

CONSULTANT'S FORUM

Determining the Effective Sample Size of a Parametric Prior

Satoshi Morita,^{1,*} Peter F. Thall,² and Peter Müller²

¹Department of Epidemiology and Health Care Research, Kyoto University

Graduate School of Medicine, Kyoto 606-8501, Japan

²Department of Biostatistics, The University of Texas, M. D. Anderson Cancer Center,
Houston, Texas 77030, U.S.A

*email: satoshi_morita@pbh.med.kyoto-u.ac.jp

SUMMARY. We present a definition for the effective sample size of a parametric prior distribution in a Bayesian model, and propose methods for computing the effective sample size in a variety of settings. Our approach first constructs a prior chosen to be vague in a suitable sense, and updates this prior to obtain a sequence of posteriors corresponding to each of a range of sample sizes. We then compute a distance between each posterior and the parametric prior, defined in terms of the curvature of the logarithm of each distribution, and the posterior minimizing the distance defines the effective sample size of the prior. For cases where the distance cannot be computed analytically, we provide a numerical approximation based on Monte Carlo simulation. We provide general guidelines for application, illustrate the method in several standard cases where the answer seems obvious, and then apply it to some nonstandard settings.

KEY WORDS: Bayesian analysis; Computationally intensive methods; Effective sample size; Epsilon-information prior; Parametric prior distribution.

1. Introduction

A fundamental question in any Bayesian analysis is the amount of information contained in the prior. For many commonly used models, the answer seems straightforward. For example, it can be argued that a beta(a, b) distribution has effective sample size (ESS) $a + b$. This is based on the fact that a binomial variable Y from a sample of size n with success probability θ following a beta(a, b) prior implies a beta($a + Y, b + n - Y$) posterior. In other words, given a sample of size n , the prior sum $a + b$ becomes the posterior sum $a + b + n$. Thus, saying that a given beta(a, b) prior has ESS $m = a + b$ requires the implicit reasoning that the beta(a, b) may be identified with a beta($c + Y, d + m - Y$) posterior arising from a previous beta(c, d) prior having a very small amount of information. A simple way to formalize this is to set $c + d = \varepsilon$ for an arbitrarily small value $\varepsilon > 0$ and solve for $m = a + b - (c + d) = a + b - \varepsilon$.

More generally, one may match a given prior $p(\theta)$ with the posterior $q_m(\theta | Y)$ arising from an earlier prior $q_0(\theta)$ that is chosen to be vague in a suitable sense and that was updated by a sample of size m , and consider m to be the ESS of $p(\theta)$. In this general formulation, $p(\theta)$, $q_0(\theta)$, and $q_m(\theta | Y)$ play roles analogous to those of the beta(a, b), beta(c, d), and beta($a + Y, b + n - Y$) distributions given above. In some cases one may find the hyperparameters of $q_m(\theta | Y)$ as a function of m , compare $q_m(\theta | Y)$ with $p(\theta)$, and solve for m analytically. For

many parametric Bayesian models, however, this analytic approach does not work, and it is not obvious how to determine the ESS of the prior. A simple example is the usual normal linear regression model where the observed response variable Y for predictor X has mean $\beta_0 + \beta_1 X$ and variance σ^2 , so that $\theta = (\beta_0, \beta_1, \sigma^2)$. A traditional, technically convenient prior is that (β_0, β_1) is bivariate normal and σ^2 is inverse chi-squared, with hyperparameters chosen either for computational convenience or by elicitation. In either case, there is no obvious answer to the question of what the ESS of the prior may be. Moreover, for many commonly used choices of $q_0(\theta)$, the joint prior $p(\theta)$ cannot be matched with $q_m(\theta | Y)$ analytically.

Understanding the prior ESS is important when applying Bayesian methods in settings with a small to moderate sample size. For example, when fitting a Bayesian model to a data set of 10 observations, an a priori ESS of 1 is reasonable, whereas a prior ESS of 20 implies that the prior, rather than the data, dominates posterior inferences. If the prior is elicited from a domain expert, then an informative prior is desirable (Chaloner and Rhame, 2001; Garthwaite, Kadane, and O'Hagan, 2005). In contrast, if the prior is only a technically convenient ad hoc choice, as is often the case in practice, then understanding the ESS may prompt the investigator to reconsider the prior choice. Thus, it is important to have a good idea of the prior's ESS when interpreting one's inferences. This is especially important from the viewpoint of

defending Bayesian methods against the concern that the prior may inappropriately introduce artificial information.

In this article, we present a definition for the ESS of a prior $p(\theta)$ in a Bayesian parametric model, and we provide methods for computing the ESS in a wide variety of settings. Our approach relies on the idea of constructing an “ ε -information” prior $q_0(\theta)$, considering a sample \mathbf{Y} of size m and the posterior $q_m(\theta | \mathbf{Y})$, and computing a distance between $q_m(\theta | \mathbf{Y})$ and $p(\theta)$ in terms of the curvature (second derivatives) of $\log \{p(\theta)\}$ and $\log \{q_m(\theta | \mathbf{Y})\}$. The value of m minimizing the distance is the prior ESS. For cases where the distance cannot be computed analytically, we provide a numerical approximation based on Monte Carlo simulations from $q_m(\theta | \mathbf{Y})$. In cases where θ is multivariate, one may compute multiple ESSs, one associated with each of several subvectors of θ .

