BIOMETRICS 56, 1249–1255
December 2000

A Bayesian Approach to Finite Mixture Models in Bioassay via Data Augmentation and Gibbs Sampling and Its Application to Insecticide Resistance

Pingping Qu

Department of Statistics, University of Kentucky,
817 Patterson Office Tower, Lexington, Kentucky 40506-0027, U.S.A.
email: quping@ms.uky.edu

Yinsheng Qu

Fred Hutchinson Cancer Research Center,
1100 Fairview Avenue North MP-702, Seattle, Washington, 98109, U.S.A.
email: yqu@fhcrc.org

SUMMARY. After continued treatment with an insecticide, within the population of the susceptible insects, resistant strains will occur. It is important to know whether there are any resistant strains, what the proportions are, and what the median lethal doses are for the insecticide. Lwin and Martin (1989, *Biometrics* **45**, 721–732) propose a probit mixture model and use the EM algorithm to obtain the maximum likelihood estimates for the parameters. This approach has difficulties in estimating the confidence intervals and in testing the number of components. We propose a Bayesian approach to obtaining the credible intervals for the location and scale of the tolerances in each component and for the mixture proportions by using data augmentation and Gibbs sampler. We use Bayes factor for model selection and determining the number of components. We illustrate the method with data published in Lwin and Martin (1989).

KEY WORDS: Bayes factor; Insecticide resistance; LD50; Marginal likelihood; Mixture probit model; Partially proper prior; Schwarz criterion; Tolerance distribution.

1. Introduction

Insecticide resistance is a big problem in public health and agriculture (Brown and Pal, 1971). Resistance is a developed attribute consequent upon continued treatment with the insecticide. With the increase of the relative frequency of the resistant strains versus the normal or standard strains, the susceptibility of the species to the insecticide declines. Before a susceptible population becomes resistant completely, the population is a mixture in tolerance distribution. Here tolerance is "defined as the level of intensity below which the response does not occur and above which the response occurs" (Govindarajulu, 1988). The tolerance varies across the population. When the insect population consists of a susceptible strain and several resistant strains, the tolerance distribution is a mixture. A very important question for a mixture model is how many components there are in the mixture. In our example, this question is how many phenotypes there are in respect of susceptibility to the insecticide. Usually the susceptibility is controlled by a single gene. Let A denote the normal susceptible gene and a denote the resistant allele. The genotypes are AA, Aa, and aa. If resistance is dominant or recessive, then only two phenotypes are observed. If the inheritance of resistance is intermediate, then there exist three phenotypes. Therefore, the number of components in the mixture is of scientific interest. However, this question is not easy to answer (McLachlan and Basford, 1988, Chapter 1). Lwin and Martin (1989) propose using the EM algorithm or a direct optimization technique to obtain the maximum likelihood estimates (MLE) for the parameters in each component of the mixture and the mixture proportions. The EM algorithm is believed to be less sensitive to the initial values for convergence than the Gauss–Newton approach. However, for the probit mixture models, to get convergence for a model with three or more components, you have to try a great number of starting values. Most of them do not converge. The standard errors for the MLEs in Lwin and Martin (1989) are given by numerical differentiation (when a direct method is used), which is not accurate (obviously too big in this data set). Another drawback in this likelihood approach is using a goodness-of-fit test to determine the number of components in the mixture, which is not appropriate. Although the bootstrap method for this purpose is valid (Feng and McCulloch, 1996), its computation

is extremely difficult. This is because the likelihood has multiple maxima and the EM algorithm leads to a local maximum if it converges. Therefore, a grid of many, say 27 (as Feng and McCulloch (1994) suggested for normal mixtures), different starting points are needed for finding the global maximum. For probit mixtures, the bootstrap method is not practical.

In the following, we propose a simple Bayesian sampling approach to obtaining the posterior distributions for the parameters, which provides much richer inference than the likelihood approach. To determine the number of components in the mixture, we fit a series of models with increasing numbers of components and select the most plausible model by computing the Bayes factor. The computation is efficient and is not sensitive to the initial values at all. We illustrate this Bayesian approach by reanalyzing the bioassay data published in Lwin and Martin (1989).

