

# A BAYESIAN $\chi^2$ TEST FOR GOODNESS-OF-FIT

VALEN E. JOHNSON

ABSTRACT. This article describes an extension of classical  $\chi^2$  goodness-of-fit tests to Bayesian model assessment. The extension, which essentially involves evaluating Pearson's goodness-of-fit statistic at a parameter value drawn from its posterior distribution, has the important property that it is asymptotically distributed as a  $\chi^2$  random variable on  $K-1$  degrees of freedom, independently of the dimension of the underlying parameter vector. By averaging over the posterior distribution of this statistic, a global goodness-of-fit diagnostic is obtained. Advantages of this diagnostic—which may be interpreted as the area under an ROC curve—include ease of interpretation, computational convenience, and favorable power properties. The proposed diagnostic can be used to assess the adequacy of a broad class of Bayesian models, essentially requiring only a finite-dimensional parameter vector and conditionally independent observations.

## 1. INTRODUCTION

Model assessment presents a challenge to Bayesian statisticians, one that has become an increasingly serious problem as computational advances have made it possible to entertain models of a complexity not considered even a decade ago. Because diagnostic methods have not kept pace with these computational advances, practitioners are often faced with the prospect of interpreting results from a model that has not been adequately validated.

Numerous solutions to this problem have been considered. The most orthodox of these depend on the specification of alternative models and the use of Bayes factors for model selection. This approach is reasonable when both a relatively broad class of models can be specified as alternatives, and when implied Bayes factors can be readily computed. Unfortunately, it often happens in practice that neither requirement is satisfied, making this approach impractical for routine application. Complicating the situation still further is the fact that Bayes factors are not defined when improper priors are used in model specification, although this difficulty may be partially circumvented through the use of intrinsic Bayes factors or related devices (e.g., Berger and Pericchi (1996), O'Hagan (1995)).

A second strategy for assessing model adequacy centers on the use of posterior-predictive model checks. This approach was initially proposed by Guttman (1967) and Rubin (1984), and was extended to more general discrepancy functions by Gelman, Meng, and Stern (1996) (Gelfand (1996) has advocated related techniques based on cross-validated predictive densities). The primary advantage of posterior-predictive model assessment is its relative ease of implementation. In many models, the output from numerical algorithms used to generate samples from the posterior distribution can be used to generate observations from the predictive model, which in turn can be used to compute  $p$  values for the discrepancy function of interest. Posterior-predictive model assessment also facilitates case-diagnostics, which, in many circumstances, are more telling in examining model fit than are global goodness-of-fit statistics. However, such approaches also have an important disadvantage. As Bayarri and Berger (2000) and Robin, van der Vaart, and Ventura

(2000) and others have noted, they do not produce  $p$  values that have (even asymptotically) a uniform distribution. Because output from predictive posterior model checks is not calibrated, using  $p$  values based on them for model assessment is difficult.

Bayarri and Berger (2000) and Robin, van der Vaart, and Ventura (2000) propose alternative distributions under which  $p$  values, and thus model diagnostics, can be calculated. These include partial posterior predictive  $p$  values and conditional predictive  $p$  values (Berger and Bayarri), and modifications to posterior predictive and “plug-in”  $p$  values (Robin, van der Vaart and Ventura). The attractive feature of each of these variations on more standard definitions of  $p$  values is that these statistics are distributed either as  $U(0, 1)$  random variables, or approach  $U(0, 1)$  random variables as sample sizes become large. Their drawback is that they can seldom be defined and calculated in realistically complex models.

The goal of this article is to present a goodness-of-fit diagnostic that bridges the gap between diagnostics that are easy to compute but whose null distributions are unknown, and diagnostics whose null distributions are known but that cannot generally be computed. The proposed diagnostic is closely related to the classical  $\chi^2$  goodness-of-fit statistic, whose properties are now briefly reviewed.

In the case of a point null hypothesis, the standard  $\chi^2$  statistic is defined as

$$R^0 = \sum_{k=1}^K \frac{(m_k - np_k)^2}{np_k},$$

where  $m_k$  represents the number of observations observed within the  $k$ th partitioning element,  $p_k$  the probability assigned by the null model to this interval,  $K$  the number of partitions or intervals specified over the sample space, and  $n$  the sample

size. For independent and identically distributed data satisfying certain regularity requirements, Pearson (1900) demonstrated that the asymptotic distribution of  $R^0$  was  $\chi^2$  on  $K - 1$  degrees of freedom.

The situation for composite hypotheses is more complicated. Assuming that bins are determined *a priori*, Cramér (1946, pages 426-434) demonstrated that the distribution of

$$R^g = \sum_{k=1}^K \frac{(m_k - np_k^g)^2}{np_k^g}$$

is that of a  $\chi^2$  random variable on  $K - s - 1$  degrees of freedom when estimation of model parameters is based on maximum likelihood for the *grouped* data or on the minimum  $\chi^2$  method. Maximum likelihood estimation for the grouped data implies maximization of the function

$$\prod_k p_k(\boldsymbol{\theta})^{m_k}$$

with respect to the  $s$ -dimensional parameter  $\boldsymbol{\theta}$ , while minimum  $\chi^2$  estimation involves the determination of a value of  $\boldsymbol{\theta}$  that minimizes a function related to  $R^g$ .

Chernoff and Lehmann (1954) considered the distribution of the  $\chi^2$  statistic in the more typical situation in which values of  $p_k$  are based on maximum likelihood estimates obtained using the raw (ungrouped) data. In this case, the distribution of the goodness-of-fit statistic is generally not one of a  $\chi^2$  distribution, but instead produces a value  $\hat{R}$  that has a distribution that falls between  $R^0$  and  $R^g$ . For models containing many parameters, the gap between the degrees of freedom associated with these two statistics is large, and, as a result, the  $\chi^2$  goodness-of-fit test based on the maximum likelihood estimate is usually not useful for assessing model fit in high-dimensional settings.

The goodness-of-fit statistic proposed here represents a modification of the  $\chi^2$  statistics considered above. The modification, denoted by  $R^B(\tilde{\theta})$  (or, more simply, by  $R^B$  when no confusion arises), is obtained by fixing the values of  $p_k$  and instead considering the bin counts  $m_k$  as random quantities. Allocation of observations to bins is made according to the value of each observation's conditional distribution function, conditionally on a parameter value sampled either from the posterior distribution or the asymptotic distribution of the maximum likelihood estimator. (Note that the statistic obtained in this way has some resemblance to the  $\chi^2$  statistics considered by, for example, Moore and Spruill (1975), although emphasis there focuses on randomized cells rather than on posterior sampling of parameter vectors.) The distinguishing feature of  $R^B(\tilde{\theta})$  is that, for many statistical models, its asymptotic distribution is  $\chi^2$  on  $K - 1$  degrees of freedom, independently of the dimension of the parameter vector  $\theta$ .

Because it is the sampling distribution of  $R^B$  that has a  $\chi^2$  distribution, one might argue that this procedure does not really represent a Bayesian goodness-of-fit diagnostic. However, sampling parameter values from a distribution for the purpose of inference occurs more naturally within the Bayesian paradigm, and for this reason it is likely that the proposed diagnostic will find more application there. In addition, the formal test statistics proposed below are based on the posterior distribution of  $R^B$ . For this reason, values of  $\tilde{\theta}$  used in the definition of  $R^B$  are assumed to represent samples from the posterior distribution on the parameter vector, rather than samples generated from the asymptotic normal distribution of the maximum likelihood estimator. However, either interpretation is valid.

A particular value of  $R^B(\tilde{\theta})$  generated from a single sampled value of  $\tilde{\theta}$  from the posterior distribution has, asymptotically, a  $\chi^2$  distribution on  $K-1$  degrees of freedom, and so provides a mechanism for defining a significance test for the validity of an assumed model. Nonetheless, it is clearly desirable to base formal tests of model adequacy on the posterior distribution of  $R^B$ . This goal can be approached from a variety of perspectives, including modifications of those suggested in Verdine and Wasserman (1998) or Robert and Rousseau (2002), but the method advocated here is to base formal significance tests on the probability that a randomly drawn value of  $R^B$  exceeds a  $\chi^2$  random variable on  $K-1$  degrees of freedom. This probability, denoted by  $A$ , is a commonly used quantity in signal detection theory and represents the area under the ROC curve for comparing the posterior distribution of  $R^S$  to a  $\chi^2_{K-1}$  random variable (e.g., Hanley and McNeil 1982). If the specified model is correct, then the expected value of  $A$  is 0.5. Values of  $A$  close to 0.5 thus reflect adequate model fit, while values of  $A$  close to 0 or 1 are indications of model inadequacy. Procedures for assessing the sampling distribution of  $A$  are described in the examples.

The remainder of the paper is organized as follows. In the next section, the Bayesian  $\chi^2$  statistic  $R^B$  is defined and its asymptotic properties are described. A corollary extending these properties from i.i.d. observations to conditionally independent observations is also presented. Following this, several examples that illustrate the application of this statistic and the derived quantity  $A$  are presented. Discussion and concluding remarks appear in Section 4. Proofs to the theorem and corollaries of Section 2 appear in the Appendix.

