frequentist metrics such as bias and mean squared error along with credible region area and coverage. In this section, we briefly describe the competing methods, detail how datasets were generated, and provide the simulation study results.

4.1 Description of Competing Methods

We compare the performance of the LC procedure in estimating μ_0^* to three reasonable alternatives via a simulation study. The description of the three competitors follows.

1. The first competing method is a random effects model with random effects assumed to originate from a MVN distribution. Specifically,

$\mathbf{y}_{ij} \stackrel{\text{iid}}{\sim} N_p(\mu_i, \boldsymbol{\Sigma}), \mu_i \stackrel{\text{iid}}{\sim} N_p(\mu_0^*, \boldsymbol{\Sigma}_0)$, and $\mu_0^* \sim N_p(\mathbf{m}, \mathbf{S})$. $\boldsymbol{\Sigma}$ and $\boldsymbol{\Sigma}_0$ are drawn from inverse-Wishart distributions. We refer to this model as the MVN model.

2. The second competing method is a random effects model with random effects assumed to originate from a multi-variate t -distribution. More specifically,

$\mathbf{y}_{ij} \stackrel{\text{iid}}{\sim} N_p(\mu_i, \boldsymbol{\Sigma}), \mu_i \stackrel{\text{iid}}{\sim} t_p(\mu_0^*, \boldsymbol{\Sigma}_0, \eta)$, and $\mu_0^* \sim N_p(\mathbf{m}, \mathbf{S})$. A uniform (0, 100) prior is used for η . The other parameters were assigned the same priors as in the MVN model. The t -distribution is often used as a robust alternative to the MVN when outliers are present and we refer to this model as the MVT model.

3. The third competing method is a random effects model with the random effects being modeled with an unknown density. A Dirichlet Process (DP) is used as a prior for the unknown density. That is, $\mathbf{y}_{ij} \stackrel{\text{ind}}{\sim} N_p(\mu_i, \Sigma)$, $\mu_i \sim \mathcal{P}$, and $\mathcal{P} \sim \text{DP}(\alpha P_0)$. Here P_0 follows a $N_p(\mu_0^*, \Sigma_0)$ and $\mu_0^* \sim N_p(\mathbf{m}, \mathbf{S})$. Similarly to the LC model we fixed $\alpha = 1$. Other parameters were assigned the same priors as in the MVN and MVT models. This model is very flexible in accommodating unusual observations and we refer to it as the DP model.

4.2 Creation of Synthetic Datasets

We consider the MVN and MVT models as data generating mechanisms. We generate μ_i vectors from a MVN or MVT distribution after fixing values for μ_0^* and Σ_0 . Then after fixing Σ , lab-specific observation vectors are generated by using a MVN with mean μ_i and covariance Σ and then log-transformed. We used the marine mammal dataset as a guide to picking appropriate values for μ_0^* , Σ_0 , and Σ . The marine mammal dataset was also used to choose the number of laboratories (30), the number of observations per laboratory (5), and the dimension of the observation vectors (15).

We fix the k th entry of μ_0^* at the median of the k th element computed across all laboratories from the marine mammal dataset. Σ_0 is fixed to be a diagonal matrix whose k th entry is set by computing the variance between the empirical lab-specific means for the k th element. For the data generation process, Σ is also a diagonal matrix whose k th entry is fixed at the average of the lab-specific sample variances of the k th element. For datasets containing outliers, 10 of the 30 labs are randomly selected to have outliers in at least one entry of μ_i . Five of the ten have one element randomly selected to be an outlier, two have two elements as outliers, two have five elements, and one has ten elements. This outlier structure is similar to that found in the marine mammal dataset. An outlier is generated by setting the mean for an outlying element to $\mu_{0k}^* + 6\sigma_{0kk}$. Figure 3 provides an example of the datasets used in the simulation study with and without outliers.

4.3 Results

Using the MVN, 500 datasets with and without outliers are generated. This is repeated using MVT, giving a total of 2000 datasets. For each dataset, posterior distributions of μ_0^* were obtained using the four procedures. Posterior distributions are summarized by posterior mean vectors and 15-dimensional 95% credible regions. The credible regions are 15-dimensional rectangles formed as the product of fifteen $(0.95)^{1/15} \times 100\%$ credible intervals, one for each of $\mu_{01}^*, \dots, \mu_{015}^*$. We refer to the posterior means as $E[\mu_0^*|\mathbf{y}]$. To compare the four procedures' performance in estimating μ_0^* , we use empirical coverage ratios and credible region volume along with two metrics related to the frequentist bias and MSE. As a type of total absolute bias the following was computed for each procedure:

$$\text{bias} = \frac{1}{D} \sum_{d=1}^D \left(\sum_{k=1}^{15} |E[\mu_0^*|\mathbf{y}]_{dk} - \mu_{0k}^*| \right).$$

Here, $E[\mu_0^*|\mathbf{y}]_{dk}$ is the k th element of the $E[\mu_0^*|\mathbf{y}]$ and d is an index for the $D = 500$ datasets that were generated. Values in the MSE column were computed using

$$\text{MSE} = \frac{1}{D} \sum_{d=1}^D \left(\sum_{k=1}^{15} \text{Var}[\mu_0^* | \mathbf{y}]_{dk} + (E[\mu_0^* | \mathbf{y}]_{dk} - \mu_{0k}^*)^2 \right).$$

This might be thought of as a type of total MSE averaged over the 500 datasets. Results can be found in Table 1.