We describe the data in Section 2 and the probability model in Section 3. In Section 4, we give an outline of the procedure of data augmentation (Tanner and Wong, 1987) and Gibbs sampling (Gelfand and Smith, 1990; Diebolt and Robert, 1994) but leave the details of the data augmentation algorithm to the Appendix. Then we present the results of Bayesian analysis of the data, emphasizing the method of model selection in Section 5. We conclude the presentation with a brief discussion on some practical issues in Section 6.

2. The *Ostertagia* spp. Egg Data

The investigators are interested in the effectiveness of the anthelminitic drug thiabendazole (TBZ) on the *Ostertagia* spp. worm eggs from sheep infected with this parasite. After 3 years of treatment with the anthelminitic, some worm eggs have developed resistance to the drug and the egg population is thought to be composed of different proportions of a susceptible strain and one or more resistant strains.

The data are taken from the experiment of year 4, which are published in Lwin and Martin (1989) and are illustrated in Figure 1. There are 18 dose groups. The first dose group is a control group without TBZ. The dose of TBZ increases from 0.0230 μg/ml in the second group to 1.500 μg/ml in the 18th group. For the ith group, d_i denotes the dose of TBZ and m_i and n_i denote the number of deaths and the total number of eggs, respectively. In the following, the number of survived eggs is denoted by l_i, which equals $n_i - m_i$.

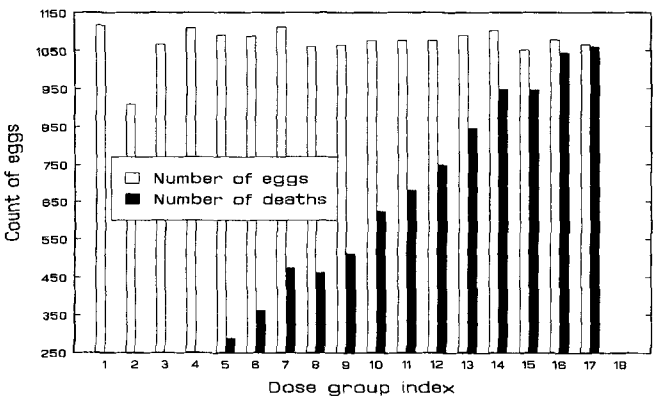

Figure 1. The number of dead and total eggs in the 18 dose groups. The doses of the drug TBZ in the 18 groups from the first to the 18th group are 0, 0.023, 0.047, 0.063, 0.078, 0.094, 0.125, 0.156, 0.188, 0.250, 0.313, 0.375, 0.500, 0.625, 0.750, 1.000, 1.250, and 1.500 μg/ml, respectively.

3. Bayesian Hierarchical Model

3.1 *The Mixture Probit Model*

In the control group, where no TBZ is applied, 79 out of 1117 eggs died. This is called the natural or control mortality (Morgan, 1992, Chapter 3). In the other groups with different dose levels, there are two kinds of mortality: natural mortality and drug-induced mortality. Let λ denote the natural mortality and ψ denote the drug-induced mortality. Then the overt mortality, H, is given by Abbott's formula, $H = \lambda + (1 - \lambda)\psi$ (Abbott, 1925).

We assume that the tolerance distribution is a mixture of k components and, in each component, the tolerance distribution is normal (Finney, 1978). This is a probit mixture model. The jth component has proportion w_j, and the tolerances in this component have a normal distribution with mean μ_j and variance σ_j^2 $(j = 1, \ldots, k)$. Let $x_i = \log d_i$ denote the dose level in the logarithmic scale for the ith dose group. Then the probability of death caused by the drug in the ith dose group and the jth component, ψ_{ij}, equals $\Phi\{(x_i - \mu_j)\sigma_j^{-1}\}$, where $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution. Obviously, μ_j is the median lethal dose (LD50), in the log scale, for the jth component at which 50% eggs exposed to the drug are expected to die.

The overt mortality in the jth component and the ith dose group H_{ij} equals $\lambda + (1-\lambda)\Phi\{(x_i - \mu_j)\sigma_j^{-1}\}$ for $i = 2, \ldots, 18$.

3.2 *The Prior*

For the mixing proportions, $(w_1, \ldots, w_k)$, we use the Dirichlet$(1, \ldots, 1)$ prior distribution. For the natural mortality rate, λ, we use the beta$(1, 1)$ prior distribution.