Section 2 presents a motivating application and defines ε -information priors and ESS. Computational methods are presented in Section 3. Section 4 gives guidelines for using ESS computations in specific settings. Applications are described in Sections 5 and 6, including discussions of connections between our proposed procedures and related methods given by Spiegelhalter, Freedman, and Parmar (1994), Ibrahim and Chen (2000), Hodges and Sargent (2001), and Spiegelhalter et al. (2002). We close with a brief discussion in Section 7.

2. Effective Sample Size

The following example illustrates why it may be useful to determine the ESS of a prior. We consider a design for a phase I trial to determine an optimal dose combination $X = (X_1, X_2)$ of two cytotoxic agents (Thall et al., 2003). The toxicity probability at X is given by the six-parameter model

$$\pi(X, \theta) = \frac{\alpha_1 X_1^{\beta_1} + \alpha_2 X_2^{\beta_2} + \alpha_3 (X_1^{\beta_1} X_2^{\beta_2})^{\beta_3}}{1 + \alpha_1 X_1^{\beta_1} + \alpha_2 X_2^{\beta_2} + \alpha_3 (X_1^{\beta_1} X_2^{\beta_2})^{\beta_3}}, \quad (1)$$

where all parameters in $\theta = (\alpha_1, \beta_1, \alpha_2, \beta_2, \alpha_3, \beta_3)$ are nonnegative. Under this model, if only agent 1 is administered at dose X_1 , with $X_2 = 0$, as in a single-agent phase I trial, then $\pi(X, \theta) = \pi_1(X_1, \theta_1) = \alpha_1 X_1^{\beta_1} / (1 + \alpha_1 X_1^{\beta_1})$ only depends on X_1 and $\theta_1 = (\alpha_1, \beta_1)$. Similarly, if $X_1 = 0$ then $\pi(X, \theta) = \pi_2(X_2, \theta_2) = \alpha_2 X_2^{\beta_2} / (1 + \alpha_2 X_2^{\beta_2})$ only depends on X_2 and $\theta_2 = (\alpha_2, \beta_2)$. The parameters $\theta_3 = (\alpha_3, \beta_3)$ characterize interactions that may occur when the two agents are used in combination. The model parameter vector thus is partitioned as $\theta = (\theta_1, \theta_2, \theta_3)$. Because phase I trials of combinations generally require that each agent be previously tested alone, it is natural to obtain informative priors on θ_1 and θ_2 , but assume a vague prior on θ_3 . Denoting by $\text{Ga}(a, b)$, the gamma distribution with mean a/b and variance a/b^2 , the elicitation process (Thall et al., 2003, Section 3) yielded the priors $\alpha_1 \sim \text{Ga}(1.74, 4.07)$, $\beta_1 \sim \text{Ga}(10.24, 1.34)$ for the effects of agent 1 alone, $\alpha_2 \sim \text{Ga}(2.32, 5.42)$, $\beta_2 \sim \text{Ga}(15.24, 1.95)$ for the effects of agent 2 alone, and $\alpha_3 \sim \text{Ga}(0.33, 0.33)$, $\beta_3 \sim \text{Ga}(0.0008, 0.0167)$ for the interaction parameters.

Because doses must be selected sequentially in phase I trials based on very small amounts of data, an important question is: Which ESS may be associated with the prior? Our proposed methods (Section 5, below) show that the overall ESS of this prior is $m = 1.5$. However, informative priors on θ_1 and θ_2 were obtained and a vague prior on θ_3 was desired, it also is important to determine the prior ESS of each sub-

vector. Applying our proposed methods yielded prior ESSs, $m_1 = 547.3$ for θ_1 , $m_2 = 756.8$ for θ_2 , and $m_3 = 0.01$ for θ_3 . The small value for m_3 confirms that the prior on θ_3 reflects little information about the interaction of the two agents. The large numerical discrepancy between $m = 1.5$ and $(m_1, m_2) = (547.3, 756.8)$ is desirable. It reflects the fact that, for each $i = 1, 2$, “ θ_i ” has a very different meaning in the submodel $\pi_i(X_i, \theta_i)$ parameterized by θ_i alone versus its meaning in the full six-parameter model $\pi(X, \theta)$. See, for, example, Berger and Pericchi (2001). From a geometric viewpoint, if $\pi(X, \theta)$ is thought of as a response surface varying as a function of the two-dimensional dose (X_1, X_2) , because the edges of the surface correspond to the submodels $\pi_1(X_1, \theta_1)$ where $X_2 = 0$ and $\pi_2(X_2, \theta_2)$ where $X_1 = 0$, the large values of m_1 and m_2 indicate that the locations of the edges were well known, whereas the small overall ESS $m = 1.5$ says that otherwise very little was known about the surface. In practice, one would report m_1, m_2, m_3 , and m to the clinician from whom the priors were elicited. The clinician could then judge whether m_1 and m_2 are reasonable characterizations of his/her prior information about the single agents, and compare m to the trial’s sample size. In the motivating application, a trial of gemcitabine and cyclophosphamide for advanced cancer, the large values of m_1 and m_2 were appropriate because there was substantial clinical experience with each single agent, and the small overall ESS also was appropriate because no clinical data on the two agents used together were available and a sample size of 60 patients was planned.