When no outliers are present and the data are generated using the MVN model, results from the LC procedure are very much comparable to the other four procedures in terms of coverage, volume, bias, and MSE. The fact that these intervals did not attain the nominal 95% coverage is to be expected as we are working with Bayesian posterior probability intervals which have a different interpretation than confidence intervals in a finite sample setting. When outliers are introduced, the MVN and DP models are more negatively affected than the MVT and LC models and the LC model performs much better in all four metrics. When the data are generated using MVT and there are no outliers, the MVT model performs the best out of the four procedures in terms of bias and MSE. This is to be expected. However, when outliers are introduced, the LC procedure has a lower MSE and bias than the MVT. It appears that the LC procedure has a clear advantage over the other three procedures at estimating μ_0^* when outlying laboratories are present. In other simulation results (not shown) the advantages of using the LC model (in terms of smaller bias and MSE) compared to the MVT model become more obvious as the ratio of outlying labs to participating labs increases or as the distance between an outlying lab mean and the majority component mean increases.

To demonstrate the utility of a contamination model with more than one cluster, we report in Table 2 several summaries of the outlying clusters identified by the LC model. At each iteration of the MCMC algorithm we count the number of clusters that contained an element that had been allocated to the contamination. The “clusters occupied” column reports the median of this value across MCMC iterates and the 500 datasets. For data generated using the MVN with outliers, approximately three clusters of the contamination were occupied while four were occupied for data generated from MVT with outliers. We also enumerated the total number of elements that were classified as outliers [using $E(\gamma_{ik} | \mathbf{y}) < 0.1$ as the outlier criterion]. The values found in the column “outlying elements” is the median number of elements allocated to the contamination across the 500 datasets. Recall that 29 elements were randomly selected to be outliers for datasets that contained outliers. The median number of labs that had at least one outlying element is also consistent with the true value of 10 or 0 labs (as is seen in the “outlying labs” column). Table 2 indicates that the LC model is fairly accurate in its allocation of outlying elements.

5. DATA ANALYSIS OF THE NIST/NOAA SPONSORED INTER-LABORATORY STUDY

In this section we briefly describe the data that were produced from the marine mammal inter-laboratory experiment, provide results of the analysis from the LC model along with the three other methods outlined in the previous section, and assess model fit using cross-validation.

5.1 Description of NIST/NOAA Inter-Laboratory Study Data

In 2005 a NIST/NOAA sponsored inter-laboratory study was conducted to improve the quality of trace element measurements in marine environmental systems. The NIST prepared fresh-frozen marine mammal control materials [white-sided dolphin liver

homogenate (QC04LH4)]. A glass jar containing approximately 8–10 grams of the frozen material was distributed to 33 participating laboratories. Each lab was asked to keep the material in an environment that would preserve its authenticity and to divide the material into five aliquots. Measurements of 15 trace elements (Ag, As, Cd, Co, Cs, Cu, Fe, Hg, Mn, Mo, Rb, Se, Sn, V, and Zn) using in-house analytical techniques were to be taken on each aliquot. The raw measurement results were submitted to the NIST. Figure 1 provides a graphical display of measurement results for As and Se. Additional details are provided in the article by Christopher et al. (2007).

Some labs did not measure all 15 trace element concentrations on each aliquot. We imputed values for the missing data under the missing at random (MAR) assumption (Gelman et al. 2004) within our proposed MCMC algorithm. The core assumption of MAR is that the missingness mechanism does not depend on the missing data. Information regarding the rationale behind a lab's decision not to measure the concentration of all 15 trace elements on each aliquot is not available. However, there is no obvious indicator of a MAR assumption violation and it seems completely plausible that the missingness mechanism did not depend on the trace element.

Algorithms to impute missing values and update posterior draws for the four procedures were written in the C programming language. For each procedure 10,000 posterior draws were collected after a burn-in of 20,000 and thinning of 5. With regard to computation time, the computer code associated with the two simpler models required slightly less time (MVN, 584 seconds and MVT, 602 seconds) to run than that of the more complicated models (DP, 980 seconds and LC, 818 seconds). The added flexibility afforded by the LC model comes at a minimal computational cost.

5.2 Results From the Marine Mammal Data Analysis

Prior to analysis the data were normalized and log-transformed. After transforming back to the original scale, marginal posterior means and 95% credible intervals were calculated for each trace element and are provided in Table 3.

From Table 3 it can be seen that the 95% credible interval widths are generally the smallest under the LC model. In fact, the average credible interval width is smaller for the LC (1.58) compared to the DP model (6.10), the MVT model (1.97), and the MVN model (5.49).

Notice that in terms of estimating the location of μ_{0k}^* , the LC and MVT procedures are very similar. The advantage that the LC procedure has over the MVT is mainly in the estimation of the variability associated with μ_{0k}^* .

We set $L = 10$ when fitting the LC model to the marine mammal dataset. Table 4 gives the posterior probability distribution on the number of occupied contamination clusters. Five occupied contamination clusters turned out to have the highest posterior probability.