For the tolerance means, μ, and the standard deviations, σ, we use the partial prior proposed by Roeder and Wasserman (1997) and utilized by Carroll, Roeder, and Wasserman (1999). It does not need subjective input for μ and σ but restricts the order of the tolerance means so that $\mu_1 < \mu_2 < \cdots < \mu_k$ to avoid unidentifiability by label switching (Redner and Walker, 1984).

For the mixture model with unequal tolerance variances, we use a common scale parameter, A, to shrink the σ_j toward a common value and choose the improper prior $\pi(A) \propto A^{-1}$. Conditional on A, the component precision $\tau_j = \sigma_j^{-2}$ follows the gamma distribution $\Gamma(\alpha = \nu/2, \beta = A/2)$. We choose the hyperparameter $\nu = 1$.

For the mixture model with equal tolerance variance, we use the gamma distribution for the prior of the common precision, $\pi(\tau \mid A) = \Gamma(\alpha = \nu/2, \beta = A/2)$.

For the tolerance mean given precision τ, we use a Markov prior,

$$\pi(\mu \mid \tau) \propto \pi(\mu_1 \mid \tau)\pi(\mu_2 \mid \mu_1, \tau), \ldots, \pi(\mu_k \mid \mu_{k-1}, \tau),$$

where $\pi(\mu_1 \mid \tau) \propto 1$ and

$$\pi(\mu_j \mid \mu_{j-1}, \tau) = N^+_{\mu_{j-1}}\left\{\mu_{j-1}, 2M^2(\tau_{j-1} + \tau_j)^{-1}\right\}.$$

Here $N_a^+(\cdot,\cdot)$ is a normal distribution truncated to be greater than a and M is a hyperparameter. We choose $M = 5$, as Roeder and Wasserman (1997) suggested.

4. Data Augmentation and Gibbs Sampling

Our algorithm for the data augmentation and Gibbs sampling consists of three steps: generating numbers, generating tolerances, and generating parameters.

4.1 *Generating Numbers of Dead and Survived Eggs for Each Component*

Based on the observed data (x_i, n_i, m_i) and the current estimate of parameters $(\mu, \sigma, w, \lambda)$, we first generate the number of dead eggs, m_{ij}, in dose group i and component j such that $\Sigma_j\, m_{ij} = m_i$ and generate the number of survivors, l_{ij}, such that $\Sigma_j\, l_{ij} = l_i = n_i - m_i$. Then we generate the number of drug-induced deaths, p_{ij}, and the number of natural deaths, $q_{ij} = m_{ij} - p_{ij}$. The augmented numbers for the ith dose group are $\{(l_{ij}, p_{ij}, q_{ij}); j = 1, \ldots, k\}$. For details in data augmentation, see the Appendix.

4.2 *Generating Tolerances*

For the ith dose group and the jth component, we generate l_{ij} tolerances for the survivals from $N_{x_i}^+(\mu_j, \sigma_j^2)$, which are greater than x_i, and generate p_{ij} tolerances for the dead from $N_{x_i}^-(\mu_j, \sigma_j^2)$, which are equal to or less than x_i.

4.3 *Generating the Parameters*

Let $C_j = \Sigma_i\, (l_{ij} + p_{ij} + q_{ij})$ denote the total number of eggs in the jth component. Then the posterior distribution for the mixing proportions, w, is Dirichlet$(C_1 + 1, \ldots, C_k + 1)$. The posterior distribution for the natural death rate, λ, is beta$(Q + 1, P + 1)$, where $Q = \Sigma_i\, \Sigma_j\, q_{ij}$ is the total natural deaths and $P = \Sigma_i\, \Sigma_j\, p_{ij}$ is the total drug-induced deaths.

The component means, μ, and variances, σ^2, in the models with unequal variances are generated using the conditional sampling method proposed by Roeder and Wasserman (1997). For the common variance in the models with equal variances, we sample a scale parameter, A, from the conditional density $p(A \mid \text{rest}) = \Gamma(\alpha = k\nu/2, \beta = \Sigma_j\, \tau_j/2)$ and then we draw $\tau = \sigma^{-2}$ from the gamma distribution with $\alpha = (\nu + \Sigma_j\, N_j)/2$ and $\beta = [A + \Sigma_j\, \{(N_j - 1)s_j^2 + N_j(\bar{y}_j - \mu_j)^2\}]/2$, where $N_j = \Sigma_i\, (l_{ij} + p_{ij})$ is the number of generated tolerances, $\bar{y}_j$ is the sample mean, and s_j^2 the sample variance for the tolerances in the jth component.

