

# Bayesian Inferences on Umbrella Orderings

Chris Hans<sup>1,\*</sup> and David B. Dunson<sup>2</sup>

November 3, 2004

<sup>1</sup> Institute of Statistics and Decision Sciences, Duke University, Durham, NC

<sup>2</sup> Biostatistics Branch, National Institute of Environmental Health Sciences, MD A3-03,

P.O. Box 12233, Research Triangle Park, NC 27709

\* *email*: [hans@stat.duke.edu](mailto:hans@stat.duke.edu)

**SUMMARY.** In regression applications with categorical predictors, interest often focuses on comparing the null hypothesis of homogeneity to an ordered alternative. This article proposes a Bayesian approach for addressing this problem in the setting of normal linear and probit regression models. The regression coefficients are assigned a conditionally conjugate prior density consisting of mixtures of point masses at zero and truncated normal densities, with a (possibly unknown) changepoint parameter included to accommodate umbrella ordering. Two strategies of prior elicitation are considered: (1) a Bayesian Bonferroni approach in which the probability of the global null hypothesis is specified and local hypotheses are considered independent; and (2) an approach which treats these probabilities as random. A single Gibbs sampling chain can be used to obtain posterior probabilities for the different hypotheses and to estimate regression coefficients and predictive quantities, either by model averaging or under the preferred hypothesis. The methods are applied to data from a carcinogenesis study.

**KEY WORDS:** Bathtub shapes; Downturn order; Gibbs sampler; Mixture prior; Order constraint; Probit model; Shape restriction; Variable selection.

## **1. Introduction**

In many applications, there is interest in assessing evidence of homogeneity against an ordered alternative. For example, in studies of environmental exposures, the focus is often on demonstrating an increasing dose response trend, with a possible downturn at high doses (Simpson and Margolin, 1986). Secondary interests include dose response estimation and inferences on pairwise comparisons and the low observed adverse effect level (LOAEL), corresponding to the lowest dose at which there is an increase relative to background. Ideally, each of these interests could be addressed using a single methodology, while accounting for important covariates and incorporating information from previous studies.

Motivated by applications to bioassay studies, Peddada, Prescott and Conaway (2001)

and Mancuso, Ahn and Chen (2001) independently developed generalizations of the Cochran-Armitage trend test to accommodate order-restricted alternatives without assuming known dose scores. Mancuso et al. (2001) rely on an isotonic regression likelihood ratio test, while Peddada et al. (2001) use the procedure of Hwang and Peddada (1994) to avoid problems with restricted maximum likelihood estimators. For a review of classical order restricted inference, refer to Robertson, Wright and Dykstra (1988).

This article proposes a Bayesian alternative to these approaches. Although it is typically straightforward to incorporate strict constraints in parametric Bayesian models by choosing a prior with support on the restricted space (Gelfand, Smith, and Lee, 1992; Albert, 1994), inferences on uncertain orderings require carefully-specified priors. Much of the Bayesian literature on order restrictions has focused on strict constraints in nonparametric models (Gelfand and Kuo, 1991; Ramgopal, Laud, and Smith, 1993; Lavine and Mockus, 1995; Gelfand and Kottas, 2001; Hoff, 2003). A notable exception is the variable selection approach of Geweke (1996), which allowed for uncertainty in the predictors to be included in a linear model, while placing sign constraints on the regression coefficients.

Considering inferences on simple orderings, Dunson and Neelon (2003) proposed an approximate Bayesian approach based on projecting draws from an unconstrained posterior density to the restricted space. Gunn and Dunson (2003) later modified this approach to allow umbrella restrictions on individual-specific means in a hierarchical model. Such projection approaches have the disadvantage of lacking an explicitly specified prior distribution for the restricted parameters and hypotheses of interest, and the properties of the resulting *quasi*-posterior distribution are still under investigation. In this article, we propose a fully Bayesian approach with explicit priors specified for the parameters and hypotheses of interest. The fully Bayesian approach has advantages when interest focuses on hypothesis testing (as opposed to estimation) in high dimensional cases, as transformation approaches like Gunn and Dunson (2003) can assign very low probability to the null hypothesis of no

difference between groups even if the hypothesis holds.

Our approach uses mixture priors with constrained support for the regression coefficients in a carefully-parameterized linear model. In particular, we specify a product mixture of point masses and truncated normal densities, following a structure related to the variable selection priors proposed by Geweke (1996). Then, to accommodate umbrella orderings, we choose a multinomial prior for the location of a change in the sign restriction. When the data follow a normal linear model, the full conditional posterior densities are conjugate and posterior computation can proceed via a simple blocked Gibbs sampling algorithm. By using the data augmentation strategy of Albert and Chib (1993), this Gibbs sampler can be applied directly when outcomes are binary or ordered categorical.

A related approach was proposed by Dunson and Herring (2003) for Bayesian order restricted inference in proportional hazards regression, though they only considered simple orderings and their prior was structured for survival data. Our prior specification is also related to priors proposed for smooth monotone curve estimation (Neelon and Dunson, 2003) and for variable selection in linear regression (George and McCulloch, 1997; Chipman, George and McCulloch, 2001). However, none of these articles considered inferences on umbrella orderings with an unknown changepoint parameter. Our method is also related to piecewise constant regression models with changepoints, handled in a Bayesian context by Holmes and Heard (2003), who allow for isotonic regression but do not address the problem of a downturn. Our approach differs from the typical Bayesian changepoint problem in the specification of the prior: we place a prior distribution on specific changepoints of interest, rather than the usual approach of placing a prior distribution first on the number of changepoints and then, conditionally, on the locations of the changepoints. As such, we formulate our model in a way that is consonant with the particular scientific interests in the applications we consider.

Using our approach, posterior probabilities for the different hypotheses, a posterior distribution for the LOAEL, and posterior summaries of the parameters, either model-averaged

or under the preferred model, can be calculated from a single run of a Gibbs sampler. We consider two strategies for elicitation of the model space prior, the first based on the *Bayesian Bonferroni* approach (Westfall, Johnson and Utts, 1997) and the second based on choosing a prior for the probability of a local change. We recommend the second approach, as it allows for correlation in the local hypotheses and induces a less severe multiple testing penalty. In this manner our approach is related to that of Gönen, Westfall and Johnson (2003), though they use a very different prior and focus on two-sample multivariate normal data.

Section 2 describes the prior specification in a variety of cases. Section 3 outlines the Gibbs sampling algorithm. Section 4 applies the approach to data from an NTP carcinogenicity study, and Section 5 discusses the results.

## 2. Prior Distributions for Order Restricted Inference

Although the strategy proposed here can be used for a broad class of order restrictions, we focus on the three cases most commonly seen in applications: (i) simple increasing order; (ii) simple tree order; and (iii) umbrella order with an unknown peak location. Change in direction of the ordering (e.g., from increasing to decreasing) is trivial, so we focus on one case for sake of brevity. We initially assume that the response variable,  $y_i$ ,  $i = 1, \dots, n$ , is normally distributed, where  $n$  is the number of subjects, and that we have a single categorical predictor,  $w_i \in \{1, \dots, k\}$ . In particular, let

$$y_i = \mu_{w_i} + \epsilon_i = \alpha + \sum_{j=1}^{w_i-1} \beta_j + \epsilon_i, \quad (1)$$

where  $\mu_j = \mathbb{E}(y_i | w_i = j)$ ,  $\alpha$  is an intercept,  $\boldsymbol{\beta} = (\beta_1, \dots, \beta_{k-1})'$  are increments on the mean, and  $\epsilon_i \sim N(0, \sigma^2)$ . Modifications to allow multiple predictors and/or binary outcomes are straightforward, as we describe in Section 2.4.