2. A BAYESIAN  $\chi^2$  STATISTIC

To begin, let  $Y_1, \dots, Y_n$  ( $= \mathbf{Y}$ ) denote scalar-valued, continuous, identically distributed, conditionally independent observations drawn from probability density function  $f(y|\boldsymbol{\theta})$  defined with respect to Lebesgue measure and indexed by a  $s$ -dimensional parameter vector  $\boldsymbol{\theta} \in \Theta \subset \mathbf{R}^s$ . Denote by  $F(\cdot|\boldsymbol{\theta})$  and  $F^{-1}(\cdot|\boldsymbol{\theta})$  the (non-degenerate) cumulative distribution and inverse distribution functions corresponding to  $f(\cdot|\boldsymbol{\theta})$ . Augment the observed sample  $\mathbf{Y}$  with an i.i.d. sample  $V_1, \dots, V_s$  from a  $U(0,1)$  distribution. Let  $p(\boldsymbol{\theta}|\mathbf{y})$  denote the posterior density of  $\boldsymbol{\theta}$  based on  $\mathbf{Y}$ , and let  $p(\boldsymbol{\theta}_j|\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_{j-1}, \mathbf{y})$  denote the marginal conditional posterior density of  $\boldsymbol{\theta}_j$  given  $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_{j-1}, \mathbf{y})$ . Define  $\tilde{\boldsymbol{\theta}}$  implicitly by

$$(1) \quad v_1 = \int_{-\infty}^{\tilde{\theta}_1} p(\theta_1|\mathbf{y})d\boldsymbol{\theta}_1, \quad \dots \quad v_s = \int_{-\infty}^{\tilde{\theta}_s} p(\theta_s|\tilde{\theta}_1, \dots, \tilde{\theta}_{s-1}, \mathbf{y})d\boldsymbol{\theta}_s.$$

That is,  $\tilde{\boldsymbol{\theta}}$  denotes a value of  $\boldsymbol{\theta}$  sampled from the posterior distribution based on  $\mathbf{Y}$ . Let  $\boldsymbol{\theta}_0$  denote the true but unknown value of  $\boldsymbol{\theta}$ . The maximum likelihood estimate of  $\boldsymbol{\theta}$  is denoted by  $\hat{\boldsymbol{\theta}}$ .

To construct the Bayesian goodness-of-fit statistic proposed here, choose quantiles  $0 \equiv a_0 < a_1, \dots < a_{K-1} < a_K \equiv 1$ , with  $p_k = a_k - a_{k-1}$ ,  $k = 1, \dots, K$ . Define  $\mathbf{z}_j(\tilde{\boldsymbol{\theta}})$  to be a vector of length  $K$  whose  $k$ th element is 0 unless

$$(2) \quad F(Y_j|\tilde{\boldsymbol{\theta}}) \in (a_{k-1}, a_k],$$

in which case it is 1. Finally, define

$$\mathbf{m}(\tilde{\boldsymbol{\theta}}) = \sum_{j=1}^n \mathbf{z}_j(\tilde{\boldsymbol{\theta}}).$$

It follows that the  $k$ th component of  $\mathbf{m}(\tilde{\boldsymbol{\theta}})$ ,  $m_k(\tilde{\boldsymbol{\theta}})$ , represents the number of observations that fell into the  $k$ th bin, where bins are determined by the quantiles of the inverse distribution function evaluated at  $\tilde{\boldsymbol{\theta}}$ . Finally, define

$$(3) \quad R^B(\tilde{\boldsymbol{\theta}}) = \sum_{k=1}^K \left[ \frac{(m_k(\tilde{\boldsymbol{\theta}}) - np_k)}{\sqrt{np_k}} \right]^2.$$

The asymptotic distribution of  $R^B$  is provided in the following theorem.

**Theorem 1.** *Assuming that the regularity conditions specified in the appendix apply,  $R^B$  converges to a  $\chi^2$  distribution on  $K - 1$  degrees of freedom as  $n \rightarrow \infty$ .*

The simplicity of Theorem 1 is somewhat remarkable given the complexity of the corresponding distribution on  $\hat{R}$ . As mentioned above, the asymptotic distribution of  $\hat{R}$  does not, in general, follow a  $\chi^2$  distribution. Instead, it has the distribution of the sum of a  $\chi^2$  random variable on  $K - s - 1$  degrees of freedom and the weighted sum of  $s$  additional squared normal deviates with weights ranging from 0 to 1.

Heuristically, the idea underlying Theorem 1 is that the degrees of freedom lost by substituting the MLE for  $\boldsymbol{\theta}$  in Pearson's  $\chi^2$  statistic are exactly recovered by replacing the MLE with a sampled value from the posterior in  $R^B$ .

As a corollary, Theorem 1 can be extended to the more general case in which the functional form of the density  $f(y|\boldsymbol{\theta})$  varies from observation to observation. Specifically, if the density of the  $j$ th observation is denoted by  $f_j(y|\boldsymbol{\theta})$ , with distribution and inverse distribution functions  $F_j$  and  $F_j^{-1}$ , respectively, then the following corollary also applies.

**Corollary 1.** *Assume the conditions referenced in Theorem 1 are extended so as to provide also for the asymptotic normality of both the posterior distribution on  $\boldsymbol{\theta}$  and*

of the maximum likelihood estimator when the likelihood function is proportional to

$$\prod_{j=1}^n f_j(y_j | \boldsymbol{\theta}).$$

Assume also that the functions  $f_j(\cdot | \boldsymbol{\theta})$  satisfy the same conditions implied in Theorem 1 for  $f(\cdot | \boldsymbol{\theta})$ . Define  $\mathbf{z}_j(\boldsymbol{\theta})$  to be 1 or 0 depending on whether or not

$$(4) \quad F_j(Y_j | \tilde{\boldsymbol{\theta}}) \in (a_{k-1}, a_k],$$

with  $\mathbf{a}$  fixed. Then the asymptotic distribution of  $R^B$  based on this revised definition of  $\mathbf{z}_j(\boldsymbol{\theta})$  is  $\chi^2$  on  $K - 1$  degrees of freedom.

Outlines of the proof of Theorem 1 and the corollary appear in the Appendix.

From a practical perspective, the corollary is important because it extends the definition of  $R^B$  to essentially all models in which observations are continuous and conditionally independent given the value of a finite-dimensional parameter vector.

The results cited above for continuous-valued random variables can be extended to discrete random variables in one of two ways. The most direct extension is to simply proceed as in the continuous case, using a randomization procedure to allocate counts to bins when the mass assigned to an observation spans the boundaries defining the bins. The second is to define fixed bins in the standard way based on the possible outcomes of the random variable, and to then evaluate the bin probabilities at sampled values of  $\boldsymbol{\theta}$  from the posterior distribution. That is, if  $f(y | \boldsymbol{\theta})$  denotes the probability mass function of a discrete random variable  $y$  and

$$(5) \quad p_k(\tilde{\boldsymbol{\theta}}) = \sum_{j=1}^n \sum_{y \in \text{bin } k} f_j(y | \tilde{\boldsymbol{\theta}}),$$

then the  $\chi^2$  statistic  $R^B$  may be redefined as

$$(6) \quad R^B(\tilde{\boldsymbol{\theta}}) = \sum_{k=1}^K \left[ \frac{(m_k - np_k(\tilde{\boldsymbol{\theta}}))}{\sqrt{np_k(\tilde{\boldsymbol{\theta}})}} \right]^2.$$

In this case, the asymptotic distribution of  $R^B(\tilde{\boldsymbol{\theta}})$  is similar to that described above in the continuous case and is detailed in the following corollary.

**Corollary 2.** *If the regularity conditions specified in Theorem 1 apply to the discrete probability mass function  $f(y|\boldsymbol{\theta})$ , then, using predefined bins and the definition of the bin probabilities given in (5), the distribution of  $R^B(\tilde{\boldsymbol{\theta}})$  as defined in (2) converges to a  $\chi^2$  distribution on  $K - 1$  degrees of freedom as  $n \rightarrow \infty$ .*

### 3. EXAMPLES

**3.1. Goodness-of-fit tests under a normal model with unknown mean and variance.** In this example, the distribution of  $R^B$  under a normal model is investigated and compared with the distribution of  $\hat{R}$  and  $R^g$ . Posterior samples of  $R^B$  generated from a single data vector are used in ROC-type analyses to generate a summary model diagnostic. The power of this test statistic is investigated and compared to the power of the test statistic  $R^g$  when data are generated under several non-normal alternatives.

Suppose that  $\mathbf{Y} = (Y_1, \dots, Y_{50})$  denotes a random sample from a standard normal distribution. For purposes of illustration, assume that the mean  $\mu$  and variance  $\sigma^2$  of the data are unknown, and that the joint prior assumed for  $(\mu, \sigma^2)$  is proportional to  $1/\sigma^2$ . Let  $\tilde{\mu}$  and  $\tilde{\sigma}^2$  denote a sampled value from the posterior distribution based on  $\mathbf{Y}$ .

For a given data vector  $\mathbf{Y}$  and posterior sample  $(\tilde{\mu}, \tilde{\sigma}^2)$ , bin counts  $m_k(\tilde{\mu}, \tilde{\sigma}^2)$  are determined by counting the number of observations  $y_i$  that fall into the interval  $(\tilde{\sigma}\Phi^{-1}(a_{k-1})+\tilde{\mu}, \tilde{\sigma}\Phi^{-1}(a_k)+\tilde{\mu})$ , where  $\Phi^{-1}(\cdot)$  denotes the standard normal quantile function. Based on these counts,  $R^B(\tilde{\mu}, \tilde{\sigma}^2)$  is calculated according to (3).