5. Analysis of the *Ostertagia* spp. Egg Data

We fit the *Ostertagia* spp. egg data with seven models: a one-component model (M_1), two-component models with equal variance (M_2) and unequal variances $(M_{2'})$, three-component models with equal variance (M_3) and unequal variances $(M_{3'})$, and four-component models with equal variance (M_4) and unequal variances $(M_{4'})$.

For each model, we take 6000 iterations and burn-in the first 1000 iterations. From the remained 5000 samples, we calculate the posterior means and the 95% Bayesian highest posterior density (HPD) intervals using the method proposed by Chen and Shao (1999). The results for the two three-component models, M_3 and $M_{3'}$, are given in Table 1. For the parameters of $M_{3'}$, we also present the maximum likelihood estimates from the EM algorithm in Lwin and Martin (1989) for comparison.

5.1 *Model Selection*

Model specification consists of two steps. In the first step, we formulate a set of candidate models. In the second step, we

Table 1
Posterior means and 95% HPD intervals in three-component models M_3 (equal variance) and $M_{3'}$ (unequal variances) from 5000 samples

Model	Parameter	Mean	95% HPD interval		Length of HPD	MLE[a]
M_3	μ_1	−1.111	−1.138	−1.085	0.053	
	μ_2	−0.580	−0.637	−0.518	0.119	
	μ_3	−0.191	−0.230	−0.152	0.077	
	σ	0.157	0.138	0.179	0.042	
	w_1	0.392	0.359	0.427	0.068	
	w_2	0.307	0.247	0.363	0.117	
	w_3	0.301	0.248	0.360	0.111	
	λ	0.078	0.067	0.088	0.021	
	LD50	0.213	0.207	0.219	0.012	
$M_{3'}$	μ_1	−1.093	−1.139	−1.043	0.095	−1.120
	μ_2	−0.532	−0.632	−0.427	0.205	−0.600
	μ_3	−0.170	−0.244	−0.112	0.132	−0.210
	σ_1	0.176	0.131	0.228	0.097	0.160
	σ_2	0.161	0.089	0.258	0.169	0.190
	σ_3	0.146	0.115	0.178	0.063	0.170
	w_1	0.418	0.354	0.499	0.146	0.380
	w_2	0.324	0.187	0.492	0.305	0.300
	w_3	0.258	0.117	0.379	0.262	0.320
	λ	0.076	0.065	0.088	0.024	0.076
	LD50	0.212	0.206	0.220	0.014	0.212

[a] Maximum likelihood estimate from the EM algorithm in Lwin and Martin (1989).

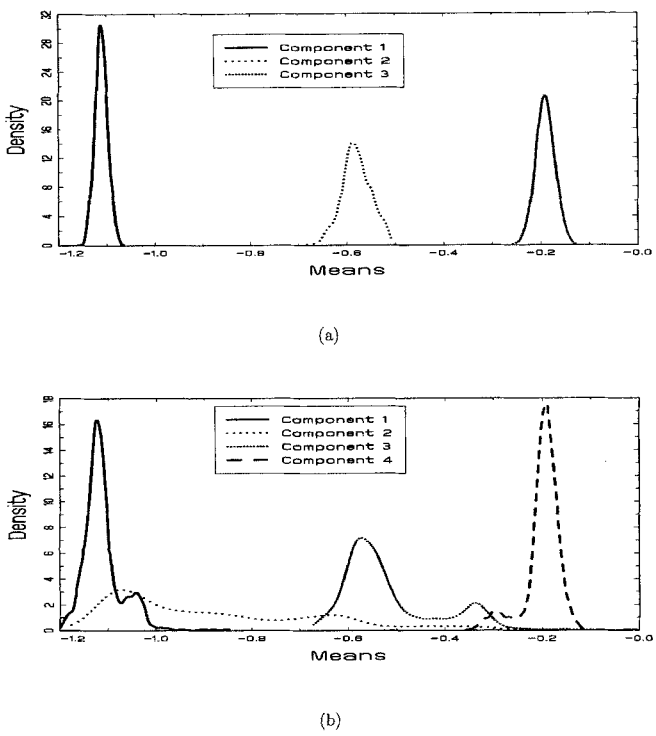

Figure 2. The tolerance means (**a**) under M_3, three components with equal variances, and (**b**) under M_4, four components with equal variances. Note that, under M_4, the tolerance mean in component 2 overlaps with component 1, component 3, and even component 4, indicating that only three components are supported by the data.