This example illustrates four key features of our proposed method, namely, that (1) ESS is a readily interpretable index of a prior’s informativeness, (2) it may be useful to compute ESSs for both the entire parameter vector and for particular subvectors, (3) ESS values may be used as feedback in the elicitation process, and (4) even when standard distributions are used, it may not be obvious how to define a prior’s ESS.

The intuitive motivation for the following construction is to mimic the rationale, given in Section 1, as to why the ESS of a beta(a, b) equals $a + b$. As a general Bayesian framework, let $f(Y | \theta)$ denote the probability distribution function (pdf) of an s -dimensional random vector Y , and let $p(\theta | \tilde{\theta})$ be the prior on the parameter vector $\theta = (\theta_1, \dots, \theta_d)$, where $\tilde{\theta}$ denotes the vector of hyperparameters. The likelihood of an independent and identically distributed (i.i.d.) sample $\mathbf{Y}_m = (Y_1, \dots, Y_m)$ is then given by $f_m(\mathbf{Y}_m | \theta) = \prod_{i=1}^m f(Y_i | \theta)$.

We define an ε -information prior $q_0(\theta | \tilde{\theta}_0)$ by requiring it to have the same mean, $E_{q_0}(\theta) = E_p(\theta)$, and correlations, $\text{Corr}_{q_0}(\theta_j, \theta_{j'}) = \text{Corr}_p(\theta_j, \theta_{j'})$, $j \neq j'$, as $p(\theta | \tilde{\theta})$, while inflating the variances of the elements of θ so that $\text{Var}_{q_0}(\theta_j) \gg \text{Var}_p(\theta_j)$, in such a way that $q_0(\theta | \tilde{\theta}_0)$ has small information but $\text{Var}_{q_0}(\theta_j)$ must exist for $j = 1, \dots, d$. Table 1 illustrates how to specify $q_0(\theta | \tilde{\theta}_0)$ for several standard parametric priors. Given the likelihood $f_m(\mathbf{Y}_m | \theta)$ and ε -information prior $q_0(\theta | \tilde{\theta}_0)$, we denote the posterior by $q_m(\theta | \tilde{\theta}_0, \mathbf{Y}_m) \propto q_0(\theta | \tilde{\theta}_0) f_m(\mathbf{Y}_m | \theta)$ and the marginal distribution under $p(\theta | \tilde{\theta})$ by

$$f_m(\mathbf{Y}_m | \tilde{\theta}) = \int f_m(\mathbf{Y}_m | \theta) p(\theta | \tilde{\theta}) d\theta. \quad (2)$$

When $\tilde{\theta}$ is fixed we write $f_m(\mathbf{Y}_m)$ for brevity. To define the ESS(s), consider the following three cases based on $p(\theta | \tilde{\theta})$. For implementation, we find it useful to distinguish between

Table 1

Examples of ε -information prior distributions. The hyperparameters c , c_1 , and c_2 are very large constants chosen to inflate the variances of the elements of θ under q_0 .

d	Distribution	$p(\theta \tilde{\theta})$	$q_0(\theta \tilde{\theta}_0)$
1	Beta	$\text{Be}(\tilde{\alpha}, \tilde{\beta})$	$\text{Be}(\tilde{\alpha}/c, \tilde{\beta}/c)$
1	Gamma	$\text{Ga}(\tilde{\alpha}, \tilde{\beta})$	$\text{Ga}(\tilde{\alpha}/c, \tilde{\beta}/c)$
1	Univariate normal with known variance	$N(\tilde{\mu}, \tilde{\sigma}^2)$	$N(\tilde{\mu}, c\tilde{\sigma}^2)$
1	Scaled inverse χ^2	$\text{Inv } \chi^2(\tilde{\nu}, \tilde{\sigma}^2)$	$\text{Inv } \chi^2(4 + c^{-1}, \tilde{\nu}\tilde{\sigma}^2/2(\tilde{\nu} - 2))$
2	Normal inverse χ^2	$N(\tilde{\mu}, \tilde{\sigma}^2/\tilde{\phi}) * \text{Inv } \chi^2(\tilde{\nu}, \tilde{\sigma}^2)$	$N(\tilde{\mu}, c\tilde{\sigma}^2/\tilde{\phi}) * \text{Inv } \chi^2(4 + c^{-1}, \tilde{\nu}\tilde{\sigma}^2/2(\tilde{\nu} - 2))$
3	Dirichlet	$\text{Dir}(\tilde{\alpha}_1, \tilde{\alpha}_2, \tilde{\alpha}_3)$	$\text{Dir}(\tilde{\alpha}_1/c, \tilde{\alpha}_2/c, \tilde{\alpha}_3/c)$
3	Multivariate normal	$\text{MVN}(\tilde{\mu}_1, \tilde{\mu}_2, \tilde{\sigma}_1^2, \tilde{\sigma}_2^2, \tilde{\sigma}_{12})$	$\text{MVN}(\tilde{\mu}_1, \tilde{\mu}_2, c_1^2\tilde{\sigma}_1^2, c_2^2\tilde{\sigma}_2^2, c_1c_2\tilde{\sigma}_{12})$