Figure 1 depicts a quantile-quantile plot of  $R^B$  values calculated for 10,000 independent samples of  $\mathbf{Y}$ . Each value of  $R^B$  depicted in this plot corresponds to a single draw of  $(\mu, \sigma^2)$  from the posterior distribution based on a single observation vector  $\mathbf{Y}$ . Bins were defined according to the vector  $\mathbf{a} = (0, .2, .4, .6, .8, 1)$ . As expected, the distribution of  $R^B$  closely mimics that of a  $\chi_4^2$  random variable.

The normal deviates used in the construction of Figure 1 were also used to compute the classical  $\chi^2$  statistic based on the maximum likelihood estimates of  $\mu$  and  $\sigma^2$  (i.e., using the ungrouped data). The quantile-quantile plot of 10,000  $\hat{R}$  values obtained from these data is displayed in Figure 2. In this plot, the  $\hat{R}$  values have been plotted against the expected order statistics from a  $\chi_2^2$  random variable. Five equal probability bins based on the standard normal distribution were used to define these  $\hat{R}$  values. As might be expected, the plotted  $\chi^2$  values display some deviation from the approximate their  $\chi_2^2$  distribution.

Grouped maximum likelihood estimates were also used to calculate  $R^g$  values using these normal samples. The corresponding quantile-quantile plot for the 10,000  $R^g$  values is displayed in Figure 3; as expected, these values demonstrate substantially better agreement with a  $\chi_2^2$  random variable than do the values depicted in Figure 2.

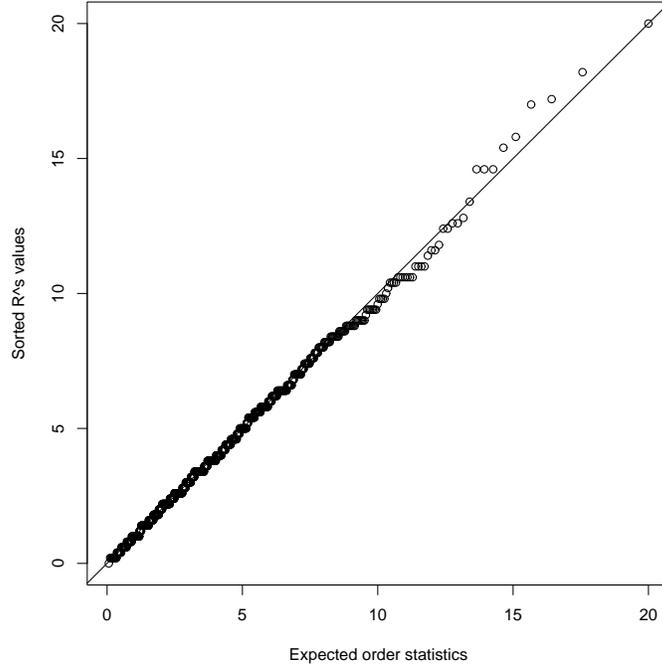


FIGURE 1. Quantile-quantile plot of  $R^B$  values for i.i.d. normal data. Values of  $R^B$  displayed in this plot were determined from independent samples of 50 standard normal deviates, and are plotted against the expected order statistics from a  $\chi_4^2$  distribution. Posterior samples of the mean and variance were estimated using reference priors and observations were binned into 5 bins of equal probability (i.e.,  $\mathbf{a} = (0, .2, .4, .6, .8, 1)$ ).

Returning to the investigation of the properties of  $R^B$ , Figure 1 demonstrates excellent agreement between this statistic and its asymptotic distribution. To illustrate its power in detecting departures from the normal model, suppose now that the experiment above is repeated with independent Student  $t$  variates substituted for the normal deviates. The degrees of freedom of the  $t$  variates used in this replication range from 1 to 10, and for each value within this range, 10,000 independent samples of size 50 were drawn.

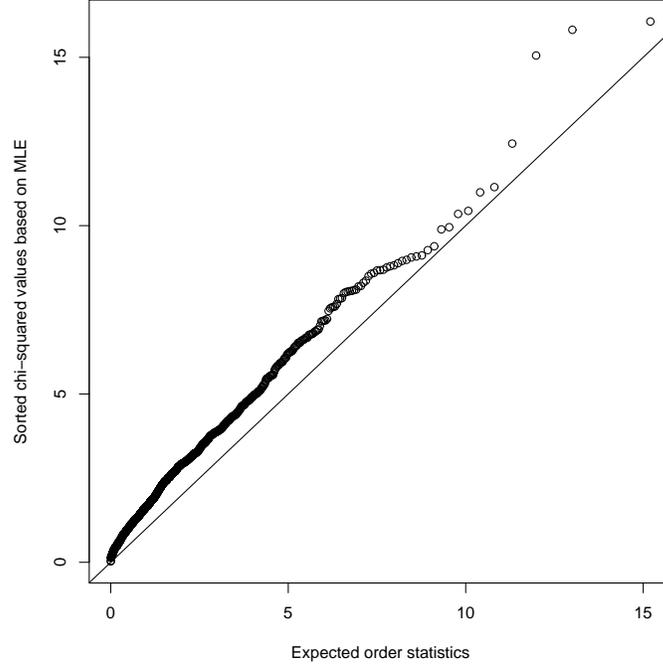


FIGURE 2. Quantile-quantile plot of  $\hat{R}$  values for i.i.d. normal data. Values of  $\hat{R}$  displayed in this plot were each determined from a separate sample of 50 standard normal deviates, and are plotted against the expected order statistics from a  $\chi^2_2$  distribution.

As before, each vector of  $t$  variates defines a posterior distribution on  $\mu$  and  $\sigma^2$ .

From this posterior, the test statistic  $A$ , defined as

$$(7) \quad A = \Pr_{\theta|\mathbf{y}}(R^B(\theta) > X), \quad X \sim \chi^2_{K-1},$$

can be approximated in a straightforward way using Monte Carlo integration.

The value of  $A$  generated from a particular sample  $\mathbf{Y}$  provides a convenient measure for model adequacy. Informally,  $A$  can be compared to its nominal value of 0.5, with values close to 0.5 suggesting adequate model fit. Values of  $A$  exceeding, say, 0.67 indicate that the posterior odds that a sampled value of  $R^B(\tilde{\theta})$  is greater

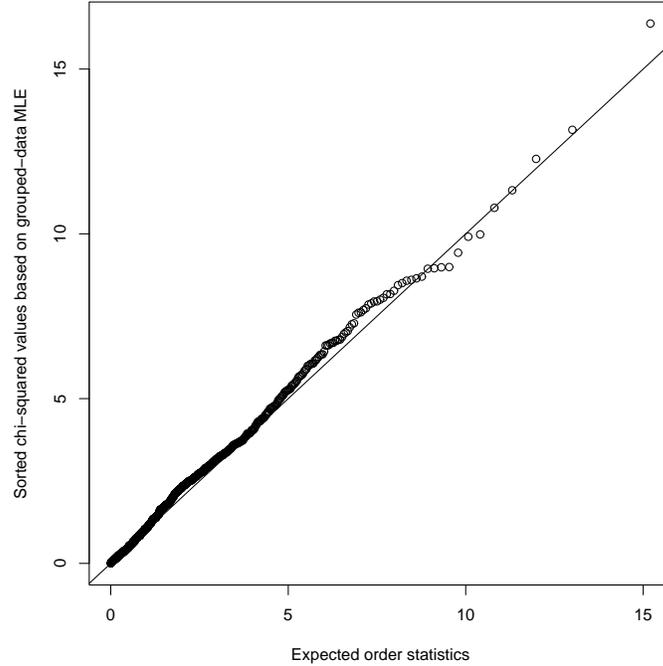


FIGURE 3. Quantile-quantile plot of  $R^g$  values for i.i.d. normal data. Values of  $R^g$  displayed in this plot were each determined from a separate sample of 50 standard normal deviates, and are plotted against the expected order statistics from their asymptotic  $\chi_2^2$  distribution.

than a random variable drawn from its nominal  $\chi^2$  distribution exceed 2:1. Values greater than 0.6 or 0.7 thus raise concern over model adequacy. These concerns are heightened when the posterior distribution of  $R^B(\tilde{\theta})$  is also highly concentrated above 0.5.

More formal model assessment can be based on approximating the sampling distribution of  $A$  using “posterior-predictive-posterior” model checks. That is, sampled values  $\tilde{\theta}$  from the posterior can be used to generate simulated data according to  $f(\cdot|\tilde{\theta})$ . Denoting the posterior-predictive data vector generated in this fashion by  $\mathbf{Y}^{pp}$ , posterior-predictive-posterior values of  $A^{pp}$  can be generated for each value

of  $\mathbf{Y}^{pp}$  by averaging  $R^B$ , computed from  $\mathbf{Y}^{pp}$ , over the posterior distribution on  $\boldsymbol{\theta}$  induced by  $\mathbf{Y}^{pp}$ . The value of  $A$  obtained for the original data vector can then be compared to the empirical distribution of the  $A^{pp}$ .

In principle, exactly this procedure can be implemented to calculate the probability that the test statistic  $A$ , based on a random sample of  $t$  variates, falls into the critical region of a test based on the empirical distribution of sampled values  $A^{pp}$ . In this case, however, it is not necessary to generate values of  $A^{pp}$  for each sample of  $t$  variates. Under this normal model, values of  $R^B$  are invariant to shifts in location and scale of the data, so the sampling distribution of  $A$ , for any future draw of 50 i.i.d. normal deviates, can be approximated by the empirical distribution of  $A$  values obtained under the normal sampling scheme used at the beginning of this example. It follows that critical regions for significance tests based on  $A$  are exact under this model, save for the Monte Carlo error encountered in the empirical approximation of their distribution. A histogram estimate of the distribution of  $A$  for this experiment is depicted in Figure 4; values of  $A$  used in the construction of this plot were based on a Monte Carlo approximation of (7) using 10,000 posterior samples of  $\mu, \sigma^2$  for each randomly-generated data vector  $\mathbf{Y}$ .