### 2.1 Simple Increasing Order

Let  $H_{0j} : \mu_j = \mu_{j+1}$  denote the local null hypothesis of equality between the  $j$ th and  $(j+1)$ st

groups, let  $H_{1j} : \mu_j < \mu_{j+1}$  denote the one-sided alternative, and let  $\mathbf{h} = (h_1, \dots, h_{k-1})'$  be an indicator vector indexing the ordering between the adjacent groups:  $h_j = 1(\beta_j > 0)$  for  $j = 1, \dots, k-1$ . For non-negative  $\beta_j$ , the global null hypothesis of homogeneity,  $H_0 : \mu_1 = \dots = \mu_k$ , corresponds to  $\mathbf{h} = \mathbf{0}$ .

To enforce the non-decreasing means restriction and allocate probability to the different possibilities for  $\mathbf{h}$ , we choose a prior distribution for  $\boldsymbol{\beta}$  which restricts  $\beta_j \geq 0$  while allocating non-zero probability mass to the boundary value of  $\beta_j = 0$ . In particular, let

$$\pi(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \text{N}(\boldsymbol{\alpha}; \boldsymbol{\alpha}_0, \sigma_0^2) \prod_{j=1}^{k-1} \text{ZI-N}^+(\beta_j; \pi_{0j}, m_j, s_j^2), \quad (2)$$

where  $\text{ZI-N}^+(\pi_{0j}, m_j, s_j^2)$  denotes the zero-inflated (ZI) positive normal density consisting of the mixture of a point mass at zero (with probability  $\pi_{0j}$ ) and a  $\text{N}(m_j, s_j^2)$  density truncated below by zero,

$$\text{ZI-N}^+(\beta_j; \pi_{0j}, m_j, s_j^2) = \pi_{0j} \delta_0(\cdot) + (1 - \pi_{0j}) \frac{1(\beta_j > 0) \text{N}(\beta_j; m_j, s_j^2)}{\int_0^\infty \text{N}(z; m_j, s_j^2) dz},$$

where  $\delta_0(\cdot)$  is the degenerate distribution with all of its mass at zero.

Under prior (2), we have  $\Pr(H_{0j}) = \pi_{0j}$  and  $\Pr(H_0) = \prod_{j=1}^{k-1} \pi_{0j}$ , which implies *a priori* independence of the local null hypotheses conditional on  $\{\pi_{01}, \dots, \pi_{0,k-1}\}$ . A reasonable simplifying assumption is to assume equal prior probabilities for the local null hypotheses,  $\pi_{0j} = \phi$ , and hence  $\Pr(H_0) = \phi^{k-1}$ . We consider two strategies of elicitation for  $\phi$ : (1) a *Bayesian Bonferroni* approach which fixes  $\Pr(H_0)$  (say at 0.5), leading to  $\phi = \Pr(H_0)^{1/(k-1)}$ ; and (2) an approach treating  $\phi$  as random with hyperprior  $\pi(\phi) = \text{Beta}(a, b)$ .

Elicitation (1) can be overly conservative in practice, because  $\phi$  approaches 1 as the number of dose groups,  $k$ , increases. Elicitation (2) tends to be much less conservative, allowing the data to inform more strongly about the degree of shrinkage towards equality ( $\beta_j = 0$ ). This is accomplished because elicitation (2) induces dependency in the different local hypotheses due to its hierarchical structure: the prior can be written  $\pi(h_j|\phi) = \phi^{1-h_j}(1-$

$\phi^{h_j}$  for  $j = 1, \dots, k - 1$ , which, after integrating over  $\pi(\phi)$ , induces dependencies between the local null hypotheses. For example,  $\Pr(H_{0j}|h_{(j)}) = (a + k - 2 - \sum_{i \neq j} h_i)/(a + b + k - 2)$  is the prior probability of  $H_{0j}$  conditional on  $h_{(j)}$ , the values of all the other local hypotheses;  $H_{0j}$  is thus more likely when many other  $H_{0i}(i \neq j)$  are true.

We specify the hyperparameters  $a$  and  $b$  so that on average  $\Pr(H_0) = 0.5$ , i.e.  $\mathbb{E}[\phi^{k-1}] = 0.5$ , implying that the hyperparameters must satisfy

$$0.5 = \int_0^1 \phi^{k-1} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \phi^{a-1} (1-\phi)^{b-1} d\phi = \frac{\Gamma(a+b)\Gamma(a+k-1)}{\Gamma(a)\Gamma(a+k-1+b)}. \quad (3)$$

We impose the constraint that  $a + b = 1$ , representing unit information in the prior and allowing us to solve (3) numerically for  $a$  and  $b$ . Each of the prior distributions described above is conditionally conjugate, and posterior sampling is described in Section 3.1.

## 2.2 Simple Tree Order

When a researcher is interested in comparing a control or reference group to several treatment groups, and the treatments are known *a priori* to not reduce the mean response, the alternative hypothesis of interest may be:

$$H_1 : \mu_1 \leq \mu_j, \text{ for } j = 1, \dots, k - 1, \text{ with at least one strict inequality,} \quad (4)$$

where  $j = 1$  indexes the control or reference group and  $H_0$  is the null hypothesis as described in Section 2.1. For this alternative hypothesis, which is known as simple tree order, one does not assume an ordering among the treatment groups.

It is convenient to reparameterize (1) so that  $\mathbb{E}(y_i | w_i = j) = \alpha$  if  $j = 1$  and  $\mathbb{E}(y_i | w_i = j) = \alpha + \beta_{j-1}$  if  $j > 1$ . Under this parameterization, expression (4) is equivalently expressed in terms of the constraint on  $\beta$  described in Section 2.1. Hence, the prior density can be chosen exactly as in (2). Conditional conjugacy also holds for simple tree order, and testing of the global null hypothesis and sub-hypotheses,  $H_{0j} : \mu_1 = \mu_j$ , for  $j = 2, \dots, k$ , can proceed in a manner parallel to that used in the simple order case.

### 2.3 Umbrella Order

The alternative hypothesis of umbrella ordering in the means can be expressed as

$$H_1 : \bigcup_{\tau=1}^k \mu_1 \leq \dots \leq \mu_\tau \geq \mu_{\tau+1} \geq \dots \geq \mu_k, \quad (5)$$

with at least one strict inequality, where  $\tau \in \{1, \dots, k\}$  is an unknown discrete-valued peak parameter. Reparameterizing expression (1) so that

$$\mathbb{E}(y_i | w_i, \tau) = \alpha + \sum_{j=1}^{w_i-1} (-1)^{1(j \geq \tau)} \beta_j, \quad (6)$$

we can use prior (2) for  $(\alpha, \boldsymbol{\beta}')$  exactly as we did in the simple order and simple tree order cases, with no additional complications in the case where  $\tau$  is known. When  $\tau$  is unknown, we need to also choose a prior for the peak location. Assuming *a priori* independence between  $(\alpha, \boldsymbol{\beta}')$  and  $\tau$ , we choose

$$\pi(\tau) = \text{Multinomial}\left(1, \dots, k; p_{01}, \dots, p_{0k}\right). \quad (7)$$

In the special case where  $\mathbf{p}_0 = (0, \dots, 0, 1)'$ ,  $H_1$  is equivalent to the simple increasing order hypothesis discussed in Section 2.1. Similarly,  $\mathbf{p}_0 = (1, 0, \dots, 0)'$  is equivalent to simple decreasing order. In the case where we wish to allow for a downturn, we regard  $p_{0j}$  as the conditional prior probability of a downturn given a dose effect at level  $j$ ,

$$\Pr(\text{Downturn at dose } j | \beta_j > 0) = \Pr(\tau = j | \beta_j > 0) = \Pr(\tau = j) = p_{0j}. \quad (8)$$

Posterior evidence of a downturn can be assessed in a similar manner, by first assessing evidence in favor of a dose effect at level  $j$  and then, conditionally on the existence of a dose effect, computing the posterior probability of  $\tau = j$ .