these cases although, formally, cases 1 and 2 are special instances of case 3.

Case 1: $d = 1$, with $p(\theta | \tilde{\theta})$ being a univariate parametric model. For this case, we will define one ESS. Examples include the beta, gamma, univariate normal with known variance, and inverse χ^2 distributions.

Case 2: $d \geq 2$ with $p(\theta | \tilde{\theta})$ being a d -variate parametric model. For this case, we will define one ESS. Examples include the Dirichlet and multivariate normal (MVN) distributions.

The following case deals with settings where it is scientifically appropriate to define two or more ESSs for $p(\theta | \tilde{\theta})$.

Case 3: $d \geq 2$ with $p(\theta | \tilde{\theta})$ written as a product of K parametric distributions, $p(\theta | \tilde{\theta}) = \prod_{k=1}^K p_k(\theta_k, | \tilde{\theta}_k, \theta_1, \dots, \theta_{k-1})$, where $\theta = (\theta_1, \dots, \theta_K)$ is partitioned into K subvectors, for $1 < K \leq d$. In this case, a vector of K ESSs, one for each subvector, may be meaningful. An example is a normal inverse χ^2 distribution where $(\theta_1, \theta_2) = (\sigma^2, \mu)$, the variance and mean of a normal sampling model, with $p(\theta_1, \theta_2) = p_1(\sigma^2) p_2(\mu | \sigma^2)$ and $\sigma^2 \sim \text{Inv } \chi^2(\tilde{\nu}, \tilde{\sigma}^2)$ and $\mu | \sigma^2 \sim N(\tilde{\mu}, \sigma^2/\tilde{\phi})$. Here $K = d = 2$ and the two subvectors of θ are the single parameters σ^2 and μ . We will discuss other examples in Sections 4 and 5.

To define the distance between $p(\theta | \tilde{\theta})$ and $q_m(\theta | \tilde{\theta}_0, \mathbf{Y}_m)$ in cases 1 and 2, the basic idea is to find the sample size, m , that would be implied by normal approximations of the prior $p(\theta)$ and the posterior $q_m(\theta | \tilde{\theta}_0, \mathbf{Y}_m)$. This led us to use the second derivatives of the log densities to define the distance. The real validation and justification of our definition, however, comes from comparing the resulting ESS values with the commonly reported ESS in standard settings. We carry out these comparisons in Section 5.

Let $\bar{\theta} = E_p(\theta)$ denote the prior mean under $p(\theta | \tilde{\theta})$. We define

$$D_{p,j}(\theta) = -\frac{\partial^2 \log\{p(\theta | \tilde{\theta})\}}{\partial \theta_j^2},$$

and

$$D_{q,j}(m, \theta, \mathbf{Y}_m) = -\frac{\partial^2 \log\{q_m(\theta | \tilde{\theta}_0, \mathbf{Y}_m)\}}{\partial \theta_j^2}, \quad j = 1, \dots, d.$$

Denote $D_{p,+}(\theta) = \sum_{j=1}^d D_{p,j}(\theta)$ and $D_{q,+}(m, \theta) = \sum_{j=1}^d \int D_{q,j}(m, \theta, \mathbf{Y}_m) f_m(\mathbf{Y}_m) d\mathbf{Y}_m$. We define a distance between $p(\theta | \tilde{\theta})$ and $q_m(\theta | \tilde{\theta}_0, \mathbf{Y}_m)$ for sample size m as the

difference of the trace of the two information matrices,

$$\delta(m, \bar{\theta}, p, q_0) = |D_{p,+}(\bar{\theta}) - D_{q,+}(m, \bar{\theta})|. \quad (3)$$

That is, we define the distance in terms of the trace of the information matrix (second derivative of the log density) of the prior $p(\theta | \tilde{\theta})$, and the expected information matrix of the posterior $q_m(\theta | \tilde{\theta}_0, \mathbf{Y}_m)$, where the expectation is with respect to the marginal $f_m(\mathbf{Y}_m)$. When $d = 1$, because the “+” subscript is superfluous, we write $D_p(\bar{\theta})$ and $D_q(m, \bar{\theta})$.