Table 1 displays the proportion of times in 10,000 draws of  $t$  samples that the value of the test statistic  $A$  was larger than the .95 quantile of the sampled values of  $A^{pp}$ . For comparison, the observed power of the test based on the grouped-maximum-likelihood  $\chi^2$  statistic  $R^g$  at the 5% level is also shown. The five bins used in the definition of the latter statistic were defined as the quintiles of a standard normal distribution.

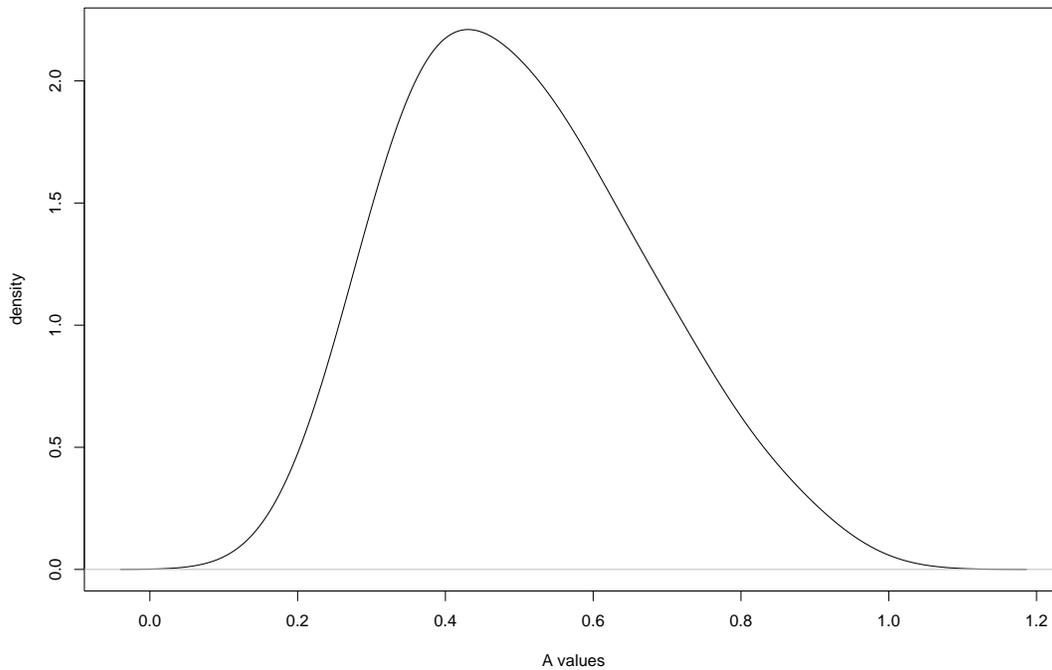


FIGURE 4. Density estimate of  $A$  for assessing goodness-of-fit for independent normal samples of size 50.

From Table 1, it is clear that the test statistic  $A$  offers substantially better power than  $R^g$  against this class of alternative models. Part of this advantage stems from the symmetry and unimodality of the alternative hypotheses, which  $R^g$  is ill-equipped to accommodate, and part from the fact that the bins used in the definition of  $R^g$  were fixed according to the null hypothesis. Perhaps a fairer comparison of the two statistics is provided in Table 2, which is based on repeatedly drawing random samples of size 50 from gamma distributions. In this case, the power of the two statistics is more comparable, though tests based on  $A$  again exhibit higher power. And, of course, tests based on  $A$  provide essentially exact levels of significance, a property not shared with  $R^g$ .

| Degrees of Freedom | Power of $A$ | Power of $R^g$ |
|--------------------|--------------|----------------|
| 1                  | 0.99         | 0.13           |
| 2                  | 0.79         | 0.05           |
| 3                  | 0.39         | 0.05           |
| 4                  | 0.32         | 0.05           |
| 5                  | 0.20         | 0.05           |
| 6                  | 0.15         | 0.05           |
| 7                  | 0.12         | 0.05           |
| 8                  | 0.10         | 0.05           |
| 9                  | 0.09         | 0.05           |
| 10                 | 0.08         | 0.05           |

TABLE 1. Power of test statistics  $A$  and  $R^g$  in detecting departures from normality when data are distributed according to  $t$  distributions. The first column in this table lists the degrees of freedom of the  $t$  distribution from which data were drawn. The second and third columns list the proportion of samples rejected in one-sided 5% tests.

| Gamma Shape | Power of $A$ | Power of $R^g$ |
|-------------|--------------|----------------|
| 1           | 0.84         | 0.61           |
| 2           | 0.50         | 0.23           |
| 3           | 0.31         | 0.15           |
| 4           | 0.24         | 0.12           |
| 5           | 0.18         | 0.09           |
| 6           | 0.16         | 0.09           |
| 7           | 0.14         | 0.08           |
| 8           | 0.13         | 0.08           |
| 9           | 0.11         | 0.07           |
| 10          | 0.11         | 0.07           |

TABLE 2. Power of test statistics  $A$  and  $R^g$  in detecting departures from normality when data are distributed according to gamma distributions. The first column in this table lists the shape parameter of the gamma distribution from which data were drawn. The second and third columns provide the proportion of samples rejected in one-sided 5% tests. Ten thousand samples of size 50 were used to estimate the power for each entry in the table. Bins used in the definition of  $R^g$  were based on quintiles of a normal distribution with mean and variance matched to the gamma distribution from which data were generated.

**3.2. A Random Effects Model.** To further assess the accuracy of the asymptotic distribution of  $R^B$  in normal error models, consider next a simple random effects model. The particular model considered can be specified as

$$\mu \sim N(0, 1) \quad \beta_i | \sigma_r^2 \sim N(0, \sigma_r^2), \quad i = 1, \dots, 25$$

$$Y_{i,j} | \beta_i, \sigma^2 \sim N(\mu + \beta_i, \sigma^2) \quad j = 1, \dots, 4 \quad \sigma_r^2, \sigma^2 \sim IG(1, 1).$$

One thousand simulations from this model were performed, resulting in 1,000  $\mathbf{Y}$  vectors each of length 100. For each of these simulated vectors, a Gibbs sampler was used to generate draws from the posterior distribution on  $\mu, \boldsymbol{\beta}, \sigma^2$  and  $\sigma_r^2$ . After a burn-in of 1,000 iterations, a parameter value was drawn from the posterior distribution and used to calculate a value of  $R^B$  according to (3) and (4). Ten equal probability bins were used to define  $R^B$ , and individual observations  $y_{i,j}$  were assigned to bin  $k$  when

$$\Phi\left(\frac{y_{ij} - \mu - \beta_i}{\sigma}\right) \in \left(\frac{k-1}{10}, \frac{k}{10}\right).$$

Figure 5 depicts the quantile-quantile plot constructed for these values of  $R^B$ . The agreement between the sampled  $R^B$  values and their asymptotic distribution appears excellent, despite the fact that only 100 observations were used to estimate 28 parameters (note that the random effect parameters are used to assign observations to bins and so have been counted as parameters), and that the values of the random effects were essentially based on only four observations each.

**3.3. Lip Cancer Data.** Spiegelhalter, Best, Carlin, and van der Linde (2002) describe a re-analysis of lip cancer incidence data originally considered by Clayton

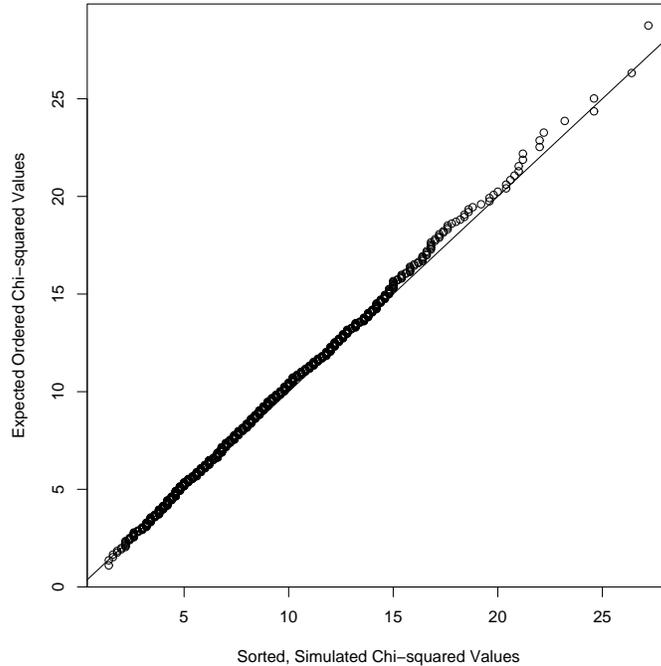


FIGURE 5. Quantile-quantile plot of  $R^B$  values for random effects model. Values of  $R^B$  displayed in this plot were each determined from an independent draw from the random effects model, and are plotted against the expected order statistics from a  $\chi_9^2$  distribution.

and Kaldor (1987). Their purpose in examining these data was to illustrate the use of the deviance information criterion (DIC) to select from among five potential models for the number of lip cancer cases,  $Y_i$ , observed in 56 Scottish districts as a function of available age and sex adjusted expected rates  $E_i$ . These data and models are reconsidered here for the related purpose of assessing which of the models provides an adequate probabilistic description of the data.