A nice characteristic of the LC model is that it is possible to estimate the posterior probability of the i th lab being part of the majority. This can be done by computing $E(\pi_i|\mathbf{y})$, the posterior mean of π_i . In addition, $E(\gamma_{ik}|\mathbf{y})$, the posterior mean of γ_{ik} , provides the posterior probability that the k th element corresponding to the i th lab is part of the majority. These probabilities are listed in Table 5. The number of trace elements for each lab such that $E(\gamma_{ik}|\mathbf{y}) < 0.1$, $E(\gamma_{ik}|\mathbf{y}) < 0.5$, and $E(\gamma_{ik}|\mathbf{y}) < 0.9$ are found in parentheses. Lab 18 reported measurements for multiple trace elements that might be considered to be far from the majority and this is reflected in its posterior probability of being an outlier. With regard to the trace elements referenced in Section 1, the posterior probability of As belonging to the majority component for Lab 17 is 0.87, for Lab 18 is 0.00, and for Lab 26 is 0.05, while the

posterior probability of Se belonging to the majority component for Lab 17 is 0.025, for Lab 18 is 0.00, and for Lab 26 is 0.80.

We assessed model fit using cross-validation. A testing partition of the marine mammal data was created by removing two randomly selected observations from 10 randomly selected laboratories. The MVN, MVT, DP, and LC models were fit to the remaining portion of the marine mammal dataset and the removed observations were imputed within the MCMC algorithm.

Typically the goal in using out-of-sample prediction to assess model fit is to determine how concentrated the posterior distribution for the predictive values is around the truth. Often this is done by computing the mean squared prediction error (MSPE). However, the MSPE only provides a measure of how close a point estimate of the predicted value is to the true value. That is, it does not consider in any way how concentrated the posterior distribution is around the truth. In addition, when outliers are present it is not completely clear that the MSPE is a good metric to assess out-of-sample prediction. Since outliers make up a small minority of the observations, overly simple models that underestimate the true uncertainty may do well in terms of MSPE as the bias introduced by the outliers is dominated by the smaller variance. Because of this we propose the following measure of how concentrated the posterior distribution for a predicted value is to the truth:

$$\frac{1}{p} \sum_{k=1}^p \left[\frac{1}{T} \sum_{t=1}^T (y_{pred_{tk}} - y_{test_k})^2 \right]. \quad (13)$$

Here $y_{pred_{tk}}$ is the t th MCMC iterate of the k th entry of the predicted vector and y_{test_k} is the k th entry of a data vector from the testing dataset. This metric incorporates each MCMC iterate to assess concentration around the truth. The value of (13) averaged over the twenty predictions under the LC model (0.062) was slightly smaller than that of the other three models [MVN (0.064), MVT (0.064), and DP (1.130)]. The average prediction interval width was also slightly smaller for the LC model (1.59) compared to the other three [MVN (1.63), MVT (1.60), and DP (5.18)].

5.3 Univariate Analyses

A very simple approach to the analysis of these types of data is to perform an independent analysis for each trace element. Indeed this was the approach the NIST originally used to establish a reference value (Christopher et al. 2007). Here we compare the performance of the LC model in establishing a reference value to that obtained by performing 15 independent univariate analyses. There are any number of approaches one might take to conduct a univariate analysis that would be robust (accommodating) to the presence of outliers. None of these approaches are exactly comparable/analogous to the LC model. Because of this we use the following univariate model that accommodates univariate outliers very flexibly:

$$\begin{aligned} y_{ij} &\sim N(\mu_i, \sigma_i^2), \\ (\mu_i, \sigma_i^2) &\sim \mathcal{P}, \\ \mathcal{P} &\sim \text{DP}(\alpha, P_0) \quad \text{with } P_0 = N(\mu_0^*, \sigma_0^{2*}). \end{aligned}$$

In addition a normal-inv-Gamma prior was used for (μ_0^*, σ_0^{2*}) and an inv-Gamma prior was used for σ_i^2 . Hyperprior values equaling those used for the LC model were used. Using a Dirichlet process prior to model the density of (μ_i, σ_i^2) provides a great deal of flexibility

and hopefully this results in the ability to make reasonable comparisons between the LC model and 15 independent univariate analyses. The above model was fit to each of the 15 trace elements found in the marine mammal dataset. Posterior means and 95% credible intervals associated with μ_0^* were computed for each trace element and are listed in Table 6.

Apart from the fact that performing separate univariate analyses effectively ignores the dependence structure and therefore can produce misleading results, there appear to be some pragmatic gains to using the LC procedure. The credible interval widths associated with the LC model are shorter compared to those coming from the 15 independent analyses. Obviously shorter credible intervals are desirable only if point estimates corresponding to the intervals are near the truth. We argue that this is indeed the case in the present setting using the results from the simulation study (the bias associated with estimates using the LC procedure was very small in the presence of outliers) and the fact that the LC model and 15 independent univariate analyses produce point estimates that are fairly similar. In the cases for which the point estimates from the two procedures differ (e.g., Se), those associated with the 15 independent univariate analyses are greater than those from the LC procedure. This could be an indication that results from the univariate analyses are still influenced by outlying labs as outlying labs produced measurements greater than the majority in the marine mammal dataset.