DEFINITION 1: *The ESS of $p(\theta | \tilde{\theta})$ with respect to the likelihood $f_m(\mathbf{Y}_m | \theta)$ is the integer m that minimizes the distance $\delta(m, \bar{\theta}, p, q_0)$.*

Algorithm 1, below, will generalize this to allow noninteger-valued m . An essential point is that the ESS is defined as a property of a prior and likelihood pair, so that, for example, a given prior might have two different ESS values in the context of two different likelihoods.

The definition of the distance (3) involves some arbitrary choices. We chose this definition after an extensive empirical investigation (not shown) of alternative formulations. Instead of evaluating the curvature at the prior mean, one could use the prior mode. Similarly, one could marginalize θ with respect to the prior, averaged over \mathbf{Y}_m with respect to $f_m(\mathbf{Y}_m | \theta)$ rather than the marginal $f_m(\mathbf{Y}_m)$, or use the determinant rather than the trace of the information matrix. One also could define $\delta(\cdot)$ in terms of Kullback–Liebler divergence, or variances. We investigated all of these alternatives and evaluated the resulting ESS in each of several standard cases, and found that the proposed distance (3) was best at matching the results that are commonly used as ESS values.

For case 3, a more general definition is required. A motivating example is the logistic regression model, $\text{logit}\{\pi(X, \theta)\} = \beta_0 + \beta_1 X_1 + \beta_2 X_2$, where $d = 3$, $\theta = (\beta_0, \beta_1, \beta_2)$ and $\beta_j \sim N(\tilde{\mu}_j, \tilde{\sigma}_j^2)$ independently with $\tilde{\theta} = (\tilde{\mu}_j, \tilde{\sigma}_j^2)$ for $j = 0, 1, 2$. In this case, the subvectors of interest are $\theta_1 = \beta_0$ and $\theta_2 = (\beta_1, \beta_2)$, so two ESS values, m_1 and m_2 , may be computed. To accommodate case 3, we generalize (3) by defining a set of K subvector-specific distances. Let γ_k be the set of indices of the elements of θ_k , and denote $D_{p,+}^k(\theta) = \sum_{j \in \gamma_k} D_{p,j}(\theta)$ and $D_{q,+}^k(m_k, \theta) = \sum_{j \in \gamma_k} \int D_{q,j}(m_k, \theta, \mathbf{Y}_{m_k}) f_{m_k}(\mathbf{Y}_{m_k}) d\mathbf{Y}_{m_k}$. For each $k = 1, \dots, K$, we define the distance between $p_k(\theta_k | \tilde{\theta}_k, \theta_1, \dots, \theta_{k-1})$ and $q_{m_k}(\theta_k | \tilde{\theta}_{0,k}, \mathbf{Y}_{m_k}, \theta_1, \dots, \theta_{k-1})$ to be

$$\delta_k(m_k, \bar{\theta}, p, q_0) = |D_{p,+}^k(\bar{\theta}) - D_{q,+}^k(m_k, \bar{\theta})|. \quad (4)$$

DEFINITION 2: Assume $p(\theta | \tilde{\theta})$ as in case 3. Let $m_k = \arg \min \delta_k(m, \bar{\theta}, p_2, q_0)$. We define (m_1, \dots, m_K) to be the ESSs for the prior $p(\theta | \tilde{\theta})$ with respect to the model $f_m(\mathbf{Y}_m | \theta)$ and the partition $\theta = (\theta_1, \dots, \theta_K)$.

3. Computational Methods

Let $\bar{\theta} = (\bar{\theta}_1, \dots, \bar{\theta}_d)$ denote the prior mean vector. With the following algorithms, we generalize Definitions 1 and 2 to allow noninteger ESS values.

Algorithm 1, for cases 1 and 2: Let M be a positive integer chosen so that, initially, it is reasonable to assume that $m \leq M$.

- Step 1. Specify $q_0(\theta | \tilde{\theta}_0)$.
- Step 2. For each $m = 0, \dots, M$, compute $\delta(m, \bar{\theta}, p, q_0)$.
- Step 3. The ESS is the interpolated value of m minimizing $\delta(m, \bar{\theta}, p, q_0)$.

In practice, step 2 is carried out either analytically or using the simulation-based numerical approximation as described in Section 3.2.

Algorithm 2, for case 3: For each $k = 1, \dots, K$, let M_k be a positive integer chosen so that, initially, it is reasonable to assume that $m_k \leq M_k$.

- Step 1. Specify $q_0(\theta | \tilde{\theta}_0) = \prod_{k=1}^K q_{0,k}(\theta_k | \tilde{\theta}_{0,k}, \theta_1, \dots, \theta_{k-1})$.
- Step 2. For each $k = 1, \dots, K$, and $m_k = 0, \dots, M_k$, compute $\delta_k(m_k, \bar{\theta}, p, q_0)$.
- Step 3. The ESS of θ_k is the interpolated value of m_k minimizing $\delta_k(m_k, \bar{\theta}, p, q_0)$.