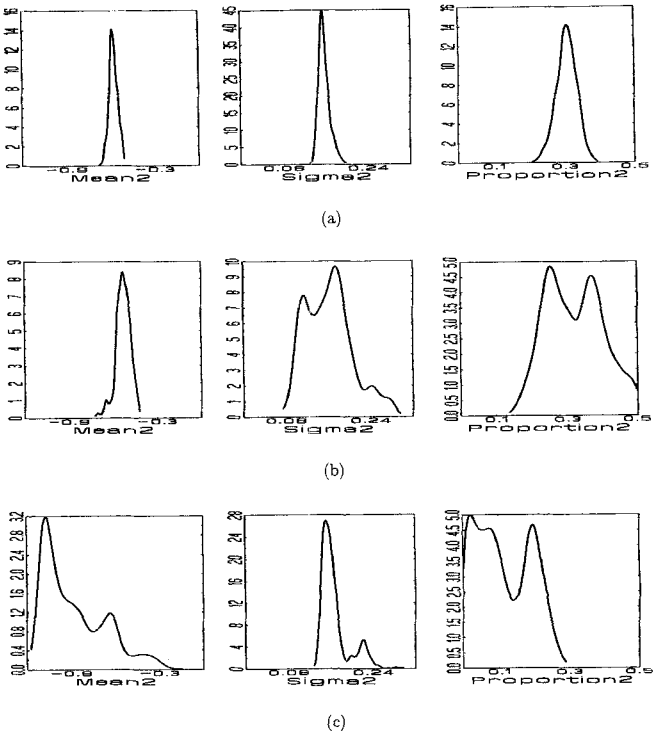

Figure 3. Marginal density curves for μ_2, σ_2, and w_2 (**a**) under M_3, three components with equal variances, (**b**) under $M_{3'}$, three components with unequal variances, and (**c**) under M_4, four components with equal variances. Note that, under $M_{3'}$ and M_4, the density curves are wide with multiple modes, indicating difficulties with convergence, a symptom of overfitting.

select a model or a small number of models to be used in making inferences.

For our data, we have seven candidate models, and no models with more than four components are included. This is because the four-component models, M_4 and $M_{4'}$, have symptoms of overfitting. So models with more than four components are not considered.

We illustrate the symptoms of overfitting in Figures 2 and 3. Figure 2 shows the tolerance means for models M_3 and M_4. Under M_3 (Figure 2a), the three means are well separated, whereas under M_4, μ_2 overlaps with all other components and μ_3 overlaps with μ_4. Obviously, the data support only three components. M_4 is overfitted (Crawford, 1994).

The marginal density curves for the tolerance mean, μ_2, the tolerance standard deviation, σ_2, and the mixing proportion, w_2, for models M_3, $M_{3'}$, and M_4 are shown in Figure 3a, 3b, and 3c, respectively. We see that, under models $M_{3'}$ and M_4, the density curves are widened with multiple modes, indicating difficulties with convergence for the parameters. This is a symptom of overfitting. Also, we find that models $M_{2'}$, $M_{3'}$, $M_{4'}$, and M_4 have nonvanishing autocorrelations, indicating overfitting (Chib, 1995).

For two models M_k and M_j in the candidate set, the evidence in favor of M_k and against M_j is given by the Bayes factor B_{kj}, which is defined as the ratio of posterior to prior odds,

$$B_{kj} = \frac{p(M_k \mid D)}{p(M_j \mid D)} \div \frac{\pi(M_k)}{\pi(M_j)} = \frac{p(D \mid M_k)}{p(D \mid M_j)} = \exp\{\log p(D \mid M_k) - \log p(D \mid M_j)\}.$$