5.4 Using a t -Distribution Residual to Accommodate Within-Lab Outliers

In addition to outlying labs, it is possible for within-lab outliers to exist. These types of outliers can easily be accommodated by changing (1) to

$$\mathbf{y}_{ij} \stackrel{\text{ind}}{\sim} t_p(\mu_i, \Sigma_i, \kappa),$$

where $t_p(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \kappa)$ denotes a shifted ($\boldsymbol{\mu}$) and scaled ($\boldsymbol{\Sigma}$) p -dimensional t -distribution with κ degrees of freedom. Conceptually this is a very straightforward modification of the LC model. Computationally, we lose conjugacy if the t density is used directly. However, the t -distribution can be represented as a scale mixture of normals and this characterization of the t preserves full conjugacy. Thus, a t -density residual can be obtained with the following hierarchy:

$$\begin{aligned} \mathbf{y}_{ij} &\stackrel{\text{ind}}{\sim} N_p(\mu_i, \lambda_{ij} \Sigma_i), \\ \lambda_{ij} &\stackrel{\text{iid}}{\sim} \text{IG}(\kappa/2, 2/\kappa). \end{aligned}$$

With this representation of the t -distribution (2) and (3) change in the following ways:

$$\begin{aligned} [\mu_i | -] &\sim N_p \left(\left[\sum_{j=1}^{n_i} \lambda_{ij}^{-1} + \mathbf{T}_{\gamma_i}^{-1} \right]^{-1} \times \left[\sum_{j=1}^{n_i} \lambda_{ij}^{-1} \mathbf{y}_{ij} + \mathbf{T}_{\gamma_i}^{-1} \delta_{\gamma_i} \right], \left[\sum_{j=1}^{n_i} \lambda_{ij}^{-1} + \mathbf{T}_{\gamma_i}^{-1} \right]^{-1} \right), \\ [\Sigma_i | -] &\sim \text{IW} \left(n_i + v, \sum_{j=1}^{n_i} \lambda_{ij}^{-1} \mathbf{y}_{ij} \text{diag}(\boldsymbol{\psi}_i)^{-1} \times (\mathbf{y}_{ij} - \mu_i)(\mathbf{y}_{ij} - \mu_i)' \text{diag}(\boldsymbol{\psi}_i)^{-1} + \mathbf{A} \right). \end{aligned}$$

In addition, the full conditional for γ_{ij} is

$$[\lambda_{ij}|y] \sim \text{IG}(\frac{1}{2}(\kappa+p), (\frac{1}{2}[(\mathbf{y}_{ij}-\mu_i)' \sum_i^{-1} (\mathbf{y}_{ij}-\mu_i)+\kappa])^{-1}).$$

One can assign a prior to but for simplicity we set $\kappa = 4$. Using the same prior values for the t -residual LC model as those used for the Gaussian residual LC model, we fit the t -residual LC model to the marine mammal dataset. Posterior means and 95% credible intervals for μ_{0k}^* , $k = 1, \dots, p$, are listed in Table 6. The results are fairly comparable to a Gaussian LC model. The average credible interval length across the 15 trace elements turned out to be 1.58 with all 15 point estimates being very similar to those from the Gaussian residual LC model. Using $E(\gamma_{ik}|\mathbf{y}) < 0.1$ as a criterion to classify the k th element of the i th lab as an outlier, there were 29 outlying elements (compared to 27 using the Gaussian residual) and 11 labs with at least one outlying element (compared to 11 using the Gaussian residual).

6. CONCLUSION

We have developed a methodology that does very well in accommodating multivariate outliers in a multilevel/hierarchical modeling framework. Although the presentation of the model is necessarily a bit notation heavy, the fundamental idea to handling multivariate outliers (locally allocating multivariate vector entries to a contamination) could not be more natural and is quite simple and intuitive. In addition to being robust to the presence of outliers, the methodology provides probabilistic inference on lab/element outlier classification. This type of information should be of interest to practitioners. Also, the methodology incorporates the uncertainty associated with lab/element outlier classification in parameter estimation and inference. These nice features are available at a minimal computational cost, as a straightforward Gibbs sampler is all that is required to implement the methodology. Computer code is available from the first author by request.

Acknowledgments

The authors are grateful to Antonio Possolo (Chief of the Statistical Engineering Division of the NIST) with help in gaining access to the marine mammal dataset, to the study's sponsor NOAA/National Marine Fisheries Service, and the Inorganic Chemical Metrology Group (Dr. Gregory Turk, group leader) in the Analytical Chemistry Division (Dr. Stephen Wise, Chief) of the Chemical Science and Technology Laboratory (Dr. Willie May, Director).

References

- Bayarri M, Morales J. Bayesian Measures of Surprise for Outlier Detection. *Journal of Statistical Planning and Inference*. 2003; 111:3–22.
- Bednarski T, Zontek S. Robust Estimation of Parameters in a Mixed Unbalanced Model. *The Annals of Statistics*. 1996; 24:1493–1510.
- Box G, Tiao G. A Bayesian Approach to Some Outlier Problems. *Biometrika*. 1968; 55:119–129. [PubMed: 5661037]
- Christopher SJ, Pugh RS, Ellis MB, Mackey EA, Spatz RO, Porter BJ, Bealer KJ, Kucklick JR, Rowles TK, Becker PR. Description and Results of the NIST/NOAA 2005 Interlaboratory Comparison Exercise for Trace Elements in Marine Mammals. *Accreditation and Quality Assurance*. 2007; 12:175–187.
- Dunson DB. Nonparametric Bayes Local Partition Models for Random Effects. *Biometrika*. 2009; 96:249–262. [PubMed: 23710074]
- Gelman, A.; Carlin, JB.; Stern, HS.; Rubin, DB. *Bayesian Data Analysis*. 2. New York: Chapman & Hall/CRC; 2004.
- Gleser LJ. Assessing Uncertainty in Measurement. *Statistical Science*. 1998; 13:277–290.
- Hoff PD. Model-Based Subspace Clustering. *Bayesian Analysis*. 2006; 1:321–344.