If the investigator is willing to assume simple ordering, but does not know *a priori* whether the order is increasing or decreasing (a common scenario in applications), they can choose  $\mathbf{p}_0 = (p_{01}, 0, \dots, 0, p_{0k})'$ , where  $p_{01} + p_{0k} = 1$ . In this important special case, the

alternative hypothesis is  $H_1 : \mu_1 \leq \dots \leq \mu_k$  or  $\mu_1 \geq \dots \geq \mu_k$ , with at least one strict inequality.

In applications such as bioassay experiments, one may want to choose the hyperparameters  $\mathbf{p}_0$  to express the evidence accrued over repeated studies that dose-response curves tend to be isotonic, with some chance of a downturn at high doses. Alternately one could choose  $\mathbf{p}_0$  based on the frequency of downturn orders observed in historical studies. The resulting inferences should be most sensitive to those dose-response shapes most commonly observed.

#### 2.4 Multiple Predictors and Generalizations

In many applications, it is important to adjust for covariates  $\mathbf{x}_i = (x_{i1}, \dots, x_{iq})'$  in conducting inferences on the association between  $y_i$  and  $w_i$ . To extend the analysis, we add the term  $\mathbf{x}_i' \boldsymbol{\alpha}$  to the linear regression model (1) (or to any of the reparameterizations). Under the assumption of *a priori* independence between  $\boldsymbol{\alpha}$  and  $\boldsymbol{\beta}$  (and  $\tau$  in the umbrella ordering case), this generalization is straightforward to implement.

First consider the case where one wishes to avoid order constraints on the regression coefficients,  $\boldsymbol{\alpha}$ . In this case, we choose a conditionally conjugate  $N(\boldsymbol{\alpha}_0, \boldsymbol{\Sigma}_0)$  prior for  $\boldsymbol{\alpha}$ . Under this structure, the full conditional distribution of  $\boldsymbol{\alpha}$  given  $\boldsymbol{\beta}$  (and potentially  $\tau$ ) follows a standard multivariate normal form, assuming a conjugate prior is used for  $\sigma^2$ .

Now suppose that instead of a single ordered categorical predictor,  $w_i$ , we have a vector of ordered categorical predictors,  $\mathbf{w}_i = (w_{i1}, \dots, w_{ir})'$ . For predictors where order restrictions are not deemed appropriate, one can simply include a vector of dichotomous indicators within  $\mathbf{x}_i$ . Hence, we focus on the case where  $\mathbf{w}_i$  consists of predictors having order constraints. In this case, we express the mean of  $y_i$  as an additive function, with the term for each  $w_{il}$  following the form described in Sections 2.1 - 2.3 (depending on the nature of the order constraint) and with  $\mathbf{x}_i' \boldsymbol{\alpha}$  added to the sum. Letting  $\boldsymbol{\beta}_l$  denote the regression coefficients characterizing the effect of  $w_{il}$ , for  $l = 1, \dots, r$ , we choose the prior  $\pi(\boldsymbol{\beta}) = \prod_{l=1}^r \pi(\boldsymbol{\beta}_l)$ , with

$\pi(\boldsymbol{\beta}_l)$  following the form shown in (2).

Although we have focused on linear regression models for normal response data, the methods apply directly to generalized linear models (McCullagh and Nelder, 1989) by simply using the expressions for  $\mathbb{E}(y_i|x_i, w_i)$  as the linear predictor in an appropriate GLM. For binary and ordered categorical responses, it is convenient to use probit models (Albert and Chib, 1993), because auxiliary variables can be then used to simplify posterior computation.

Although we focus on univariate data having a single changepoint parameter  $\tau$ , which holds for all subjects, hierarchical extensions are conceptually straightforward. For example, in time course gene expression studies, it is reasonable to assume an umbrella pattern for the gene-specific expression trajectories (Peddada et al., 2003), with genes potentially having different peaks. Instead of considering each gene separately, as done in Peddada et al. (2003), we can potentially assume a multinomial prior distribution for the gene-specific peak locations in order to borrow information across genes. Although the model and prior specification generalize automatically to such settings, a major barrier to implementation is the extreme computation involved.

### 3. Posterior Computation

This section proposes a Gibbs sampling algorithm for posterior computation, focusing on probit models for binary outcomes, first with simple increasing order constraints and then with umbrella orderings. Following Albert and Chib (1993), we introduce a latent variable  $\mathbf{z}_{n \times 1} \sim N(\boldsymbol{\eta} = \mathbf{X}\boldsymbol{\alpha} + \mathbf{W}\mathbf{T}\boldsymbol{\beta}, I)$ , where  $\mathbf{W}$  is the design matrix corresponding to the parameterization in (6) and  $\mathbf{T}$  is a  $(k-1) \times (k-1)$  diagonal matrix that encodes the changepoint  $\tau$ :  $\mathbf{T}_{ii} = 1$  if  $\tau > i$  and  $\mathbf{T}_{ii} = -1$  otherwise, with off-diagonal entries equal to zero. If we define  $y_i = 1$  if  $z_i > 0$  and  $y_i = 0$  otherwise, then this formulation results in the probit model,

$$\Pr(y_i = 1 | \mathbf{x}_i, \mathbf{w}_i) = \Phi\left(\mathbf{x}'_i \boldsymbol{\alpha} + \mathbf{w}'_i \mathbf{T} \boldsymbol{\beta}\right). \quad (9)$$

We focus first on the simple increasing order case, where  $\tau$  is fixed at  $k$  (i.e.,  $\mathbf{T}$  is the

identity matrix). The full conditional distributions of  $\boldsymbol{\alpha}$  and  $\boldsymbol{\beta}$  are available in closed form, and sampling proceeds as described in the following subsection.

### 3.1 Gibbs Sampling Algorithm

The Gibbs sampler alternately samples from the full conditional posterior distributions of the model unknowns. After choosing initial values for  $\boldsymbol{\alpha}$  and  $\boldsymbol{\beta}$ , we alternate between the following steps, for  $t = 1, \dots, T$ :

**STEP 1** For  $i = 1, \dots, n$  draw

$$z_i^{[t+1]} \Big| \boldsymbol{\alpha}^{[t]}, \boldsymbol{\beta}^{[t]}, \phi^{[t]}, \mathbf{h}^{[t]}, \mathbf{Y} \sim \begin{cases} N^+(z_i; \eta_i^{[t]}, 1) & \text{if } y_i = 1 \\ N^-(z_i; \eta_i^{[t]}, 1) & \text{if } y_i = 0, \end{cases}$$

$N^+$  and  $N^-$  denoting normal densities truncated below and above by zero, respectively.