If the hyperparameter $\tilde{\theta}$ of $p(\theta | \tilde{\theta})$ includes a degree of freedom (df) parameter $\tilde{\nu}$, as with an inverse χ^2 , inverse gamma, inverse Wishart, or t -distribution, then the corresponding hyperparameter of $q_0(\theta | \tilde{\theta}_0)$ is $\tilde{\nu}_0 = \tilde{\nu}_{\min} + \epsilon$, where $\tilde{\nu}_{\min}$ is the smallest integer that ensures the second moments of $\theta \sim q_0(\theta | \tilde{\theta}_0)$ exist and $\epsilon > 0$ is arbitrarily small. In such cases, we add $D_{q,+}(\tilde{\nu}_{\min}, \tilde{\theta}) - D_{q,+}(0, \tilde{\theta})$ to $D_{q,+}(m, \tilde{\theta})$ and add $D_{q,+}^k(\tilde{\nu}_{\min}, \tilde{\theta}) - D_{q,+}^k(0, \tilde{\theta})$ to $D_{q,+}^k(m_k, \tilde{\theta})$ to ensure that $\text{ESS} > \tilde{\nu}_{\min}$.

For each $m = 1, \dots, M$, when $\int D_{q,j}(m, \bar{\theta}, \mathbf{Y}_m) \times f_m(\mathbf{Y}_m) d\mathbf{Y}_m$ cannot be computed analytically, we use the following simulation-based approximation. Given $\bar{\theta} = E_p(\theta)$, we first simulate Monte Carlo sample $\theta^{(1)}, \dots, \theta^{(T)}$ from $p(\theta | \tilde{\theta})$ for large T , for example, $T = 100,000$. For each $t = 1, \dots, T$, simulate $Y_1^{(t)}, \dots, Y_M^{(t)}$ from $f_M(\mathbf{Y}_M | \theta^{(t)})$. Use the Monte Carlo average $T^{-1} \sum_{t=1}^T D_{q,j}(m, \bar{\theta}, \mathbf{Y}_m^{(t)})$ in place of $\int D_{q,j}(m, \bar{\theta}, \mathbf{Y}_m) f_m(\mathbf{Y}_m) d\mathbf{Y}_m$. For case 3, the same method is used to evaluate $D_{q,+}^k(m_k, \bar{\theta})$ in (4).

For regression models of Y as a function of a u -dimensional predictor X , we extend Definition 1 by augmenting the regression model with a probability distribution $g_m(\mathbf{X}_m | \xi)$ for the covariates and prior $r(\xi | \tilde{\xi})$, usually assuming independence, $g_m(\mathbf{X}_m | \xi) = \prod_{i=1}^m g(X_i | \xi)$. Then we define

$$f_m(\mathbf{Y}_m) = \int f_m(\mathbf{Y}_m | \mathbf{X}_m, \theta) g_m(\mathbf{X}_m | \xi) f(\theta | \tilde{\theta}) r(\xi | \tilde{\xi}) d\theta d\xi.$$

In this case, we simulate $\theta^{(1)}, \dots, \theta^{(T)}$ from $p(\theta | \tilde{\theta})$ and $\xi^{(1)}, \dots, \xi^{(T)}$ from $r(\xi | \tilde{\xi})$, then simulate each $X_1^{(t)}, \dots, X_M^{(t)}$

from $g_M(\mathbf{X}_M | \xi^{(t)})$, and $Y_i^{(t)}$ from $f(Y_i | \theta^{(t)}, X_i^{(t)})$ for each $i = 1, \dots, M$, to obtain $(Y_1^{(t)}, X_1^{(t)}), \dots, (Y_M^{(t)}, X_M^{(t)})$. Finally, we compute the Monte Carlo average $T^{-1} \sum_{t=1}^T D_{q,j}(m, \bar{\theta}, \mathbf{Y}_m^{(t)}, \mathbf{X}_m^{(t)})$. For case 3, the same method is used to evaluate $D_{q,+}^k(m_k, \bar{\theta})$ in (4).

4. Guidelines for Application

Before illustrating how the above methods for computing ESS may be applied in particular cases, we provide general guidelines for using ESS values in some commonly encountered settings of Bayesian inference.