- Ishwaran H, Zarepour M. Dirichlet Prior Sieves in Finite Normal Mixtures. *Statistica Sinica*. 2002; 12:941–963.
- Lischer, P. Robust Statistical Methods in Interlaboratory Analytical Studies. In: Rieder, H., editor. *Robust Statistics, Data Analysis, and Computer Intensive Methods*. New York: Springer Verlag; 1996. p. 251-265.
- Mandel J. Structure and Outliers in Interlaboratory Studies. *Journal of Testing and Evaluation*. 1995; 23:364–369.
- Muller CH, Uhlig S. Estimation of Variance Components With High Breakdown Point and High Efficiency. *Biometrika*. 2001; 88:353–366.
- Peña D, Prieto FJ. Multivariate Outlier Detection and Robust Covariance Matrix Estimation. *Technometrics*. 2001; 43:286–300.
- Penny KI, Jolliffe IT. A Comparison of Multivariate Outlier Detection Methods for Clinical Laboratory Safety Data. *The Statistician*. 2001; 50:295–308.
- Roberts G, Rosenthal J. Coupling and Ergodicity of Adaptive MCMC. *Journal of Applied Probability*. 2007; 44:458–475.
- Rocke DM. Robust Statistical Analysis of Interlaboratory Studies. *Biometrika*. 1983; 70:421–31.
- Rukhin AL. Estimating Common Vector Parameters in Interlaboratory Studies. *Journal of Multivariate Analysis*. 2007; 98:435–454.
- Rukhin AL, Vangel MG. Estimation of a Common Mean and Weighted Means Statistics. *Journal of the American Statistical Association*. 1998; 93:303–308.
- Song PX-K, Zhang P, Qu A. Maximum Likelihood Inference in Robust Linear Mixed-Effects Models Using Multivariate t Distributions. *Statistica Sinica*. 2007; 17:929–943.
- Toman B. Bayesian Approaches to Calculating a Reference Value in Key Comparison Experiments. *Technometrics*. 2007; 29:81–87.
- Verdinelli I, Wasserman L. Bayesian Analysis of Outlier Problems Using the Gibbs Sampler. *Statistics and Computing*. 1991; 1:105–117.

APPENDIX: DERIVATION OF FULL CONDITIONALS

Here we only include details regarding $[\mu_0^*|-]$ and $[\sigma_0^{2*}|-]$ as the derivations of $[\mu_h^*|-]$ and $[\sigma_h^{2*}|-]$ follow similar arguments. We begin with $[\mu_0^*|-]$:

$$\begin{aligned}
 [\mu_0^*|-] &\propto \prod_{i=1}^m \varphi(\mu_i; \delta_{\gamma_i}, \mathbf{T}_{\gamma_i}) \varphi(\delta_{\gamma_i}; \mathbf{m}_0, \mathbf{S}_0) \\
 &\propto \prod_{i=1}^m \exp\{-0.5(\mu_i - \delta_{\gamma_i})' \mathbf{T}_{\gamma_i}^{-1} (\mu_i - \delta_{\gamma_i}) - 0.5(\delta_{\gamma_i} - \mathbf{m}_0)' \mathbf{S}_0^{-1} (\delta_{\gamma_i} - \mathbf{m}_0)\} \\
 &\quad \propto \exp\left\{-0.5 \left(\sum_{i=1}^m \left[-\mu_i' \mathbf{T}_{\gamma_i}^{-1} \right. \right. \right. \\
 &\quad \quad \times \left(\mu_0^* \otimes \gamma_i + \mu_{S_i}^* \otimes (1 - \gamma_i) \right) \\
 &\quad \quad \left. \left. - \left(\mu_0^* \otimes \gamma_i + \mu_{S_i}^* \otimes (1 - \gamma_i) \right)' \mathbf{T}_{\gamma_i}^{-1} \mu_i' \right. \right. \\
 &\quad \quad \left. \left. + \left(\mu_0^* \otimes \gamma_i + \mu_{S_i}^* \otimes (1 - \gamma_i) \right) \right. \right. \\
 &\quad \left. \left. \times \mathbf{T}_{\gamma_i}^{-1} \left(\mu_0^* \otimes \gamma_i + \mu_{S_i}^* \otimes (1 - \gamma_i) \right) \right] \right. \\
 &\quad \left. + \mu_0^{*'} \mathbf{S}_0^{-1} \mu_0^* - \mu_0^{*'} \mathbf{S}_0^{-1} \mathbf{m}_0 - \mathbf{m}_0' \mathbf{S}_0^{-1} \mu_0^* \right\} \\
 &\propto \exp\left\{0.5 \left(\mu_0^{*'} \left[\sum_{i=1}^m \mathbf{T}_{\gamma_i=0}^{-1} + \mathbf{S}_0^{-1} \right] \mu_0^* - \mu_0^{*'} \left[\sum_{i=1}^m \mathbf{T}_{\gamma_i=0}^{-1} \mu_i + \mathbf{S}_0^{-1} \mathbf{m}_0 \right] - \left[\sum_{i=1}^m \mu_i' \mathbf{T}_{\gamma_i=0}^{-1} + \mathbf{m}_0' \mathbf{S}_0^{-1} \right] \mu_0^* \right) \right\} \\
 &\quad \sim N_p \left(\left[\sum_{i=1}^m \mathbf{T}_{0:\gamma_i=0}^{-1} + \mathbf{S}_0^{-1} \right]^{-1} \times \left[\sum_{i=1}^m \mathbf{T}_{0:\gamma_i=0}^{-1} \mu_i + \mathbf{S}_0^{-1} \mathbf{m}_0 \right], \left[\sum_{i=1}^m \mathbf{T}_{0:\gamma_i=0}^{-1} + \mathbf{S}_0^{-1} \right]^{-1} \right).
 \end{aligned}$$