Following the Spiegelhalter *et al* analysis of these data, begin by assuming that  $Y_i$  is Poisson with mean  $\mu_i = \exp(\theta_i)E_i$ . Five models for  $\theta_i$  are considered:

- (1)  $\theta_i = \alpha_0$ ,  $\alpha_0$  a constant,

- (2)  $\theta_i = \alpha_0 + \gamma_i$ ,  $\gamma_i$  exchangeable random effects,
- (3)  $\theta_i = \alpha_0 + \delta_i$ ,  $\delta_i$  spatial random effects with a conditional autoregressive prior (e.g., Besag 1974),
- (4)  $\theta_i = \alpha_0 + \delta_i + \gamma_i$ ,  $\delta_i$  and  $\gamma_i$  as above, and
- (5)  $\theta_i = \alpha_i$ ,  $\alpha_i$  uniform on  $(-\infty, \infty)$ .

Five thousand, thinned posterior samples of  $\boldsymbol{\mu} = \{\mu_i\}$  were generated for each of these models using WinBUGS code (Spiegelhalter, Thomas and Best 2000) kindly provided by Nicola Best. For each sampled value of  $\mu_i$ , the Poisson counts  $y_i$  were assigned to one of eight, equally spaced probability bins according to the Poisson distribution function evaluated at  $y_i$  for the given value of  $\mu_i$ . In those cases for which the probability mass function assigned to  $y_i$  spanned more than one bin, allocation to a single bin was performed randomly according to the proportion of mass assigned to the bins. Averaging over all posterior samples of  $\boldsymbol{\mu}$  for a given model yielded the values of  $A$  depicted in Table 3.

| Model | A     | DIC   |
|-------|-------|-------|
| 1     | 0.999 | 382.7 |
| 2     | 0.514 | 104.0 |
| 3     | 0.520 | 89.9  |
| 4     | 0.527 | 89.7  |
| 5     | 0.667 | 111.7 |

Table 3. Values of the goodness-of-fit statistic  $A$  for potential models of lip cancer incidence data. DIC values obtained under the “mean” parameterization are listed for comparison.

The value of  $A$  cited in Table 3 for the first model, corresponding to a constant value of  $\alpha_0$  across districts, provides a clear indication of the inadequacy of that model. Lack-of-fit in this instance can be attributed to the failure of the model to adjust for district effects; the posterior mean of the number of counts assigned to the eight bins was (14.9, 1.4, 3.7, 3.4, 3.4, 4.5, 4.0, 20.6).

Similarly, the values of  $A$  reported in rows 2–4 suggest adequate fit for these models.

The most interesting value of  $A$  in this example is that reported for the last model, which corresponds to fitting a separate Poisson model for each observation.

At first glance, one might suspect that this suspicious value of  $A$  arises from overfitting. However, the last model generates the most disperse posterior distribution of any of the models considered, since only one observation figures into the marginal posterior of each  $\mu_i$ . Instead, the difficulty with this model arises from the prior assumptions made on  $\boldsymbol{\mu}$ . The assumption of a uniform prior on  $\theta_i$  implies a prior for the mean of each Poisson observation proportional to  $1/\mu_i$ ; this prior shrinks the estimate of every  $\mu_i$  toward 0. This results in an overabundance of counts in the higher bins and larger than expected values of  $R^B$ . The posterior mean of the bin counts for this model was (5.0, 5.7, 6.2, 6.6, 7.0, 7.5, 8.1, 9.8). Re-fitting the fifth model with noninformative priors proportional to  $1/\sqrt{\mu_i}$  yielded a value of  $A = 0.492$ .

It is also interesting to compare the values of  $A$  with those provided for the DIC. Both statistics suggest inadequacy of the first model, though for different reasons. For the first model, the high value of  $A$  indicates that the data do not follow Poisson distributions with a common scaling of adjusted expected rates. The value of the

DIC statistic suggests either that the model does not fit the data or is not as precise in predicting the data as the other models considered. An advantage of  $A$  in this case is that its value is interpretable without fitting alternative models.

The comparatively large value of the DIC statistic for the second model can be attributed to greater dispersion in its posterior as compared to posterior dispersion of the third and fourth models, even though the exchangeable model appears to adequately represent variation in the observed data. The comparatively large value of DIC reported for the fifth model reflects some combination of lack-of-fit and a posterior that is more disperse than others considered.

**3.4. Two-by-two tables.** Although not proposed for application in  $2 \times 2$  tables, it is interesting to compare the performance of the Bayesian  $\chi^2$  test statistic and the derived statistic  $A$  to the test statistic  $R^g$  in this setting, its most common application.

Because a thorough examination of the power and veracity of reported significance levels for the broad class of contingency tables is beyond the scope of this paper, attention here is restricted to the analysis of realizations from two  $2 \times 2$  tables. The cell probabilities for these tables, one of which satisfies the independence assumption and one which does not, are provided in Table 4. Table 5 lists the achieved significance levels and power of the statistics  $R^g$  and  $A$  when samples of size 12 are repeatedly drawn from multinomial distributions defined by these  $2 \times 2$  tables. The critical region of tests based on  $A$  were determined from the empirical distribution of  $A^{pp}$  as described in Example 1. Specifically, for each realization of a table, posterior-predictive samples of counts were generated according to the posterior distribution on the marginal cell probabilities obtained under a uniform prior.

For each posterior-predictive sample so obtained, a value of  $A^{PP}$  was generated by averaging  $R^B$  over the posterior distribution based on the posterior-predictive table.

|                |               |               |
|----------------|---------------|---------------|
| $\frac{1}{12}$ | $\frac{1}{6}$ | $\frac{1}{4}$ |
| $\frac{1}{4}$  | $\frac{1}{2}$ | $\frac{3}{4}$ |
| $\frac{1}{3}$  | $\frac{2}{3}$ |               |

|               |               |               |
|---------------|---------------|---------------|
| $\frac{1}{6}$ | $\frac{1}{6}$ | $\frac{1}{3}$ |
| $\frac{1}{6}$ | $\frac{1}{2}$ | $\frac{2}{3}$ |
| $\frac{1}{3}$ | $\frac{2}{3}$ |               |

(a)
(b)

Table 4. Two-by-two tables: (a) independence model, (b) dependence model.

| Independence Model |      |          | Dependence Model |      |
|--------------------|------|----------|------------------|------|
| $R^g$              | A    | $\alpha$ | $R^g$            | A    |
| .008               | .009 | .010     | .046             | .060 |
| .046               | .050 | .050     | .153             | .171 |
| .107               | .102 | .100     | .277             | .248 |

Table 5. Achieved significance levels and power of the classical  $\chi^2$  test for independence ( $R^g$ ) and the test statistic  $A$  for the  $2 \times 2$  tables of Table 4. The nominal sizes of the tests are listed in the third column, labeled  $\alpha$ . Based on samples from 16 million  $2 \times 2$  tables with 12 observations each; 8,000 posterior samples were used to estimate  $A$  for each of these observations; 8,000 values of  $A^{PP}$  were generated for each posterior sample in order to approximate the sampling distribution of the corresponding  $A$ .

The values in Table 5 suggest that tests based on both  $R^B$  and  $A$  achieve approximately the nominal levels of significance for the null model of Table 4a, but

that tests based on  $A$  provide somewhat more favorable power properties against the alternative model specified in Table 4b.

Of course, the real advantages of tests based on the statistic  $A$  are only realized in more complicated models, where it becomes necessary to define bins separately for individual observations, where the dimension of the parameter vector is large, or where grouped maximum likelihood estimation is neither desirable nor feasible.

#### 4. DISCUSSION

Goodness-of-tests based on the statistics  $R^B$  and  $A$  provide a simple way of assessing the adequacy of model fit in many Bayesian models. Essentially, the only requirement for their use is that observations be conditionally independent. From a computational perspective, neither statistic requires specialized software, and both can be calculated in a straightforward way using output from existing MCMC algorithms.

Approximating the sampling distribution of  $A$ , though conceptually straightforward, does introduce an additional computational burden, but is necessary only when the achieved value of  $A$  is “significantly” larger than 0.5. Significance of  $A$  in this context has a natural interpretation in terms of the posterior probability that a sampled value of  $R^B$  exceeds a random variable drawn from its nominal  $\chi^2$  distribution. In this regard, values of  $A$  that are close to 0.5 may indicate adequate model fit for the purposes of a given analysis even if when the sampling distribution of  $A^{pp}$  would permit rejection of the model in a significance test.

Aside from applications in Bayesian model assessment, the  $\chi^2$  statistic proposed here can be extended, albeit somewhat awkwardly, to models estimated using maximum likelihood. In that setting, parameter values can be sampled from their

asymptotic normal distribution and used as if they were sampled from a posterior distribution. Although not entirely palatable from a classical perspective, such a procedure does provide a mechanism for conducting a (sub-optimal) goodness-of-fit test for complicated models in which alternative tests may be difficult to perform.

Monitoring values of  $R^B$  generated within a MCMC algorithm also provides a rudimentary convergence diagnostic for slow-mixing chains. In fact, exceedances of  $R^B$  over a pre-specified quantile from its null distribution can be incorporated formally into the convergence diagnostics proposed in Raftery and Lewis (1992). To the extent that such exceedances are adequately described by a two-state Markov chain, the use of  $R^B$  in this context eliminates the requirement to assess convergence on a parameter-by-parameter basis, as is normally done in Raftery and Lewis's diagnostic scheme. It also provides a natural mechanism for determining whether burn-in has occurred.