**STEP 2** Draw  $\boldsymbol{\alpha}$  from its full conditional distribution

$$\boldsymbol{\alpha}^{[t+1]} \Big| \mathbf{z}^{[t+1]}, \boldsymbol{\beta}^{[t]}, \phi^{[t]}, \mathbf{h}^{[t]}, \mathbf{Y} \sim N(\boldsymbol{\alpha}; \hat{\boldsymbol{\alpha}}^{[t]}, \hat{\boldsymbol{\Sigma}}_\alpha),$$

$$\hat{\boldsymbol{\Sigma}}_\alpha = (\boldsymbol{\Sigma}_0^{-1} + \mathbf{X}'\mathbf{X})^{-1}, \quad \hat{\boldsymbol{\alpha}}^{[t]} = \hat{\boldsymbol{\Sigma}}_\alpha \left( \boldsymbol{\Sigma}_0^{-1}\boldsymbol{\alpha}_0 + \mathbf{X}'(\mathbf{z}^{[t+1]} - \mathbf{W}\boldsymbol{\beta}^{[t]}) \right).$$

**STEP 3** Draw from  $p(\mathbf{h}, \boldsymbol{\beta} | \mathbf{z}^{[t+1]}, \boldsymbol{\alpha}^{[t+1]}, \phi^{[t]}, \mathbf{Y}) = p(\boldsymbol{\beta} | \mathbf{h}, \mathbf{z}^{[t+1]}, \boldsymbol{\alpha}^{[t+1]}, \phi^{[t]}, \mathbf{Y})p(\mathbf{h} | \mathbf{z}^{[t+1]}, \boldsymbol{\alpha}^{[t+1]}, \phi^{[t]}, \mathbf{Y})$  in two steps:

**STEP 3a** Draw the indicator vector from  $p(\mathbf{h} | \mathbf{z}^{[t+1]}, \boldsymbol{\alpha}^{[t+1]}, \phi^{[t]}, \mathbf{Y})$ , which is a multinomial distribution with  $2^{k-1}$  elements (details in Appendix 1).

**STEP 3b** Conditional on the model drawn in STEP 3a, draw the nonzero components of  $\boldsymbol{\beta}^{[t+1]}$  from  $p(\boldsymbol{\beta} | \mathbf{h}^{[t+1]}, \mathbf{z}^{[t+1]}, \boldsymbol{\alpha}^{[t+1]}, \phi^{[t]}, \mathbf{Y})$ , which is a multivariate normal distribution truncated such that each element is positive (details in Appendix 1).

**STEP 4** Draw  $\phi^{[t+1]} \sim \text{Beta}(\phi; a + k - 1 - N^{[t+1]}, b + N^{[t+1]})$ , where  $N^{[t+1]} = \sum_{i=1}^{k-1} h_i^{[t+1]}$ .

This step is unnecessary if  $\phi$  is specified *a priori* via the *Bayesian Bonferroni* approach.

As the number of categories,  $k$ , increases it becomes infeasible to compute the  $2^{k-1}$  probabilities for sampling  $\mathbf{h}$  at each iteration. One could instead sample  $\beta_j$  one at a time or update  $\beta$  in blocks. Such approaches work well in the applications motivating this article.

### 3.2 Generalization to Umbrella Order

Generalizing posterior computation to the umbrella order case with unknown changepoint,  $\tau$ , can be done in several ways. The most straightforward extension would be to add an extra step to the Gibbs sampler described above to sample  $\tau$  conditional on all other parameters. Under the prior specification in (7), the full conditional for  $\tau$  is multinomial and the updated probabilities are easy to compute. Unfortunately the resulting chain is likely to mix very slowly due to high correlation between  $\beta$  and  $\tau$ .

We propose two different sampling schemes which do not suffer from slow mixing, both of which involve a block update of  $(\beta, \mathbf{h}, \tau)$ . The first method is an extension of STEP 3 of the sampler described above. We first draw from the conditional posterior of  $(\mathbf{h}, \tau)$ , which is a multinomial distribution with  $k2^{k-1}$  components, where the factor of  $k$  comes from the possible locations of the changepoint. Given the model defined by  $(\mathbf{h}, \tau)$ , the nonzero components of  $\beta$  are sampled from the distributions described in Appendix 2.

An alternative sampling scheme which simultaneously updates  $\beta$ ,  $\tau$  and  $\mathbf{h}$  but which is practical for large  $k$  is a Metropolis algorithm that uses an independence sampler as a proposal for  $\tau$ . First  $\tau$  is drawn from a multinomial proposal density, perhaps the prior  $\pi(\tau)$  or a normalized profile likelihood for  $\tau$  computed in advance. Once a candidate value of  $\tau$  is drawn,  $\mathbf{h}$  is drawn conditionally on  $\tau$ . Finally, the nonzero components of  $\beta$  are drawn (componentwise or as a block) conditionally on  $\tau$  and  $\mathbf{h}$ . The vector  $(\beta, \mathbf{h}, \tau)$  is then accepted or rejected jointly via a Metropolis ratio.

### 3.3 Computing Posterior Model Probabilities

We can obtain a Monte Carlo estimate of the posterior probability of the null hypothesis by

averaging model indicators across the Gibbs iterations as follows:

$$\widehat{\Pr}(H_0|\mathbf{Y}) = \frac{1}{T-B} \sum_{t=B+1}^T 1(\beta_1^{[t]} = \dots = \beta_{k-1}^{[t]} = 0), \quad (10)$$

where  $B$  is the number of iterations in the burn-in interval. In some instances this estimator can be improved by using Rao-Blackwellization (Gelfand and Smith, 1990; Liu, Wong and Kong, 1994). In particular, instead of (10), we compute

$$\begin{aligned} \widehat{\Pr}_{RB}(H_0|\mathbf{Y}) &= \frac{1}{T-B} \sum_{t=B+1}^T \mathbb{E} \left[ 1(\beta_1 = \dots = \beta_{k-1} = 0) \middle| \boldsymbol{\alpha}^{[t]}, \phi^{[t]}, \mathbf{Y}, \mathbf{z}^{[t]} \right] \\ &= \frac{1}{T-B} \sum_{t=B+1}^T \Pr(\mathbf{h} = (0, \dots, 0) \middle| \boldsymbol{\alpha}^{[t]}, \phi^{[t]}, \mathbf{Y}, \mathbf{z}^{[t]}) \equiv \tilde{\pi}_0^{RB}, \end{aligned} \quad (11)$$

where we call (11) the Rao-Blackwellized estimate of  $\Pr(H_0|\mathbf{Y})$ . The posterior probability of the alternate hypothesis is thus  $\tilde{\pi}_1^{RB} = 1 - \tilde{\pi}_0^{RB}$ . We can use a similar approach to obtain posterior probabilities for sub-null hypotheses of equalities in specific groups. These posterior probabilities can be used as a basis for inference, as we illustrate in Section 4.

### 3.4 Computing the Posterior Distribution of the LOAEL

The posterior distribution of the LOAEL is computed using the probabilities of the  $\mathbf{h}$  and  $\tau$  configurations, which were recorded at each iteration of the Gibbs sampler. For an analysis with no downturn allowed in the model we have

$$\Pr(\text{LOAEL} = j|\mathbf{Y}) = \Pr(\beta_1 = \dots = \beta_{j-1} = 0, \beta_j > 0|\mathbf{Y}),$$

for  $j \leq k-1$ . If we allow for a downturn we must estimate

$$\Pr(\text{LOAEL} = j|\mathbf{Y}) = \Pr(\beta_1 = \dots = \beta_{j-1} = 0, \beta_j > 0, \tau \neq j|\mathbf{Y}),$$

again for  $j \leq k-1$ . Estimates of these quantities can be found using Rao-Blackwellization as described in the previous section.