Now $[\sigma_0^{2*}|-]$ can be obtained by

$$\begin{aligned}
[\sigma_{0k}^{2*} | -] &\propto \prod_{i=1}^m \varphi(\mu_i; \delta_{\gamma_i}, \mathbf{T}_{\gamma_i}) \text{IG}(\sigma_{0k}^{2*}; a_{\sigma_k}, b_{\sigma_k}) \\
&\propto (\sigma_{0k}^{2*})^{-0.5 \sum_{i=1}^m \gamma_i - a_{\sigma_k} - 1} \\
&\times \exp \left\{ -0.5 \sum_{i=1}^m \gamma_{ik} (1/\sigma_{0k}^{2*}) (\mu_{ik} - \mu_{0k}^*)^2 - 1/(b_{\sigma_k} \sigma_{0k}^{2*}) \right\} \\
&\sim \text{IG} \left(a_{\sigma_k} + \frac{1}{2} \sum_{i=1}^m \gamma_{ik}, \left[\frac{1}{b_{\sigma_k}} + \frac{1}{2} \sum_{i=1}^m \gamma_{ik} (\mu_{ik} - \mu_{0k}^*)^2 \right]^{-1} \right).
\end{aligned}$$

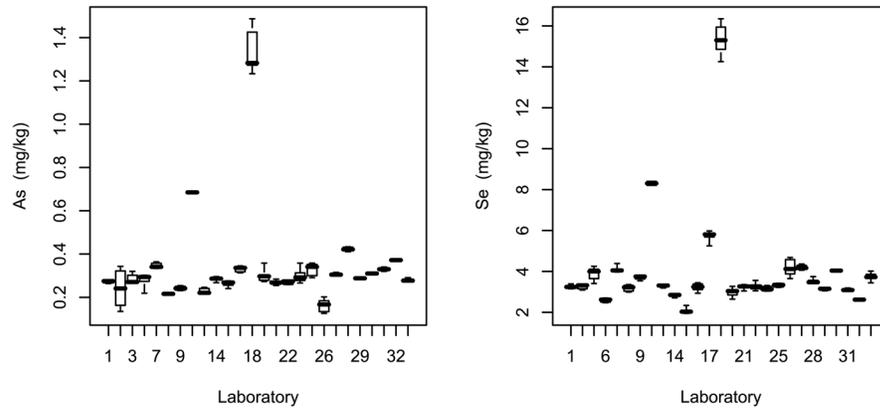


Figure 1. Boxplots for concentrations of As and Se measured in the NIST/NOAA sponsored inter-laboratory study and recorded as mass fraction mg per gram.

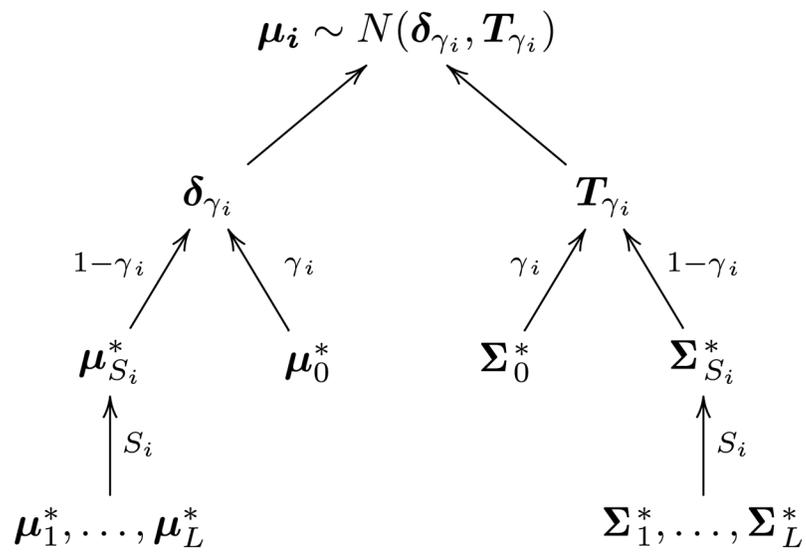


Figure 2. Graphical display of how the mean and covariance matrix of μ_i is constructed. Recall that $S_i = \ell$ with probability ν_ℓ for $\ell = 1, \dots, L$ and $(\delta_{\gamma_i k} = \mu_{0k}^*, t_{\gamma_i k}^2 = \sigma_{0k}^{*2})$ with probability π_i for $i = 1, \dots, m$.