1. *Prior elicitation.* When eliciting a prior from an area expert, ESS values may be provided as a readily interpretable form of feedback. The area expert may use this as a basis to modify his/her judgments, if desired, and this process may be iterated. For example, in the motivating example of Section 2, we would report the ESS values $m_1 = 547$ and $m_2 = 756$ to the investigator planning the study. If his/her prior were based on earlier single-agent trials with around 100 patients each, (s)he would be prompted to revise the replies to the prior elicitation questions.
2. *Formalizing uninformative priors.* Often an investigator wishes to formalize vague prior information. The ESS can be used to confirm that the chosen prior carries little information, as desired. For example, in the motivating example in Section 2.1, the reported ESS $m_3 = 0.01$ for the interaction parameter confirms that this prior is vague.
3. *Reviewing others' analyses.* When interpreting or formally reviewing a Bayesian data analysis, the ESS of the analyst's prior provides a tool for evaluating the reasonableness of the analysis. In particular, if it is claimed that a vague or uninformative prior was used, the ESS provides an objective index to evaluate this claim. If appropriate, one may alert the analyst if a prior appears to be overly informative. Similarly, if an informative prior based on historical data is used in the analysis, reporting the ESS enables the reviewer to verify that the prior data are given appropriate weight.
4. *Sensitivity analyses.* In performing a conventional Bayesian sensitivity analysis in which prior parameters are varied and corresponding posterior values of interest are computed, the ESS of each prior may be computed to enhance interpretation of this analysis. The ESS itself may be used as an index of prior informativeness in such a sensitivity analysis.
5. *Designing outcome-adaptive experiments.* When formulating a prior as part of a Bayesian model to be used in a sequentially outcome-adaptive experiment, the ESS may be used to calibrate the prior to ensure that the data, rather than the prior, will dominate early decisions during the trial.
6. *Reviewing Bayesian designs.* When interpreting or formally reviewing a Bayesian design, such as that given in a clinical trial protocol, the ESS of the prior provides a tool for determining the extent to which the prior may influence the design's decisions. Currently, an important reservation about using Bayesian inference in a

Table 2

Prior, likelihood, and corresponding posterior q_m with respect to the ε -information prior q_0 , and traditionally reported prior effective sample size, ESS, for some common models. In line three, we denote $s^2 = \sum_{i=1}^m (Y_i - \tilde{\nu}_0)^2$.

$p(\theta \tilde{\theta})$	$f(\mathbf{Y}_m \theta)$	$q_m(\theta \tilde{\theta}, \mathbf{Y}_m)$	ESS
$\text{Be}(\tilde{\alpha}, \tilde{\beta})$	$\text{Bin}(n, \theta)$	$\text{Be}(c^{-1}\tilde{\alpha} + Y, c^{-1}\tilde{\beta} + m - Y)$	$\tilde{\alpha} + \tilde{\beta}$
$\text{Ga}(\tilde{\alpha}, \tilde{\beta})$	$\text{Exp}(\theta)$	$\text{Ga}(c^{-1}\tilde{\alpha} + m, c^{-1}\tilde{\beta} + \sum Y_i)$	$\tilde{\alpha}$
$\text{Inv } \chi^2(\tilde{\nu}, \tilde{\sigma}^2)$	$\text{N}(0, \sigma^2)$	$\text{Inv } \chi^2(\tilde{\nu}_0 + m, \frac{\tilde{\nu}_0 \tilde{\sigma}^2 + s^2}{\tilde{\nu}_0 + m})$	$\tilde{\nu}$
$\text{Dir}(\tilde{\alpha})$	$\text{Mn}(n, \theta)$	$\text{Dir}(c^{-1}\tilde{\alpha} + \mathbf{S})$	$\sum \tilde{\alpha}_j$

regulatory environment, such as the planning of clinical trial protocols, is the difficulty of evaluating and judging the appropriateness of prior distributions in complex probability models. The ESS provides a useful tool to mitigate such concerns.

5. Validation with Standard Models

We validate the proposed definition of ESS by computing the implied sample sizes in standard models (Table 2) for which commonly reported prior-equivalent sample sizes exist. Following Gelman et al. (2004), we denote $\text{Be}(\alpha, \beta)$, $\text{Bin}(n, \theta)$, $\text{Ga}(\alpha, \beta)$, $\text{Exp}(\theta)$, $\text{N}(\mu, \sigma^2)$, $\text{Inv } \chi^2(\nu, s^2)$, $\text{Dir}(\alpha_1, \dots, \alpha_J)$, $\text{Mn}(n, \theta_1, \dots, \theta_J)$, and $\text{BeBin}(n, \alpha, \beta)$ for the beta, binomial, gamma, exponential, normal, scaled inverse χ^2 , Dirichlet, multinomial, and beta-binomial distributions. The corresponding ε -information priors are given in Table 1. For each model in Table 2, the reported ESS matches the obvious choice.

Example 1. Beta/binomial model. $\delta(m, \bar{\theta}, p, q_0) = \{(\tilde{\alpha} - 1)\bar{\theta}^{-2} + (\tilde{\beta} - 1)(1 - \bar{\theta})^{-2}\} - \{(\tilde{\alpha}/c + \sum_{Y=0}^m Y f_m(\mathbf{Y}_m) - 1)\bar{\theta}^{-2} + (\tilde{\beta}/c + m - \sum_{Y=0}^m Y f_m(\mathbf{Y}_m) - 1)(1 - \bar{\theta})^{-2}\}$, where $f_m(\mathbf{Y}_m) = \text{BeBin}(n, \tilde{\alpha}, \tilde{\beta})$ and $\bar{\theta} = \text{E}_p(\theta) = \tilde{\alpha}/(\tilde{\alpha} + \tilde{\beta})$. Figure 1 shows a plot of $\delta(m, \bar{\theta}, p, q_0)$ against m in the case $\bar{\theta} = (\tilde{\alpha}, \tilde{\beta}) = (3, 7)$. Using $\theta \sim \text{Be}(3, 7)$, the computed ESS is 10, matching the commonly reported ESS in this case. Analogous plots (not