## 4. Application to Bioassay Data

### 4.1 Data Structure and Model

We apply the approach to data from a National Toxicology Program (NTP) toxicology and tumorigenesis study of tumeric oleoresin (NTP, 1993). Tumeric oleoresin is commonly used as a spice and coloring agent in Indian food and curries, and there is concern about the potential carcinogenic effects. We focus on data on the occurrence of hepatocellular adenomas in female mice, data analyzed previously by Peddada et al. (2001). Mice were randomly assigned to a control group, containing 50 mice, or to one of three treatment groups. Treated mice were fed diets containing 2,000, 10,000 or 50,000 ppm of tumeric oleoresin. There were 50 mice in the 2,000 and 50,000 ppm groups and 51 mice in the 10,000 ppm group. In addition to data from the tumeric oleoresin study, historical control data on the occurrence of hepatocellular adenomas in female mice were available from three recent studies (NTP study numbers 394, 419 and 439) that used the same feed regimen and were conducted at the same laboratory. Summary statistics are shown in Table 1.

Data for each mouse consist of an indicator of the presence or absence of hepatocellular adenomas (a type of tumor) and the age at death. It is important to include age at death when investigating treatment effects on tumor incidence to avoid biases that can arise when the chemical decreases survival and hence the time at risk of developing the tumor. Analyses of NTP data focus on assessing evidence of an overall dose-response trend, comparing tumor incidence in specific groups, and identifying a low observed adverse effect level (LOAEL). In the historical database, there are many chemicals that show evidence of an increasing monotonic trend at lower dose levels, with a downturn in the high dose group. Such a downturn may be due to toxic effects of the chemical that inhibit tumor development.

We assume the following model for the probability of tumor occurrence:

$$\Pr(y_i = 1 | w_i, x_i, \boldsymbol{\alpha}, \boldsymbol{\beta}, \tau) = \Phi\left(\mathbf{x}'_i \boldsymbol{\alpha} + \sum_{j=1}^{w_i-1} (-1)^{1(j \geq \tau)} \beta_j\right), \quad (12)$$

where  $w_i \in \{1, 2, 3, 4\}$  is a categorical variable indexing the dose group,  $\mathbf{x}_i = (1, x_i)'$ ,  $x_i$  is a monotone transformation of age at death (chosen so that  $x_i$  is closer to standard normal

distributed), and  $\boldsymbol{\alpha}$  are unknown regression coefficients. Under this model, a positive value of  $\beta_j$  implies that the probability of tumor is higher in group  $j + 1$  than in group  $j$ , for a fixed age at death. Under the assumption that  $\beta_j \geq 0$ , (12) restricts the probability of tumor occurrence to be non-decreasing up to dose group  $\tau$  and non-increasing thereafter.

We choose a prior for  $\boldsymbol{\alpha}, \boldsymbol{\beta}, \tau$  as described in Section 2, with

$$\pi(\boldsymbol{\alpha}) = N(\boldsymbol{\alpha}_0, \boldsymbol{\Sigma}_0), \quad \pi(\boldsymbol{\beta}) = \prod_{j=1}^3 \text{ZI-N}^+(\beta_j; \phi, m_j, s_j^2), \quad \text{and} \quad \pi(\tau) = \text{Mult}(1, 2, 3, 4; \mathbf{p}_0),$$

where  $\boldsymbol{\alpha}_0$  is the posterior mode and  $\boldsymbol{\Sigma}_0$  is 5 times the posterior covariance matrix from an analysis of  $\boldsymbol{\alpha}$  using the historical control data, and with  $\phi \sim \text{Beta}(\phi; 0.67, 0.33)$ , determined according to (3), to assign on average equal prior probability to  $H_0 : \beta_1 = \beta_2 = \beta_3 = 0$  and  $H_1 : \beta_j \geq 0$  (with at least one strict inequality). In addition, we let  $m_j = 0$  and  $s_j^2 = 4$  to correspond to our vague (on the scale of the linear predictor) prior about the slope, and we let  $\mathbf{p}_0 = (0, 0, 0.5, 0.5)'$ , which, conditional on there being a non-zero dose effect at the highest level, corresponds to prior ignorance about a possible downturn. We evaluate the sensitivity of the results to these choices of prior distributions in Section 4.4.

## 4.2 The Analysis

We implemented the Gibbs sampler described at the end of Section 3.1 for posterior computation, updating each  $\beta_j$  one at a time, and including a Metropolis step for  $\tau$ . For each analysis, we ran the chain for 310,000 iterations, discarding the first 10,000 iterations as a burn-in. The resulting posterior probabilities of the twelve model configurations are shown in Table 2. We estimate the posterior probability of the global null hypothesis  $H_0$  to be 0.006, providing strong evidence for a dose-response trend. Estimates of the posterior probability of sub-null hypotheses  $H_{0j} : \beta_j = 0$  and alternative hypotheses consistent with  $H_1 : \beta_j \geq 0$  for  $j = 1, 2, 3$  can be computed via marginalization from the values in Table 2. In particular, we estimate  $\Pr(\beta_j = 0 | \mathbf{Y}) = 0.903, 0.010$  and  $0.845$  for  $j = 1, 2, 3$  respectively, providing strong evidence for a dose effect at the 10,000 ppm group, and little support for dose effects

for the 2,000 and 50,000 ppm groups.

To assess evidence of a downturn at the highest dose level, we computed the conditional probability of a downturn given a dose effect,  $\Pr(\tau = 3 | \beta_3 > 0, \mathbf{Y}) = 0.797$ , providing some evidence of a downturn. However, the evidence is weak, particularly considering that the marginal probability of a dose effect at the highest level is only  $\Pr(\beta_3 > 0 | \mathbf{Y}) = 0.155$ . The posterior distribution of the LOAEL can be computed as described in Section 3.4, and we find these probabilities to be 0.097, 0.896 and  $< 0.001$  for the 2,000, 10,000, and 50,000 ppm groups, respectively, providing good evidence for the 10,000 ppm group to be the LOAEL.

Estimates of the posterior probability of tumor occurrence at each dose level can be easily computed from the output of the Gibbs sampler. Figure 1 shows a fitted dose-response curve, along with 95% credible intervals for the probability of tumor occurrence in each dose group. These results are averaged over all twelve possible models, and so the possible downturn at the highest dose level is only slightly apparent, as the large posterior probability of no dose effect dominates. Our model averaged estimator tends to shrink dose group differences towards zero in the absence of evidence in the data of a dose effect, an appealing property which tends to minimize investigator over-interpretation of noise as evidence in favor of non-monotone dose response. Estimates under the highest probability model or other models of interest can be similarly computed.

#### 4.3 Comparison to Other Models

We compared our analysis of the NTP data to three other models: (i) an unconstrained Bayesian approach; (ii) the Gunn and Dunson (GD) (2003) approach described in Section 1; and (iii) the method of Peddada et al. (2001), a frequentist method. The unconstrained Bayesian analysis replaces (12) with

$$\Pr(y_i = 1 | w_i, x_i, \boldsymbol{\alpha}, \boldsymbol{\beta}) = \Phi\left(\mathbf{x}'_i \boldsymbol{\alpha} + \sum_{j=1}^{w_i-1} \beta_j\right),$$

where the priors are now unconstrained normal densities

$$\pi(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \pi(\boldsymbol{\alpha})\pi(\boldsymbol{\beta}) = \text{N}(\boldsymbol{\alpha}_0, \boldsymbol{\Sigma}_0) \prod_{j=1}^{k-1} \text{N}(\beta_j; m_j, s_j^2).$$

This approach assigns zero prior and posterior probability to hypotheses of no dose effect. To perform inferences we applied the GD projection approach, allowing for a possible downturn only at the highest dose group.