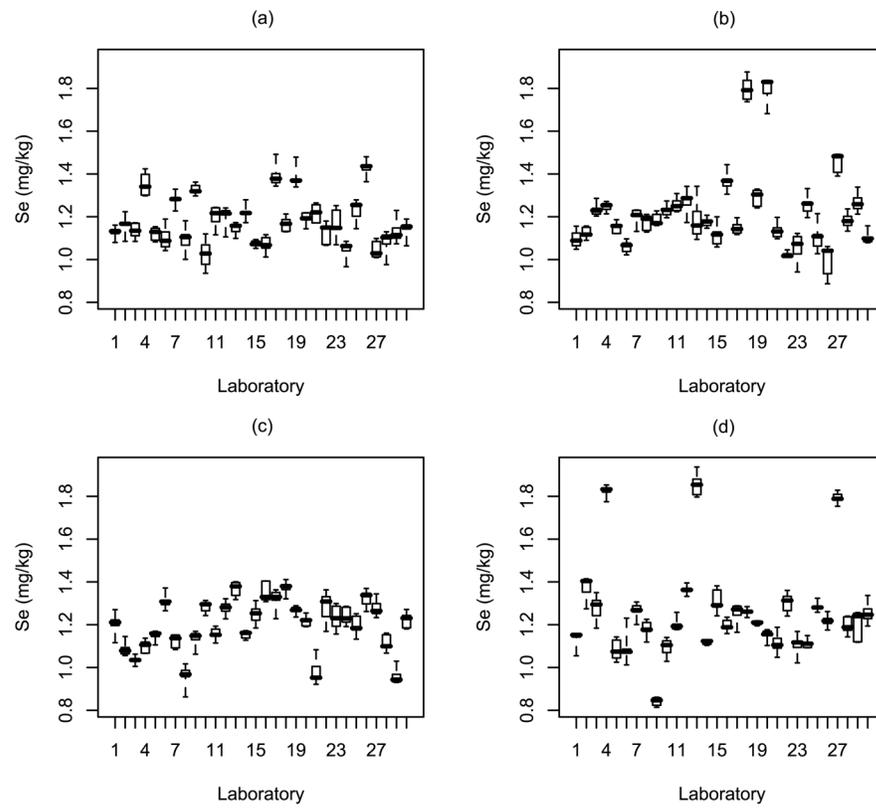


Figure 3. Boxplots of concentrations of Se from a randomly generated dataset used in the simulation study. Plots (a) and (c) contain no outliers while plots (b) and (d) do. The data in plots (a) and (b) are generated using the MVN model and those in plots (c) and (d) are generated using MVT. In plot (b) two of the ten outlier labs have outlying Se values. Three of the ten outlying labs have outlying Se values in plot (d).

Table 1

Results from the simulation study. Entries correspond to averages computed across the 500 generated datasets

True model	Procedure	Coverage	Volume $\times 10^{11}$	Bias	MSE	
MVN	No outliers	MVN model	0.902	0.001	0.302	0.034
		MVT model	0.908	0.002	0.318	0.038
		DP model	0.916	0.003	0.316	0.038
		LC model	0.890	0.001	0.311	0.035
	10 outliers	MVN model	0.890	0.647	0.799	0.156
		MVT model	0.944	0.026	0.410	0.065
MVT		DP model	0.922	1.410	0.722	0.145
		LC model	0.920	0.001	0.312	0.038
	No outliers	MVN model	0.926	1.500	0.406	0.065
		MVT model	0.946	0.005	0.318	0.042
		DP model	0.912	0.224	0.410	0.067
		LC model	0.906	0.007	0.338	0.045
	10 outliers	MVN model	0.856	31.41	0.855	0.194
		MVT model	0.960	0.159	0.443	0.077
		DP model	0.886	19.29	0.796	0.188
		LC model	0.932	0.039	0.375	0.056

Table 2

Results from the simulation study with regard

	Procedure	Clusters occupied	Outlying elements	Outlying labs
MVN	No outliers	2	0	0
	10 outliers	3	29	10
MVT	No outliers	3	10	4
	10 outliers	4	35	12

NOTE: Outlier allocation accuracy and contamination cluster occupancy.

Table 3

Posterior means and 95% credible intervals for the fifteen trace elements measured in the marine mammal inter-laboratory study using the four procedures

	MVN		MVT		DP		LC	
	Mean	95% CI						
Ag	0.50	(0.42, 0.60)	0.48	(0.43, 0.54)	0.52	(0.38, 0.66)	0.47	(0.43, 0.50)
As	0.30	(0.26, 0.34)	0.28	(0.26, 0.30)	0.30	(0.25, 0.35)	0.29	(0.26, 0.31)
Cd	0.25	(0.20, 0.30)	0.22	(0.20, 0.23)	0.26	(0.18, 0.36)	0.22	(0.20, 0.23)
Co	0.03	(0.00, 0.06)	0.02	(0.01, 0.03)	0.03	(0.00, 0.08)	0.01	(0.01, 0.02)
Cs	0.03	(0.02, 0.05)	0.03	(0.02, 0.04)	0.03	(0.02, 0.05)	0.03	(0.03, 0.03)
Cu	5.52	(4.84, 6.16)	5.26	(4.99, 5.53)	5.56	(4.60, 6.48)	5.25	(5.04, 5.45)
Fe	366.2	(333.7, 398.9)	355.3	(344.1, 367.6)	360.6	(325.3, 398.0)	354.7	(344.7, 364.1)
Hg	4.00	(3.42, 4.62)	3.67	(3.42, 3.95)	4.05	(3.08, 5.03)	3.62	(3.43, 3.82)
Mn	3.20	(2.86, 3.55)	3.10	(2.97, 3.23)	3.07	(2.69, 3.45)	3.13	(3.03, 3.23)
Mo	0.64	(0.25, 1.09)	0.42	(0.32, 0.52)	0.41	(0.17, 0.70)	0.39	(0.32, 0.47)
Rb	1.25	(1.14, 1.36)	1.18	(1.13, 1.24)	1.23	(1.09, 1.39)	1.19	(1.14, 1.24)
Se	3.58	(3.13, 4.03)	3.32	(3.09, 3.54)	3.68	(3.07, 4.33)	3.32	(3.09, 3.57)
Sn	0.12	(0.04, 0.21)	0.08	(0.05, 0.11)	0.15	(0.03, 0.30)	0.09	(0.05, 0.13)
V	0.05	(0.04, 0.06)	0.04	(0.04, 0.05)	0.05	(0.04, 0.07)	0.05	(0.04, 0.06)
Zn	29.49	(24.26, 35.50)	30.90	(29.04, 32.77)	30.38	(24.89, 36.07)	30.92	(29.69, 32.17)