Here $\pi(M_k)$ is the prior probability and $p(D \mid M_k)$ the marginal likelihood for model M_k, which is defined as the mean likelihood averaging over the range of the parameter θ_k (Kass and Raftery, 1995),

$$p(D \mid M_k) = \int L(\theta_k \mid D, M_k)\pi(\theta_k)d\theta_k,$$

where $L(\theta_k \mid D, M_k)$ is the likelihood function. The marginal likelihoods are usually approximated by either the Laplace method, $p(D \mid M_k) \approx (2\pi)^{d_k/2}|\Sigma|^{1/2}L(\tilde{\theta}_k \mid D, M_k)\pi(\tilde{\theta}_k)$, where $\tilde{\theta}_k$ is the posterior mode, d_k is the dimension of θ_k, and Σ is asymptotically equal to the posterior covariance matrix evaluated at $\tilde{\theta}_k$; or the Schwarz criterion (Schwarz, 1978), which is also known as the Bayes information criterion (BIC), i.e., $p(D \mid M_k) \approx n^{-d_k/2}L(\hat{\theta}_k \mid D, M_k)$, where $\hat{\theta}_k$ is the MLE under M_k and n is the sample size. We use BIC to approximate the marginal likelihood.

The posterior probability $p(M_k \mid D)$ is given by

$$\hat{p}(M_k \mid D) = \frac{\hat{p}(D \mid M_k)\pi(M_k)}{\sum_u \hat{p}(D \mid M_u)\pi(M_u)}.$$

Table 2
Values of log-marginal likelihood and probabilities for different models

Model M_j	$\log \hat{p}(D \mid M_j)$	$\hat{p}(M_j \mid D)$
M_1 : one component	−8733.980	0.000
M_2 : two components, equal variances	−8671.355	0.007
$M_{2'}$: two components, unequal variances	−8672.524	0.002
M_3 : three components, equal variances	−8666.462	0.904
$M_{3'}$: three components, unequal variances	−8669.612	0.039
M_4 : four components, equal variances	−8669.410	0.047
$M_{4'}$: four components, unequal variances	−8672.927	0.001

We assume that the prior probabilities for the seven models are equal. The estimated log-marginal likelihoods and the posterior probabilities for different models are listed in Table 2.

Under M_1, the log-marginal likelihood is −8733.980, whereas under M_2, the log-marginal likelihood is −8671.355. Thus, the Bayes factor B_{21} in favor of M_2 and against M_1 is exp(62.625), which is very strong evidence (Raftery, 1996). The Bayes factor B_{32} in favor of the three-component model M_3 against the two-component model M_2 is exp(4.893) = 133.353, which is strong evidence. For the four-component model M_4, the log-marginal likelihood reduces to −8669.410. The Bayes factor B_{34} is 19.068, which is positive evidence against M_4. Finally, the Bayes factor $B_{33'}$ in favor of the three-component model with common variance M_3 against the three-component model with different variances $M_{3'}$ equals 23.336, which is positive evidence against model $M_{3'}$.

5.2 *Statistical Conclusions of the Analysis*

We select the three-component model with equal variance, M_3. This result suggests that the resistance is inherited additively, i.e., the extremely resistant strain (component 3) corresponds to the homozygous genotype aa with double resistance mutant gene a, and the medium resistant strain (component 2) corresponds to the heterozygous genotype Aa with one resistance gene. Also, the mutation changes the location of the tolerance, not the scale. One plausible interpretation is that the mutant gene produces an unknown substance that neutralizes the insecticide to a certain extent and, consequently, shifts the tolerance to the right side for a certain distance.

In Figure 2a, we see that the three components are well separated. The leftmost component, corresponding to the susceptible or normal strain, has a median lethal dose ($LD50_1$) equal to $10^{-1.111} = 0.077$ μg/ml. The other two components, corresponding to the two resistant strains, have $LD50_2 = 0.263$ μg/ml for the moderately resistant strain and $LD50_3 = 0.644$ μg/ml for the severely resistant strain. The proportions of the three components are about 0.4, 0.3, and 0.3. The overall LD50 defined by $\log(LD50) = \Sigma_j w_j \mu_j$ is 0.213 μg/ml. The 95% Bayesian highest posterior density interval for LD50 is (0.207, 0.219). The posterior mean for the natural mortality is 0.078, with the 95% HPD interval equal to (0.067, 0.088) (Table 1).

6. Discussion

The Bayesian analysis enriches the statistical inferences and, by using data augmentation and Gibbs sampling, the computation is efficient and easy to implement. However, many problems have to be solved in order to obtain a satisfactory sampler.