Figure 1 compares the fitted dose-response curves for all three analyses. Noting the narrower credible intervals for the constrained analysis, we expect in general that including the parameter constraints will result in narrower credible intervals compared to those obtained from an unconstrained analysis. In the case where the constraints are not supported by the data, we expect to see higher posterior probability for the local null hypotheses, rather than inflated credible intervals.

Comparing our inferences to the GD results, one notable difference is that the posterior probabilities of the local null hypotheses tend to be much smaller under the GD approach. The final column of Table 2 displays the posterior probabilities for the twelve possible models based on the projection approach. Almost all of the mass is allocated to models with zero or one equalities between dose groups. As these results are only quasi-Bayesian in the sense that they are not based on an explicit prior distribution, we feel that they may overstate the dose-effect relationships and may not provide the adequate shrinkage needed to account for multiple comparisons that our prior distributions achieve. This phenomenon also explains why the GD approach assigns much more posterior probability to a downturn than our approach (0.809 compared to 0.123), as our approach assigns much more mass to no dose effect at the highest level (0.845 compared to 0.012). As noted in Section 4.2, when we condition on there being a dose effect at the highest level, our posterior probability of a downturn (0.797) is more concordant with the GD results.

The point estimates from the Peddada et al. (2001) approach are shown in Figure 1, and

the p-value for a dose-response trend (provided by Shyamal Peddada) is 0.0045. The Peddada estimates show a much more dramatic downturn at the highest dose. This discrepancy may be due partly to their very different survival adjustment, based on the poly-3 correction to the sample size. More importantly, their estimator conditions on an estimate for the peak location, while the Bayesian estimators average across the posterior for the peak. It is likely that both the frequentist and the Bayesian preferred model estimator may overestimate the height of the peak.

#### 4.4 Sensitivity to Prior Distributions

In order to assess the sensitivity to the prior, we performed analyses with  $s_j^2 = 2$  and  $s_j^2 = 8$ , which are within the range of reasonable values for  $s_j^2$ . Table 2 compares these priors with  $s_j^2 = 4$ , which we used in our main analysis, in terms of the posterior probabilities of the hypotheses of interest. While slight differences are apparent, scientific conclusions based on these analyses would be the same. In general, the results are robust to changes in the prior, although including prior information about the dose effect parameters, especially with respect to their covariance structure, would be beneficial. Such priors could be constructed based on previous analyses of related chemicals.

As discussed in Section 4.1, our main analysis used historical control data in order to construct an informative prior distribution for the baseline tumor rate and time of death coefficients. In order to assess the impact of this prior, we ran a separate analysis using  $\Sigma_0 = 5I_2$  as the prior covariance matrix and  $\boldsymbol{\alpha}_0 = (0, 0)'$  as the prior mean for  $\boldsymbol{\alpha}$ . Table 2 compares the posterior probabilities of the 12 possible configurations of  $\mathbf{h}$  under this prior to the original analysis. One noticeable effect of this flat prior is that  $\Pr(H_0|\mathbf{Y}) = 0.027$ , which is much larger than it was under the informative prior for  $\boldsymbol{\alpha}$ , underscoring the importance of using historical control data when available.

Finally, we performed two analyses to assess the effect of the prior on  $\phi$ . First we fixed

$\phi = 0.5^{1/3}$  as described in Section 2.1. This specification lead to more posterior mass being allocated to the top model, as seen in Table 2 under the column  $\phi$ . Adding more flexibility, we performed a second analysis using  $\pi(\phi) = 1$ , the uniform distribution. The results are in Table 2 under the column  $\pi(\phi)$ , and are more concordant with our main analysis.

## 5. Discussion

Assessing evidence of order constraints in regression parameters is often of interest in applications with categorical predictor variables. Typically there are many orderings that can be ruled out as not scientifically plausible *a priori* (e.g. a dose-response pattern with multiple changes in direction), and one often has specific information about monotonicity or umbrella ordering. Incorporating this information into the analysis can reduce uncertainty about hypotheses and parameters of interest. In this article we have proposed a Bayesian approach which simultaneously incorporates the order constraints as prior information, reasonably accounts for multiple comparisons, and is flexible enough to accommodate simple increasing order, simple tree order and umbrella ordering. Information from historical studies is easily incorporated in the analysis via the prior distribution, and we adjust for a mix of continuous and discrete covariates. Perhaps the most attractive feature of our model is that inferences on all parameters and hypotheses of interest can be computed from a single run of a Gibbs sampler, avoiding the need to fit separate models for each hypothesis.

While we used the toxicology application to motivate our approach, there are broad applications to essentially any study having ordinal predictors, because some prior information is almost always available about which orderings are plausible. Certainly this is the case in clinical trials with multiple treatment groups and most epidemiological studies.

## ACKNOWLEDGEMENTS

We would like to thank Dr. Joe Haseman for providing the data for the example, Dr. Laura Gunn for providing code used to implement the isotonic projection approach, and Dr.

Shyamal Peddada for analyzing our data with his method.

## REFERENCES

- Albert, J.H. (1994). A Bayesian approach to estimation of GPAs of University of Iowa freshman under order restrictions. *Journal of Educational Statistics* **19**, 1-21.
- Albert, J.H. and Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* **88**, 669-679.
- Chipman, H., George, E.I. and McCulloch, R.E. (2001). The practical implementation of Bayesian model selection. *IMS Lecture Notes - Monograph Series* **38**.
- Dunson, D.B. and Herring, A.H. (2003). Bayesian inferences in the Cox model for order restricted alternatives. *Biometrics* **59**, 918-925.
- Dunson, D.B. and Neelon, B. (2003). Bayesian inferences on order-constrained parameters in generalized linear models. *Biometrics* **59**, 286-295.
- Gelfand, A.E. and Kottas, A. (2001). Nonparametric Bayesian modeling for stochastic order. *Annals of the Institute of Statistical Mathematics* **53** (4), 865-876.
- Gelfand, A.E. and Kuo, L. (1991). Nonparametric Bayesian bioassay including ordered polytomous response. *Biometrika* **78**, 657-666.
- Gelfand, A.E. and Smith, A.F.M. (1990). Sampling-Based Approaches to Computing Marginal Densities. *Journal of the American Statistical Association* **85**, 398-409.
- Gelfand, A.E., Smith, A.F.M., and Lee, T.M. (1992). Bayesian analysis of constrained parameter and truncated data problems using Gibbs sampling. *Journal of the American Statistical Association* **87**, 523-532.

- George, E.I. and McCulloch, R.E. (1997). Approaches for Bayesian variable selection. *Statistica Sinica* **7**, 339-373.
- Geweke, J. (1996). Variable selection and model comparison in regression. In *Bayesian Statistics* **5**, J.O. Berger, J.M. Bernardo, A.P. Dawid, and A.F.M. Smith (eds.). Oxford University Press, 609-620.
- Gönen, M., Westfall, P.H., and Johnson, W.O. (2003). Bayesian multiple testing for two-sample multivariate endpoints. *Biometrics* **59**, 76-82.
- Gunn, L.H. and Dunson, D.B. (2003). Bayesian inferences on shape constrained hormone trajectories in the menstrual cycle. ISDS Discussion Paper 03-24, Duke University.
- Hoff, P.D. (2003). Bayesian methods for partial stochastic orderings. *Biometrika* **90**, 303-317.
- Holmes, C.C. and Heard, N.A. (2003). Generalized monotonic regression using random changepoints. *Statistics in Medicine* **22**, 623-638.
- Hwang, J.T.G. and Peddada, S.D. (1994). Confidence interval estimation subject to order restrictions. *Annals of Statistics* **22**, 67-93.
- Lavine, M. and Mockus, A. (1995). A nonparametric Bayes method for isotonic regression. *Journal of Statistical Planning and Inference* **46**, 235-248.
- Liu, J.S., Wong, W.H. and Kong, A. (1994). Covariance structure of the Gibbs sampler with applications to the comparisons of estimators and augmentation schemes. *Biometrika* **81**, 27-40.
- Mancuso, J.Y., Ahn, H. and Chen, J.J. (2001). Order-restricted dose-related trend tests. *Statistics in Medicine* **20**, 2305-2318.