A less obvious, but perhaps equally important, use of the  $R^B$  statistic involves code verification. Many practitioners currently fit models using customized code written for their specific application, a practice that frequently results in coding errors that are difficult to detect. This problem can be largely overcome by simply monitoring the distribution of  $R^B$ , which, in my experience, tends to deviate substantially from its null distribution when a model has been misspecified or mis-coded.

#### APPENDIX: OUTLINES OF PROOFS OF THEOREMS AND COROLLARIES

The proof of Theorem 1 and Corollary 1 are based largely on the proof given in Chernoff and Lehman (1954) in establishing the asymptotic distribution of  $\hat{R}$ .

Assume that conditions specified in Cramér (1946, pages 426-427) and Chen (1995) apply. Cramér specifies conditions that are sufficient for establishing the distribution of the  $\chi^2$  goodness-of-fit statistic when evaluated at the parameter vector maximizing the likelihood estimate based on the grouped data, whereas Chen's conditions are sufficient for establishing the asymptotic normality of the posterior distribution. Essentially, these conditions require that the likelihood be a smooth function of the parameter vector  $\boldsymbol{\theta}$  in an open interval containing  $\boldsymbol{\theta}_0$ , that the posterior distribution concentrates around a point in this interval, that the information contained in the observations increases with sample size, and that the prior assign non-negligible mass to the interval containing  $\boldsymbol{\theta}$ . In addition, assume that all third-order partial derivatives of  $f(y|\boldsymbol{\theta})$  (or, in the case of the corollary,  $f_j(y|\boldsymbol{\theta})$ ) with respect to the components of  $\boldsymbol{\theta}$  exist and are bounded in an open interval containing  $\boldsymbol{\theta}_0$ . This condition is sufficient for guaranteeing (15). Finally, note that all expectations and statements regarding probabilistic orders of magnitude described below are computed under the sampling distribution of  $\mathbf{Y}$  given  $\boldsymbol{\theta}_0$ .

The following lemmas are needed.

**Lemma 1.** *Under the conditions stated above,*

$$(8) \quad \frac{1}{\sqrt{n}} \left( m_k(\tilde{\boldsymbol{\theta}}) - m_k(\hat{\boldsymbol{\theta}}) \right) = \frac{1}{\sqrt{n}} \left( m_k^*(\tilde{\boldsymbol{\theta}}) - m_k^*(\hat{\boldsymbol{\theta}}) \right) + o_p(1)$$

$$(9) \quad = \frac{1}{\sqrt{n}} \sum_{i=1}^s \frac{\partial m_k^*(\hat{\boldsymbol{\theta}})}{\partial \theta_i} (\tilde{\theta}_i - \hat{\theta}_i) + o_p(1),$$

where

$$m_k^*(\boldsymbol{\theta}) = n \mathbf{E} \left[ \text{Ind} \left( Y \in [F^{-1}(a_{k-1} | \boldsymbol{\theta}), F^{-1}(a_k | \boldsymbol{\theta})] \right) \right].$$

*Proof of Lemma 1:*

Expanding  $m_k^*(\tilde{\theta})$  in a Taylor series expansion about  $m_k^*(\hat{\theta})$  yields

$$(10) \quad m_k^*(\tilde{\theta}) - m_k^*(\hat{\theta}) = \sum_{i=1}^s \frac{\partial m_k^*(\hat{\theta})}{\partial \theta_i} (\tilde{\theta}_i - \hat{\theta}_i) + O_p(1/n).$$

Define

$$\Delta z_{k,j} = z_{k,j}(\tilde{\theta}) - z_{k,j}(\hat{\theta}).$$

Then

$$\begin{aligned} |\Delta z_{k,j}| \leq & \text{Ind} \left( Y_j \in \left[ \min(F^{-1}(a_{k-1} | \tilde{\theta}), F^{-1}(a_{k-1} | \hat{\theta})), \max(F^{-1}(a_{k-1} | \tilde{\theta}), F^{-1}(a_{k-1} | \hat{\theta})) \right] \right) \\ & + \text{Ind} \left( Y_j \in \left[ \min(F^{-1}(a_k | \tilde{\theta}), F^{-1}(a_k | \hat{\theta})), \max(F^{-1}(a_k | \tilde{\theta}), F^{-1}(a_k | \hat{\theta})) \right] \right). \end{aligned}$$

Because  $(\hat{\theta} - \tilde{\theta})$  is  $O_p(1/\sqrt{n})$ ,  $\sqrt{n}\Delta z_{k,j} = O_p(1)$ . By the strong law of large numbers,

$$\sqrt{n} \sum_j \Delta z_{k,j} / n = \frac{(m_k(\tilde{\theta}) - m_k(\hat{\theta}))}{\sqrt{n}} = \frac{(m_k^*(\tilde{\theta}) - m_k^*(\hat{\theta}))}{\sqrt{n}} + o_p(1).$$

Substituting this expression into (10) yields (9).

**Corollary 3.** *The previous lemma also applies if  $\theta_0$  is substituted for  $\tilde{\theta}$ . That is,*

$$\begin{aligned} \frac{1}{\sqrt{n}} (m_k(\theta_0) - m_k(\hat{\theta})) &= \frac{1}{\sqrt{n}} (m_k^*(\theta_0) - m_k^*(\hat{\theta})) + o_p(1) \\ &= \frac{1}{\sqrt{n}} \sum_{i=1}^s \frac{\partial m_k^*(\hat{\theta})}{\partial \theta_i} (\tilde{\theta}_i - \hat{\theta}_i) + o_p(1). \end{aligned}$$

**Lemma 2.** *Define*

$$(11) \quad \hat{p}_k = F \left[ F^{-1}(a_k | \boldsymbol{\theta}_0) | \hat{\boldsymbol{\theta}} \right] - F \left[ F^{-1}(a_{k-1} | \boldsymbol{\theta}_0) | \hat{\boldsymbol{\theta}} \right] = \int_{F^{-1}(a_{k-1} | \boldsymbol{\theta}_0)}^{F^{-1}(a_k | \boldsymbol{\theta}_0)} f(y | \hat{\boldsymbol{\theta}}) dy.$$

*Then under the conditions stated above,*

$$(12) \quad \hat{p}_k - p_k = \frac{1}{n} \left( m_k^*(\boldsymbol{\theta}_0) - m_k^*(\hat{\boldsymbol{\theta}}) \right) + O_p \left( \frac{1}{n} \right)$$

*Proof of Lemma 2:*

For notational simplicity, define

$$G(\gamma, \delta; c) = F \left[ F^{-1}(c | \gamma) | \delta \right]$$

and

$$H_i(\gamma; c) = \frac{\partial G(\gamma, \delta; c)}{\partial \delta_i} \Big|_{\delta=\gamma}.$$

Then, noting that  $m_k^*(\boldsymbol{\theta}_0) = np_k = G(\boldsymbol{\theta}, \boldsymbol{\theta}, a_k) - G(\boldsymbol{\theta}, \boldsymbol{\theta}, a_{k-1})$ ,

$$\begin{aligned}
(\hat{p}_k - p_k) - \frac{1}{n} \left( m_k^*(\boldsymbol{\theta}_0) - m_k^*(\hat{\boldsymbol{\theta}}) \right) &= \left[ G(\boldsymbol{\theta}_0, \hat{\boldsymbol{\theta}}; a_k) - G(\boldsymbol{\theta}_0, \hat{\boldsymbol{\theta}}; a_{k-1}) \right] \\
&\quad + \left[ G(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0; a_k) - G(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0; a_{k-1}) \right] - 2p_k \\
&= \left[ \sum_i H_i(\boldsymbol{\theta}_0; a_k)(\hat{\theta}_i - \boldsymbol{\theta}_{0,i}) \right. \\
&\quad \left. - \sum_i H_i(\boldsymbol{\theta}_0; a_{k-1})(\hat{\theta}_i - \boldsymbol{\theta}_{0,i}) \right] \\
&\quad + \left[ \sum_i H_i(\hat{\boldsymbol{\theta}}; a_k)(\boldsymbol{\theta}_{0,i} - \hat{\theta}_i) \right. \\
&\quad \left. - \sum_i H_i(\hat{\boldsymbol{\theta}}; a_{k-1})(\boldsymbol{\theta}_{0,i} - \hat{\theta}_i) \right] + O_p\left(\frac{1}{n}\right) \\
&= \sum_i \left[ H_i(\boldsymbol{\theta}_0; a_k) - H_i(\hat{\boldsymbol{\theta}}; a_k) \right] (\hat{\theta}_i - \boldsymbol{\theta}_{0,i}) \\
&\quad - \sum_i \left[ H_i(\boldsymbol{\theta}_0; a_{k-1}) - H_i(\hat{\boldsymbol{\theta}}; a_{k-1}) \right] (\hat{\theta}_i - \boldsymbol{\theta}_{0,i}) + O_p\left(\frac{1}{n}\right) \\
&= \sum_h \sum_i \left[ \frac{\partial H_i(\boldsymbol{\theta}_0; a_k)}{\partial \boldsymbol{\theta}_{0,h}} - \frac{\partial H_i(\boldsymbol{\theta}_0; a_{k-1})}{\partial \boldsymbol{\theta}_{0,h}} \right] (\hat{\theta}_h - \boldsymbol{\theta}_{0,h})(\hat{\theta}_i - \boldsymbol{\theta}_{0,i}) \\
&\quad + O_p\left(\frac{1}{n}\right) \\
&= O_p\left(\frac{1}{n}\right)
\end{aligned}$$