The first problem is how to sample the latent data for the probit model. We have tried the Albert and Chib (1993) approach to sampling the latent data for binary response. For the ith dose group and the jth component, the drug-induced mortality can be written as $\psi_{ij} = \Phi(a_j + b_j x_i)$. According to Albert and Chib (1993), given (a_j, b_j), we sample the latent data from the truncated normal distribution $N_0^+(a_j + b_j x_i, 1)$ (greater than zero) for the dead and from $N_0^-(a_j + b_j x_i, 1)$ (less than zero) for the survivor. Then we update (a_j, b_j) as regression parameters in the linear models. We find that this sampler has a high autocorrelation and consequently a slow mixing rate. After reparameterizing the probit model to $\psi_{ij} = \Phi\{(x_i - \mu_j)\sigma_j^{-1}\}$, we sample the tolerances from $N_{x_i}^-(\mu_j, \sigma_j^2)$ for the dead and from $N_{x_i}^+(\mu_j, \sigma_j^2)$ for the survivors. We update the tolerance mean and variance as observed normal data. This time, the autocorrelations reduce and the mixing rate increases. (For more details in the reparameterization for the purpose of increasing mixing rate, see Gilks and Roberts (1996).)

Another problem is how to choose the priors. Because prior information is not available and we do not use the improper priors (Diebolt and Robert, 1994), we have tried a variety of proposals, such as the "just proper" priors, the weakly informative priors (Richardson and Green, 1997), and so on. Finally, we choose the partially proper prior proposed by Roeder and Wasserman (1997) and utilized by Carroll et al. (1999). The advantage of this prior is that it does not require subjective input for the μ_j and σ_j, yet the posterior is proper. Also, in this approach, only two hyperparameters need to be specified. So we choose this. However, if prior information is available, informative priors are preferred. (See Racine et al. (1986) for a detailed discussion.)

Another issue is how to compute the marginal likelihood and the Bayes factor for model selection. Which approximation methods should we use—the Laplace method, the BIC, or the harmonic mean estimator (Newton and Raftery, 1994)? We have tried all of the methods with real and simulated data. Our experience is that, in practice, BIC gives the best approximate results as long as the priors are not strong. Roeder and Wasserman (1997) conduct simulation studies and conclude that BIC does provide very good approximation to the marginal likelihood. Our experience indicates that, even

if the model is overfitted, BIC is still good for estimating the marginal likelihood and calculating the Bayes factor. This is because, when there is overfitting, the likelihood quickly stabilizes even though the parameters may never reach convergence.

Finally, the models showing overfitting symptoms are not suitable for making inferences about the parameters because their posterior marginal densities are remarkably widened with multiple modes. For example, the three-component model with unequal variances, $M_{3'}$, has much wider marginal density curves (Figure 3b) and 95% HPD intervals (Table 1) than the three-component model with equal variances M_3. Therefore, $M_{3'}$ should not be used for making inferences.

Acknowledgements

We wish to thank the editor, the associate editor, and the reviewer for their valuable comments and suggestions.

Résumé

Après un traitement continu avec un insecticide dans la population des insectes susceptibles, des souches résistantes apparaissent. Il est important de savoir s'il y a des souches résistantes, en quelles proportions, et quelles sont les doses léthales médianes pour l'insecticide. Lwin et Martin (1989) proposent un modèle de mélange probit, et utilisent l'algorithme EM pour obtenir les estimations du maximum de vraisemblance pour les paramètres. Cette aproche a des difficultés à estimer les intervalles de confiance, et à tester le nombre de composantes. Nous proposons une approche bayesienne dans chaque composante, et pour les proportions du mélange, en utilisant l'augmentation de données et l'échantillonneur de Gibbs. Nous utilisons le facteur de Bayes pour la sélection de modèle et la détermination du nombre de composantes. Nous illustrons la méthode avec des données publiées dans Lwin et Martin (1989).

References

Abbott, W. S. (1925). A method of computing the effectiveness of an insecticide. *Journal of Economic Entomology* **18,** 265–267.

Albert, J. H. and Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* **88,** 669–679.

Brown, A. W. A. and Pal, R. (1971). *Insecticide Resistance in Arthropods,* 2nd edition. Geneva: World Health Organization.

Carroll, R. J., Roeder, K., and Wasserman, L. (1999). Flexible parameter measurement error models. *Biometrics* **55,** 44–54.

Chen, M.-H. and Shao, Q.-M. (1999). Monte Carlo estimation of Bayesian credible and HPD intervals. *Journal of Computational and Graphical Statistics* **8,** 69–92.

Chib, S. (1995). Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association* **90,** 1313–1321.