- McCullagh, P. and Nelder, J.A. (1989). *Generalized Linear Models, Second Edition*. New York: Chapman & Hall/CRC.
- Neelon, B. and Dunson, D.B. (2004). Bayesian isotonic regression and trend analysis. *Biometrics* **60**, 398-406.
- NTP (1993). Toxicology and carcinogenesis studies of turmeric oleoresin in F344/N rats and B6C3F1 mice. *Technical Report 427*. National Toxicology Program, Research Triangle Park, North Carolina.
- Peddada, S.D., Prescott, K.E., and Conaway, M. (2001). Tests for order restrictions in binary data. *Biometrics* **57**, 1219-1227.
- Peddada, S.D., Lobenhofer, E.K., Li, L.P., Afshari, C.A., Weinberg, C.R. and Umbach, D. (2003). Gene selection and clustering for time-course and dose-response microarray experiments using order-restricted inference. *Bioinformatics* **19**, 834-841.
- Ramgopal, P., Laud, P.W. and Smith, A.F.M. (1993). Nonparametric Bayesian bioassay with prior constraints on the shape of the potency curve. *Biometrika* **80**, 489-498.
- Robertson, T., Wright, F. and Dykstra, R. (1988). *Order Restricted Statistical Inference*. New York: Wiley.
- Simpson, D. and Margolin, B. (1986). Recursive nonparametric testing for dose-response relationships subject to downturns at higher dose. *Biometrika* **73**, 589-596.
- Westfall, P.H., Johnson, W.O. and Utts, J.M. (1997). A Bayesian perspective on the Bonferroni adjustment. *Biometrika* **84**, 419-427.

## APPENDIX 1

### Calculations for Gibbs Sampler

To perform STEP 3 of the Gibbs sampler described in Section 3.1, we must sample from

$$p(\boldsymbol{\beta}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}) \propto \text{N}(\mathbf{z}; \boldsymbol{\eta}, \mathbf{I}) \prod_{j=1}^{k-1} \text{ZI-N}^+(\beta_j; \pi_{0j}, m_j, s_j^2), \quad (13)$$

which is a mixture distribution with  $2^{k-1}$  components, corresponding to either  $H_0$  or some combination of the local null hypotheses  $H_{0j}$  and  $H_{1j}$ . Let  $S(\mathbf{h}) = \sum_{j=1}^{k-1} h_j$  be the number of  $\beta_j$  which are allowed to be nonzero under the hypothesis corresponding to  $\mathbf{h}$ . We can write the full conditional posterior (13) as

$$\begin{aligned} p(\boldsymbol{\beta}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}) &= \sum_{l=0}^{k-1} \sum_{\mathbf{h}: S(\mathbf{h})=l} p(\boldsymbol{\beta}, \mathbf{h}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}) \\ &= \sum_{l=0}^{k-1} \sum_{\mathbf{h}: S(\mathbf{h})=l} p(\boldsymbol{\beta}|\mathbf{h}, \mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}) p(\mathbf{h}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}). \end{aligned}$$

Thus in the Gibbs sampler we draw a value of  $\mathbf{h}$  from  $p(\mathbf{h}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y})$  at each iteration and then sample a new value of  $\boldsymbol{\beta}$  from  $p(\boldsymbol{\beta}|\mathbf{h}, \mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y})$  conditional on the chosen model.

Write  $p(\mathbf{h}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}) = q(\mathbf{h}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y})/Q$ , where  $Q = \sum_{\mathbf{h}} q(\mathbf{h}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y})$ . For  $\mathbf{h} = (0, \dots, 0)'$  we have  $q(\mathbf{h}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}) = \phi^{k-1}$  and  $p(\boldsymbol{\beta}|\mathbf{h}, \mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}) = \prod_{j=1}^{k-1} 1(\beta_j = 0)$ , while for  $\mathbf{h}$  not equal to the zero vector we have

$$q(\mathbf{h}|\mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}) = \frac{P_{S(\mathbf{h})}(\mathbf{t} \geq 0; \hat{\boldsymbol{\mu}}_{\mathbf{h}}, \hat{\boldsymbol{\Sigma}}_{\mathbf{h}})}{N_{S(\mathbf{h})}(\mathbf{0}; \hat{\boldsymbol{\mu}}_{\mathbf{h}}, \hat{\boldsymbol{\Sigma}}_{\mathbf{h}})} \prod_{j=1}^{k-1} \phi^{1-h_j} \left[ (1-\phi) \frac{N_1(0; m_j, s_j^2)}{P_1(t > 0; m_j, s_j^2)} \right]^{h_j}, \quad (14)$$

$$p(\boldsymbol{\beta}|\mathbf{h}, \mathbf{z}, \boldsymbol{\alpha}, \phi, \mathbf{Y}) = \frac{N_{S(\mathbf{h})}(\boldsymbol{\beta}_{\mathbf{h}}; \hat{\boldsymbol{\mu}}_{\mathbf{h}}, \hat{\boldsymbol{\Sigma}}_{\mathbf{h}})}{P_{S(\mathbf{h})}(\mathbf{t} \geq 0; \hat{\boldsymbol{\mu}}_{\mathbf{h}}, \hat{\boldsymbol{\Sigma}}_{\mathbf{h}})} \left[ \prod_{j=1}^{k-1} 1(\beta_j = 0)^{1-h_j} 1(\beta_j > 0)^{h_j} \right]. \quad (15)$$

$N_p(\mathbf{0}; \cdot, \cdot)$  is the density function for a  $p$  dimensional normal random variable evaluated at the zero vector,  $P_p(\mathbf{t} \geq 0; \cdot, \cdot) = \int_0^\infty \dots \int_0^\infty N_p(\mathbf{t}; \cdot, \cdot) dt$ , and  $\boldsymbol{\beta}_{\mathbf{h}}$  is a  $S(\mathbf{h})$  by 1 dimensional vector containing the nonzero elements of  $\boldsymbol{\beta}$  for a given  $\mathbf{h}$ .

The precision matrix for  $\beta_{\mathbf{h}}$ ,  $\widehat{\Sigma}_{\mathbf{h}}^{-1}$ , has diagonal elements  $1/s_j^2 + \mathbf{w}'_j \mathbf{w}_j$  for  $j$  such that  $h_j = 1$ , and  $\mathbf{w}_j$  is the corresponding column of the design matrix  $\mathbf{W}$ . The off-diagonal elements of  $\widehat{\Sigma}_{\mathbf{h}}^{-1}$  are  $\mathbf{w}'_j \mathbf{w}_{j'}$ . The mean vector for  $\beta_{\mathbf{h}}$  is  $\widehat{\mu}_{\mathbf{h}} = \widehat{\Sigma}_{\mathbf{h}} \mathbf{v}_{\mathbf{h}}$ , where  $\mathbf{v}_{\mathbf{h}}$  is a  $S(\mathbf{h})$  by 1 vector whose  $j$ th element is  $m_j/s_j^2 + (\mathbf{z} - \mathbf{X}\alpha)' \mathbf{w}_j$  for  $j$  such that  $h_j = 1$ .