**Corollary 4.**

$$\sqrt{n}(\hat{p}_k - p_k) = \frac{1}{\sqrt{n}} \left( m_k(\boldsymbol{\theta}_0) - m_k(\hat{\boldsymbol{\theta}}) \right) + O_p\left(\frac{1}{\sqrt{n}}\right)$$

**Proof of Theorem 1:**

Decompose the terms appearing in (3) as follows:

$$(13) \quad \frac{m_k(\tilde{\boldsymbol{\theta}}) - np_k}{\sqrt{np_k}} = \frac{m_k(\tilde{\boldsymbol{\theta}}) - m_k(\hat{\boldsymbol{\theta}})}{\sqrt{np_k}} - \frac{m_k(\boldsymbol{\theta}_0) - m_k(\hat{\boldsymbol{\theta}})}{\sqrt{np_k}} + \frac{m_k(\boldsymbol{\theta}_0) - np_k}{\sqrt{np_k}}.$$

From the first lemma and corollary, the first two terms on the right side of (13) are asymptotically equivalent to

$$(14) \quad \frac{\sum_i \frac{\partial m_k^*(\hat{\boldsymbol{\theta}})}{\partial \theta_i} (\tilde{\theta}_i - \hat{\theta}_i)}{\sqrt{np_k}} \quad \text{and} \quad \frac{\sum_i \frac{\partial m_k^*(\hat{\boldsymbol{\theta}})}{\partial \theta_i} (\boldsymbol{\theta}_{0,i} - \hat{\theta}_i)}{\sqrt{np_k}}.$$

Also,  $(\tilde{\boldsymbol{\theta}} - \hat{\boldsymbol{\theta}})$  is asymptotically normal with mean  $\mathbf{0}$  and covariance matrix equal to the negative inverse of the information matrix (Chen 1985). So, too, is  $(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)$ , and

$$(15) \quad \mathbf{E}[(\tilde{\boldsymbol{\theta}} - \hat{\boldsymbol{\theta}})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)] = o(1/n^2)$$

(e.g., Olver 1974, Cox 1974).

Following Chernoff and Lehmann (1954), define  $\boldsymbol{\epsilon}$  to be a  $K \times 1$  vector with components

$$\epsilon_k = \frac{m_k(\boldsymbol{\theta}_0) - np_k}{\sqrt{np_k}},$$

and let  $\hat{\boldsymbol{\nu}}$  be the vector with components

$$\hat{\nu}_k = \sqrt{n}(\hat{p}_k - p_k)/\sqrt{p_k}.$$

It follows from their results that

$$(16) \quad \hat{\boldsymbol{\nu}} = \mathbf{D}(\tilde{\mathbf{J}} + \mathbf{J}^*)^{-1}(\mathbf{D}'\boldsymbol{\epsilon} + \sqrt{n}\mathbf{A}^*) + o_p(1),$$

where  $\mathbf{J}^*$  is the matrix whose  $(i, j)$ th component is

$$\mathbf{E} \left[ \frac{\partial \log g(y | z, \theta)}{\partial \theta_i} \frac{\partial \log g(y | z, \theta)}{\partial \theta_j} \right],$$

$g(y | \mathbf{z}, \boldsymbol{\theta})$  is the conditional distribution of  $y$  given  $\mathbf{z}$  and  $\boldsymbol{\theta}$ ,  $\tilde{\mathbf{J}} \equiv \mathbf{D}'\mathbf{D}$  is the matrix with elements

$$\sum_{k=1}^K \frac{1}{p_k} \frac{\partial p_k}{\partial \theta_a} \frac{\partial p_k}{\partial \theta_b},$$

and  $\mathbf{A}^*$  is the vector whose  $a$ th component is

$$\frac{1}{n} \sum_{j=1}^n \frac{\partial \log g(y | \mathbf{z}_j, \theta)}{\partial \theta_a}.$$

From the second corollary, the right-hand side of (16) also describes the large sample distribution of  $(m_k(\boldsymbol{\theta}_0) - m_k(\hat{\boldsymbol{\theta}}))/\sqrt{np_k}$ .

Taking  $\boldsymbol{\eta} = \sqrt{n}\mathbf{A}^*$  and invoking the central limit theorem, Chernoff and Lehman note that the asymptotic distribution of  $(\epsilon, \eta)$  is

$$(17) \quad N \left[ 0, \begin{pmatrix} \mathbf{I} - \mathbf{q}\mathbf{q}' & 0 \\ 0 & \mathbf{J}^* \end{pmatrix} \right],$$

where  $\mathbf{q}$  is the vector with components  $\sqrt{p_k}$ . Letting  $\varepsilon$  denote a variable having the same distribution as  $\epsilon$ , and  $\tau$  a variable having the same distribution as  $\eta$ , with all four variables distributed independently, it follows that the  $R^B$  has the asymptotic distribution

$$(\mathbf{T}\varepsilon + \mathbf{S}\tau - \mathbf{T}\epsilon - \mathbf{S}\eta + \epsilon)' (\mathbf{T}\varepsilon + \mathbf{S}\tau - \mathbf{T}\epsilon - \mathbf{S}\eta + \epsilon),$$

where  $\mathbf{S} = \mathbf{D}(\tilde{\mathbf{J}} + \mathbf{J}^*)^{-1}$  and  $\mathbf{T} = \mathbf{S}\mathbf{D}'$ . Noting that  $\mathbf{D}'\mathbf{q} = \mathbf{0}$ , the asymptotic distribution of  $(\mathbf{T}\varepsilon + \mathbf{S}\tau - \mathbf{T}\epsilon - \mathbf{S}\eta + \epsilon)'$  is  $N(\mathbf{0}, \mathbf{I} - \mathbf{q}\mathbf{q}')$ . The result follows.

**Proof of Corollary 1:**

Because the proof of this corollary is similar to Theorem 1, only an outline is presented here.

To begin, note that Lemma 1 and Corollary 3 extend to this setting if  $m_k^*(\boldsymbol{\theta})$  is redefined as

$$m_k^*(\boldsymbol{\theta}) = \sum_{j=1}^n \mathbf{E} [\text{Ind}(Y_j \in (F_j^{-1}(a_{k-1} | \boldsymbol{\theta}), F_j^{-1}(a_k | \boldsymbol{\theta})))] .$$

Next, Lemma 2 applies if (11) is modified so that

$$(18) \quad \hat{p}_{j,k} = F_j [F_j^{-1}(a_k | \boldsymbol{\theta}_0) | \hat{\boldsymbol{\theta}}] - F_j [F_j^{-1}(a_{k-1} | \boldsymbol{\theta}_0) | \hat{\boldsymbol{\theta}}] = \int_{F_j^{-1}(a_{k-1} | \boldsymbol{\theta}_0)}^{F_j^{-1}(a_k | \boldsymbol{\theta}_0)} f_j(y | \hat{\boldsymbol{\theta}}) dy,$$

where  $p_{j,k}$  and related estimates refer to the probability that the  $j$ th observation falls into the  $k$ th bin. Then

$$(19) \quad \hat{p}_{j,k} - p_{j,k} = \left( z_{j,k}^*(\boldsymbol{\theta}_0) - z_{j,k}^*(\hat{\boldsymbol{\theta}}) \right) + O_p\left(\frac{1}{n}\right)$$

where

$$z_{j,k}^*(\boldsymbol{\theta}) = \mathbf{E} [\text{Ind}(Y_j \in [F_j^{-1}(a_{k-1} | \boldsymbol{\theta}), F_j^{-1}(a_k | \boldsymbol{\theta}))]] .$$

Corollary 4 generalizes to

$$\frac{1}{\sqrt{n}} \sum_{j=1}^n (\hat{p}_{j,k} - p_{j,k}) = \frac{1}{\sqrt{n}} \sum_{j=1}^n \left( z_{k,j}(\boldsymbol{\theta}) - z_{k,j}(\hat{\boldsymbol{\theta}}) \right) + O_p\left(\frac{1}{\sqrt{n}}\right) .$$