Table 4

Fraction of MCMC iterates for which the number of contamination clusters were occupied

	Number of occupied contamination clusters						
	3	4	5	6	7	8	9
Posterior probability	0.0509	0.2717	0.3888	0.2253	0.0564	0.0065	0.0004

Table 5

The posterior probability of each laboratory being part of the majority component ($E(\pi_i|\mathbf{y})$). The values in parentheses are the number of trace elements for each lab whose posterior means of γ_{ik} are less than 0.1, 0.5, and 0.9. The posterior mean of γ_{ik} is the probability that the k th element corresponding to the i th lab is part of the majority component

Lab	$E(\pi_i \mathbf{y})$
1	0.98 (0, 0, 0)
2	0.98 (0, 0, 0)
3	0.98 (0, 0, 0)
4	0.97 (0, 0, 0)
5	0.96 (0, 0, 0)
6	0.97 (0, 0, 0)
7	0.97 (0, 0, 0)
8	0.97 (0, 0, 0)
9	0.97 (0, 0, 0)
10	0.97 (0, 0, 0)
11	0.68 (4, 5, 14)
12	0.92 (1, 1, 2)
13	0.97 (0, 0, 0)
14	0.97 (0, 0, 0)
15	0.87 (1, 1, 6)
16	0.97 (0, 0, 0)
17	0.89 (1, 2, 4)
18	0.46 (12, 13, 14)
19	0.91 (1, 1, 3)
20	0.97 (0, 0, 0)
21	0.97 (0, 0, 0)
22	0.98 (0, 0, 0)
23	0.90 (1, 2, 3)
24	0.98 (0, 0, 0)
25	0.98 (0, 0, 0)
26	0.82 (1, 2, 13)
27	0.96 (0, 0, 1)

Lab	$E(\pi_{ij y})$
28	0.75 (3, 5, 9)
29	0.98 (0, 0, 0)
30	0.92 (1, 1, 1)
31	0.97 (0, 0, 1)
32	0.92 (1, 1, 2)
33	0.94 (0, 0, 1)

Table 6

Posterior means and 95% credible intervals for the fifteen trace elements measured in the marine mammal inter-laboratory study using 15 independent univariate analyses, t -residual LC model, and the normal residual LC model

	Univariate		t -residual LC model		LC model	
	Mean	95% CI	Mean	95% CI	Mean	95% CI
Ag	0.48	(0.45, 0.53)	0.46	(0.44, 0.48)	0.47	(0.43, 0.50)
As	0.34	(0.26, 0.43)	0.29	(0.27, 0.31)	0.29	(0.26, 0.31)
Cd	0.26	(0.21, 0.30)	0.22	(0.21, 0.23)	0.22	(0.20, 0.23)
Co	0.02	(0.01, 0.04)	0.01	(0.01, 0.02)	0.01	(0.01, 0.02)
Cs	0.03	(0.02, 0.05)	0.03	(0.03, 0.03)	0.03	(0.03, 0.03)
Cu	5.44	(4.89, 6.35)	5.18	(5.01, 5.35)	5.25	(5.04, 5.45)
Fe	355.8	(333.8, 378.7)	354.8	(344.7, 364.3)	354.6	(344.7, 364.1)
Hg	3.71	(3.36, 4.14)	3.61	(3.44, 3.79)	3.62	(3.43, 3.82)
Mn	3.16	(2.94, 3.39)	3.10	(3.01, 3.19)	3.13	(3.03, 3.23)
Mo	0.41	(0.38, 0.45)	0.39	(0.34, 0.45)	0.39	(0.32, 0.47)
Rb	1.35	(1.20, 1.53)	1.18	(1.12, 1.22)	1.19	(1.14, 1.24)
Se	3.99	(3.06, 5.05)	3.35	(3.09, 3.60)	3.32	(3.09, 3.57)
Sn	0.11	(0.05, 0.19)	0.07	(0.05, 0.09)	0.09	(0.05, 0.13)
V	0.05	(0.04, 0.07)	0.05	(0.04, 0.05)	0.05	(0.04, 0.06)
Zn	30.03	(27.84, 32.32)	30.96	(29.76, 32.14)	30.92	(29.69, 32.17)