Let  $\mathcal{H} = \{0, 1\}^{k-1}$  be the  $2^{k-1}$  possible values of  $\mathbf{h}$ . Thus to update the mixture component at iteration  $t$  in the Gibbs sampler, one samples  $\mathbf{h}^{[t+1]} \in \mathcal{H}$  with probability  $p(\mathbf{h}^{[t+1]} = \mathbf{h} | \mathbf{z}, \alpha, \phi, \mathbf{Y}) = q(\mathbf{h} | \mathbf{z}, \alpha, \phi, \mathbf{Y}) / Q$ . Given the updated mixture component, the nonzero elements of  $\beta^{[t+1]}$ , namely  $\beta_{\mathbf{h}}^{[t+1]}$ , are drawn from  $p(\beta^{[t+1]} | \mathbf{h}, \mathbf{z}, \alpha, \phi, \mathbf{Y})$ , which is an  $S(\mathbf{h})$  dimensional truncated normal distribution with mean variance described above; no draw needs to be made when  $\beta^{[t+1]}$  is the zero vector because each element of  $\beta^{[t+1]}$  is set to zero.

## APPENDIX 2

### Modifications to the Gibbs Sampler for Umbrella Orderings

To sample from  $p(\beta, \mathbf{h}, \tau | \mathbf{z}, \alpha, \phi, \mathbf{Y})$  we need to first sample from the  $k 2^{k-1}$  possible combinations of  $\tau$  and  $\mathbf{h}$ . Define the diagonal matrix  $\mathbf{T}_{\tau}$  with  $[\mathbf{T}_{\tau}]_{ij} = 0$  for  $i \neq j$  and

$$[\mathbf{T}_{\tau}]_{ii} = \begin{cases} 1 & \text{if } i < \tau \\ -1 & \text{o.w.} \end{cases}$$

To compute  $q(\mathbf{h}, \tau = l | \mathbf{z}, \alpha, \phi, \mathbf{Y})$  we need only redefine  $\widehat{\mu}_{\mathbf{h}}$  to now depend on  $\tau$ . Define  $\widehat{\mu}_{\mathbf{h}, \tau} = \widehat{\Sigma}_{\mathbf{h}} \mathbf{v}_{\mathbf{h}, \tau}$ , where the  $j$ th element of  $\mathbf{v}_{\mathbf{h}, \tau}$  is defined as  $m_j/s_j^2 + (\mathbf{z} - \mathbf{X}\alpha)' [\mathbf{W}\mathbf{T}_{\tau}]_j$ , and  $[\mathbf{W}\mathbf{T}_{\tau}]_j$  is the  $j$ th column of  $\mathbf{W}\mathbf{T}_{\tau}$ . Sampling proceeds as described in Appendix 1 with  $\widehat{\mu}_{\mathbf{h}}$  replaced by  $\widehat{\mu}_{\mathbf{h}, \tau}$ .

NTP Study	Group	Tumors	$n$	% Mortality
394	control	3	49	44.90
419	control	5	50	14.00
439	control	6	49	24.49
427	control	7	50	22.00
427	2000	8	50	18.00
427	10000	19	51	33.33
427	50000	14	50	16.00

Table 1: Summary of data described in Section 4. The first three rows contain the historical data used to construct the prior distribution for  $\alpha$ , and the last four rows are the data from our main analysis. Doses are given in ppm in the “Group” column. “Tumors” gives the number of hepatocellular adenomas observed in each dose group, and  $n$  is the total number of mice in each group. The final column gives the percentage of mice in each group who died before the terminal sacrifice.

Model			Posterior Probability						
$\beta_1$	$\beta_2$	$\beta_3$	$s_j^2 = 2$	$s_j^2 = 4$	$s_j^2 = 8$	$\pi(\alpha)$	$\phi$	$\pi(\phi)$	GD
=	>	=	0.709	0.780	0.831	0.770	0.923	0.742	0.012
=	>	<	0.115	0.093	0.074	0.095	0.037	0.114	0.624
>	>	=	0.073	0.057	0.043	0.049	0.023	0.069	*
>	>	<	0.054	0.030	0.018	0.026	0.001	0.032	0.185
=	>	>	0.030	0.023	0.017	0.023	0.009	0.028	0.138
>	>	>	0.013	0.008	0.004	0.006	*	0.008	0.040
=	=	=	0.004	0.006	0.010	0.027	0.004	0.002	*
>	=	=	0.002	0.002	0.003	0.003	0.002	0.003	*
=	=	>	*	*	*	*	*	*	*
>	=	>	*	*	*	*	*	*	*
>	=	<	*	*	*	*	*	*	*
=	=	<	*	*	*	*	*	*	*

Table 2: Estimated posterior probabilities of the twelve models for the analyses described in Section 4. In the first three columns, “=” means  $\beta_j = 0$ , “>” means  $\beta_j > 0$  and “<” corresponds to a downturn,  $\beta_j < 0$ . The column labeled  $s_j^2 = 4$  is from the main analysis in Sections 4.1 and 4.2, while the other columns are the sensitivity analyses described in Section 4.4, except for the final column, which refers to the Gunn-Dunson approach in Section 4.3. The character “\*” denotes probability  $< 0.001$ .

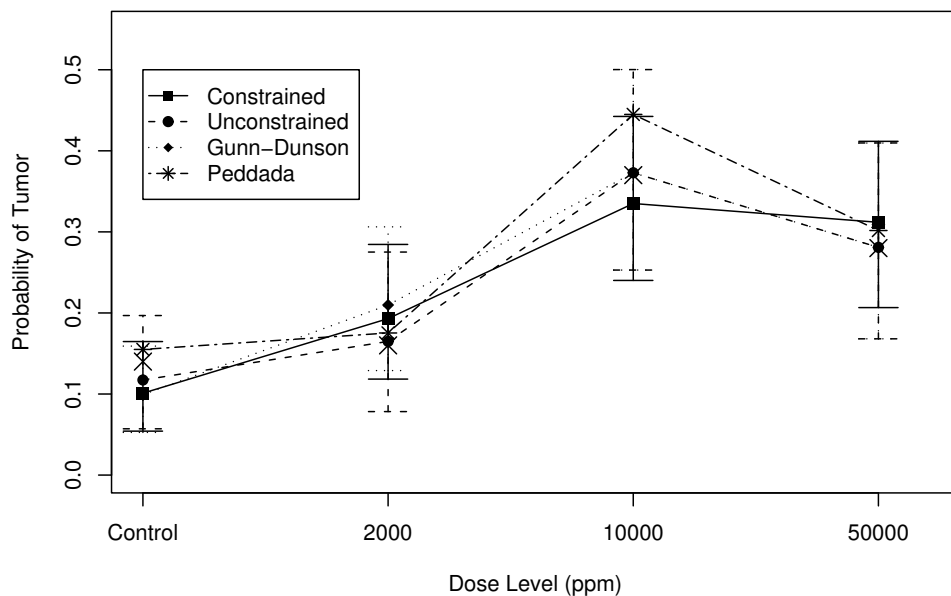


Figure 1: Model averaged dose response curves and posterior distributions for the analyses described in Section 4.2 (constrained) and Section 4.3 (unconstrained). The solid colored characters are posterior means, and the intervals are 95% credible regions. The “X” characters represent the empirical proportion of tumors in each dose group, and the “\*” characters are the point estimates from the frequentist analysis described in Section 4.3.