Extending Chernoff and Lehman's (1954) result to the case of non-identically distributed random variables requires the following modifications of the definitions of variables used in the i.i.d. case. Let

$$\epsilon_j = \left( \frac{z_{j,1} - p_{j,1}}{\sqrt{np_{j,1}}}, \dots, \frac{z_{j,K} - p_{j,K}}{\sqrt{np_{j,K}}} \right)' \quad \epsilon = (\epsilon'_1, \dots, \epsilon'_n)',$$

$$\tilde{\mathbf{J}} = \left\| \sum_{\alpha=1}^n \sum_{r=1}^K \frac{1}{p_{\alpha,r}} \frac{\partial p_{\alpha,r}}{\partial \theta_i} \frac{\partial p_{\alpha,r}}{\partial \theta_j} \right\|,$$

$$\mathbf{D} = \begin{pmatrix} \frac{1}{\sqrt{p_{1,1}}} \frac{\partial p_{1,1}}{\partial \theta_1} & \dots & \frac{1}{\sqrt{p_{1,1}}} \frac{\partial p_{1,1}}{\partial \theta_s} \\ \vdots & \vdots & \\ \frac{1}{\sqrt{p_{1,K}}} \frac{\partial p_{1,K}}{\partial \theta_1} & \dots & \frac{1}{\sqrt{p_{1,K}}} \frac{\partial p_{1,K}}{\partial \theta_s} \\ \frac{1}{\sqrt{p_{2,1}}} \frac{\partial p_{2,1}}{\partial \theta_1} & \dots & \frac{1}{\sqrt{p_{2,1}}} \frac{\partial p_{2,1}}{\partial \theta_s} \\ \vdots & \vdots & \\ \frac{1}{\sqrt{p_{n,K}}} \frac{\partial p_{n,K}}{\partial \theta_1} & \dots & \frac{1}{\sqrt{p_{n,K}}} \frac{\partial p_{n,K}}{\partial \theta_s} \end{pmatrix},$$

$$\mathbf{P} = \left( \underbrace{\mathbf{I}_k \dots \mathbf{I}_k}_{n \text{ times}} \right),$$

$$\mathbf{J}^* = \left\| \mathbf{E} \left[ \left( \sum_{\alpha=1}^n \frac{\partial \log g_{\alpha}(y|z, \theta)}{\partial \theta_i} \right) \cdot \left( \sum_{\beta=1}^n \frac{\partial \log g_{\beta}(y|z, \theta)}{\partial \theta_j} \right) \right] \right\|,$$

$$A_i^* = \frac{1}{n} \sum_{j=1}^n \frac{\partial \log g_j(y|z, \theta)}{\partial \theta_i}, \quad \text{and} \quad \hat{\nu}_{j,r} = \frac{\hat{p}_{j,r} - p_{j,r}}{\sqrt{np_{j,r}}}.$$

Then

$$\hat{\nu} = \mathbf{D}(\tilde{\mathbf{J}} + \mathbf{J}^*)^{-1}(\mathbf{D}'\epsilon + \sqrt{n}\mathbf{A}^*) + o_p(1).$$

The covariance matrix of  $\epsilon$  may be written

$$\frac{1}{n}\mathbf{I}_{n \times K} - \frac{1}{n} \begin{pmatrix} \mathbf{q}_1 \mathbf{q}_1' & \mathbf{0} & \dots & \mathbf{0} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{0} & \dots & \mathbf{0} & \mathbf{q}_n \mathbf{q}_n' \end{pmatrix},$$

where  $\mathbf{q}_i$  is the vector whose  $j$ th component is  $\sqrt{p_{i,j}}$ . Denote the rightmost matrix in this equation by  $\mathbf{Q}$ . Similarly, define  $\boldsymbol{\eta} = \sqrt{n}\mathbf{A}$ . Then the asymptotic distribution of  $\boldsymbol{\eta}$  has mean  $\mathbf{0}$  and covariance matrix equal to  $\mathbf{J}^*$ , and is uncorrelated with  $\epsilon$ .

Letting  $\hat{\mathbf{r}}$  denote the vector with components  $(z_{k,j}(\boldsymbol{\theta}) - z_{k,j}(\hat{\boldsymbol{\theta}}))/(\sqrt{np_{j,k}})$ , it follows from the generalization of Corollary 4 that the distribution of  $\mathbf{P}\hat{\mathbf{r}}$  is asymptotically the same as  $\mathbf{P}\hat{\nu}$ . Letting  $\tilde{\mathbf{r}}$  denote the vector with components  $(z_{k,j}(\tilde{\boldsymbol{\theta}}) - z_{k,j}(\hat{\boldsymbol{\theta}}))/(\sqrt{np_{j,k}})$ , then  $\mathbf{P}\tilde{\mathbf{r}}$  and  $\mathbf{P}\hat{\mathbf{r}}$  are, for large  $n$ , uncorrelated and identically distributed. Noting that

$$R^B = (\epsilon - \hat{\mathbf{r}} + \tilde{\mathbf{r}})' \mathbf{P}' \mathbf{P} (\epsilon - \hat{\mathbf{r}} + \tilde{\mathbf{r}})$$

and that  $\mathbf{D}'\mathbf{Q} = \mathbf{0}$ , some algebra and application of the central limit theorem yields the desired result.

### Proof of Corollary 2:

Expanding the components of  $R^B(\tilde{\boldsymbol{\theta}})$  yields

$$(20) \quad \frac{m_k - np_k(\tilde{\boldsymbol{\theta}})}{\sqrt{n}} = \frac{m_k - np_k(\boldsymbol{\theta}_0)}{\sqrt{n}} - \frac{p_k(\tilde{\boldsymbol{\theta}}) - p_k(\boldsymbol{\theta}_0)}{\sqrt{n}} - \frac{p_k(\boldsymbol{\theta}_0) - p_k(\hat{\boldsymbol{\theta}})}{\sqrt{n}}.$$

Asymptotically, Taylor series expansions show that the second term on the right side of this equation has the distribution of  $\mathbf{T}\epsilon + \mathbf{S}\eta$  described in the proof of Theorem 1, while the third term has the distribution of  $\mathbf{T}\epsilon + \mathbf{S}\tau$ . The result follows using methodology in the proof of Theorem 1.

## REFERENCES

- Bayarri, M.J. and Berger, J.O. (2000), “ $P$  Values for Composite Null Models,” *Journal of the American Statistical Association*, **95**, 1127-1142.
- Berger, J.O. and Pericchi, L.R. (1996), “The intrinsic Bayes factor for model selection and prediction,” *Journal of the American Statistical Association*, **91**, 109-122.
- Besag, J. (1974), “Spatial interaction and the statistical analysis of lattice systems (with discussion),” *Journal of the Royal Statistical Society, Series B*, **36**, 192-236.
- Chen, C.F. (1985), “On asymptotic normality of limiting density functions with Bayesian Implications,” *Journal of the Royal Statistical Society, Series B*, **97**, 540-546.
- Chernoff, H. and Lehmann, E.L. (1954), “The use of maximum likelihood estimates in  $\chi^2$  tests for goodness of fit,” *Annals of Mathematical Statistics*, **25**, 579-586.
- Clayton, D.G. and Kaldor, J. (1987), “Empirical Bayes estimates of age-standardized relative risks for use in disease mapping,” *Biometrics*, **43**, 671-681.
- Cox, D.R. (1974), *Theoretical Statistics*, Chapman and Hall: London.
- Cramér, H. (1946), *Mathematical Methods of Statistics*, Princeton University Press: Princeton, N.J.

De la Horra, J. and Rodríguez, M.T. (1997), "Asymptotic behavior of the posterior predictive  $P$ -value," *Communications in Statistics, Part A—Theory and Methods*, **26**, 2689-2699.

Gelfand, A.E. (1996), "Model determination using sampling-based methods", in *Markov Chain Monte Carlo in Practice*, editors W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, Chapman & Hall: London, 145-162.

Gelman, A. and Meng, X.-L. (1996), "Model checking and model improvement," in *Markov Chain Monte Carlo in Practice*, editors W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, Chapman & Hall: London, 189-202.

Gelman, A., Meng, X.-L., and Stern, H. (1996), "Posterior predictive assessment of model fitness via realized discrepancies," with discussion, *Statistica Sinica*, **6**, 733-807.

Guttman, I. (1967), "The use of the concept of a future observation in goodness-of-fit problems," *Journal of the Royal Statistical Society, Series B*, **29**, 83-100.

Hanley, J.A. and McNeil, B.J. (1982), "The Meaning and Use of the Area under a Receiver Operating Characteristic (ROC) Curve," *Radiology*, 29-36.

Moore, D.S. and Spruill, M.C. (1975), "Unified large-sample theory of general chi-squared statistics for tests of fit," *The Annals of Statistics*, **3**, 599-616.

O'Hagan, A. (1995), "Fractional Bayes factors for model comparison," with discussion, *Journal of the Royal Statistical Society, Series B*, **57**, 99-118.

Olver, F.W.J. (1974), *Asymptotics and Special Functions*, Academic Press: New York.

Pearson, K. (1900), "On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be

reasonably supposed to have arisen from random sampling”, *Philosophy Magazine*, 5th Series, **50**, 157-172.

Raftery, A.E. and Lewis, S. (1992), “How many iterations in the Gibbs sampler?,” in *Bayesian Statistics 4*, eds. J.M. Bernardo, J. Berger, A.P. Dawid, and A.F.M. Smith, Oxford University, U.K.: Oxford University Press, 763-773.

Robert, C.P. and Rousseau, J. (2002), “A Mixture Approach to Bayesian Goodness of Fit,” preprint.

Robins, J.M., van der Vaart, A., and Ventura, V. (2000), “Asymptotic distribution of  $P$  values in Composite Null Models,” *Journal of the American Statistical Association*, **95**, 1143-1159.

Rubin, D.B. (1984), “Bayesianly justifiable and relevant frequency calculations for the applied statistician,” *The Annals of Statistics*, **12**, 1151-1172.

Spiegelhalter, D., Best, N., Carlin, B. and van der Linde, A. (2002), “Bayesian measures of model complexity and fit,” *Journal of the Royal Statistical Society, Series B*, **64**, 583-616.

Spiegelhalter, D., Thomas, A. and Best, N. (2000), *WinBUGS Version 1.3 Users Manual*, Cambridge: Medical Research Council Biostatistics Unit. (Available from <http://www.mrc-bsu.cam.ac.uk/bugs>.)

Verdinelli, I. and Wasserman, L. (1998), “Bayesian goodness-of-fit testing using infinite-dimensional exponential families,” *The Annals of Statistics*, **26**, 1215-1241.