Crawford, S. (1994). An application of the Laplace method to finite mixture distributions. *Journal of the American Statistical Association* **89,** 259–267.

Diebolt, J. and Robert, C. P. (1994). Estimation of finite mixture distributions through Bayesian sampling. *Journal of the Royal Statistical Society, Series B* **56,** 363–375.

Feng, Z. and McCulloch, C. E. (1994). On the likelihood ratio test statistic for the number of components in a normal mixture with unequal variances. *Biometrics* **50,** 1158–1162.

Feng, Z. and McCulloch, C. E. (1996). Using bootstrap likelihood ratios in finite mixture models. *Journal of the Royal Statistical Society, Series B* **58,** 609–617.

Finney, D. J. (1978). *Statistical Method in Biological Assay,* 3rd edition. London: Charles Griffin.

Gelfand, A. E. and Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* **85,** 398–409.

Gilks, W. R. and Roberts, G. O. (1996). Strategies for improving MCMC. In *Markov Chain Monte Carlo in Practice,* W. R. Gilkes, S. Richardson, and D. J. Spiegelhalter (eds), 89–114. London: Chapman & Hall.

Govindarajulu, Z. (1988). *Statistical Techniques in Bioassay.* Basel: Karger.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association* **90,** 773–795.

Lwin, T. and Martin, P. J. (1989). Probits of mixtures. *Biometrics* **45,** 721–732.

McLachlan, G. J. and Basford, K. E. (1988). *Mixture Models: Inference and Applications to Clustering.* New York: Marcel Dekker.

Morgan, B. J. T. (1992). *Analysis of Quantal Response Data.* London: Chapman and Hall.

Newton, M. A. and Raftery, A. E. (1994). Approximate Bayesian inference by the weighted likelihood bootstraps. *Journal of the Royal Statistical Society, Series B* **56,** 3–48.

Racine, A., Grieve, A. P., Fluhler, H., and Smith, A. F. M. (1986). Bayesian methods in practice. *Applied Statistics* **35,** 93–150.

Raftery, A. E. (1996). Approximate Bayes factors and accounting for model uncertainty in generalised linear models. *Biometrika* **83,** 251–266.

Redner, R. A. and Walker, H. F. (1984). Mixture densities, maximum likelihood, and the EM algorithm. *SIAM Review* **26,** 195–239.

Richardson, S. and Green, P. J. (1997). On Bayesian analysis of mixtures with unknown number of components. *Journal of the Royal Statistical Society, Series B* **59,** 731–792.

Roeder, K. and Wasserman, L. (1997). Practical Bayesian density estimation using mixture of normals. *Journal of the American Statistical Association* **92,** 894–902.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics* **6,** 461–464.

Tanner, M. A. and Wong, W. H. (1987). The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* **82,** 528–546.

Received July 1999. *Revised May* 2000.
Accepted May 2000.

APPENDIX

Data Augmentation

Computing the Probability of Allocation in Each Component

From the current estimates of the parameters, μ_j, σ_j^2, w, and λ, we first calculate the overt death rate, $H_{ij} = \lambda + (1 - \lambda)\Phi\{(x_i - \mu_j)\sigma_j^{-1}\}$, for the ith dose group and the jth subpopulation. Then we compute the probability that a dead egg is from the jth component,

$$\xi_{ij|d} = \frac{w_j H_{ij}}{\sum_u w_u H_{iu}},$$

and the probability that a survivor is from the jth component,

$$\xi_{ij|s} = \frac{w_j(1 - H_{ij})}{\sum_u w_u(1 - H_{iu})}.$$

Generating the Number of Deaths and Survivals for Each Component

Letting $\xi_{i|d} = (\xi_{i1|d}, \ldots, \xi_{ir|d})'$ and $\xi_{i|s} = (\xi_{i1|s}, \cdots, \xi_{ir|s})'$, we generate the number of deaths, m_{ij}, from the multinomial distribution multin$(m_i, \xi_{i|d})$ the number of survivors, l_{ij}, from the multinomial distribution multin$(l_i, \xi_{i|s})$.

Generating the Number of Natural Deaths

We generate the number of drug-induced deaths, p_{ij}, from the binomial distribution bin(m_{ij}, p), where $p = (1 - \lambda)\psi_{ij}/H_{ij}$ is the probability that a death is caused by the drug. The number of natural deaths is $q_{ij} = m_{ij} - p_{ij}$.