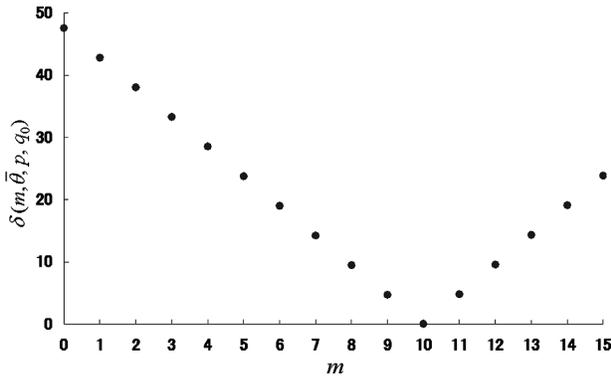


Figure 1. Plot of $\delta(m, \bar{\theta}, p, q_0)$ against m for the beta/binomial model with $\bar{\theta} = (\tilde{\alpha}, \tilde{\beta}) = (3, 7)$.

shown) in all other cases examined below are very similar in appearance to Figure 1.

Example 2. Gamma/exponential model. $\delta(m, \bar{\theta}, p, q_0) = (\tilde{\alpha} - 1)\bar{\theta}^{-2} - (\tilde{\alpha}/c + m - 1)\bar{\theta}^{-2}$, where $\bar{\theta} = \tilde{\alpha}/\tilde{\beta}$, and the ESS is found analytically to be $\tilde{\alpha}$, as desired.

Example 3. Univariate normal with known variance. For $Y | \theta \sim \text{N}(\theta, \sigma^2)$ with σ^2 known and prior $\theta | \tilde{\theta} \sim \text{N}(\tilde{\mu}, \tilde{\sigma}^2)$, so that $\bar{\theta} = (\tilde{\mu}, \tilde{\sigma}^2)$, one may compute analytically $D_p(\theta) = -\partial^2 \log\{p(\theta | \tilde{\theta})\} / \partial \theta^2 = 1/\tilde{\sigma}^2$, and similarly, $D_q(m, \theta) = m/\sigma^2$. Thus, $\delta(m, \bar{\theta}, p, q_0) = |1/\tilde{\sigma}^2 - m/\sigma^2|$, so the ESS = $\sigma^2/\tilde{\sigma}^2$, the ratio of the known variance in the likelihood to the prior variance of θ . In applying this model to a clinical trial setting where θ is the difference between two treatment effects, Spiegelhalter et al. (1994, Section 3.1.2) propose assuming that $\tilde{\sigma}^2 = \sigma^2/n_0$ to obtain a prior that "... is equivalent to a normalized likelihood arising from a (hypothetical) trial of n_0 patients with an observed value $\tilde{\mu}$ of the treatment difference statistic." Thus, in this case, the two methods for defining prior ESS agree.

Example 4. Inverse χ^2 /normal model. We find analytically that $D_p(\theta) = -(\sigma^2)^{-2}(\tilde{\nu} + 2)/2 + (\sigma^2)^{-3}\tilde{\nu}\tilde{\sigma}^2$, whereas $\int D_q(m, \bar{\theta}, \mathbf{Y}_m) f_m(\mathbf{Y}_m) d\mathbf{Y}_m$ is obtained by simulation. As explained in Section 3, the adjustment factor $\{D_q(4, \bar{\theta}) - D_q(0, \bar{\theta})\}$ is added to $D_q(m, \bar{\theta})$. For $\bar{\theta} = (\tilde{\nu}, \tilde{\sigma}^2) = (20, 1)$, ESS = 20 = $\tilde{\nu}$, as desired.

Example 5. Dirichlet/multinomial model. From Table 1, denote $\tilde{\alpha} = (\tilde{\alpha}_1, \dots, \tilde{\alpha}_J)$, $\theta = (\theta_1, \dots, \theta_J)$ and $\mathbf{S} = (S_1, \dots, S_J)$ with $S_j = \sum_{i=1}^m Y_{ji}$. Compute $D_{q,j}(m, \theta)$ analytically, as with the beta-binomial. For $d = 3$ and $\bar{\theta} = (10, 15, 25)$, ESS = 50 = $\sum \tilde{\alpha}_j$, as desired.

Example 6. Power priors. Ibrahim and Chen (2000) propose a class of "power priors" based on an initial prior $p_0(\theta | c_0)$, a likelihood $L(\theta | D_0)$ of historical data D_0 , and a scalar prior parameter a_0 . The power prior is $p(\theta | D_0, a_0) \propto L(\theta | D_0)^{a_0} p_0(\theta | c_0)$, so that a_0 weights the historical data relative to the data that will be obtained in the future. To see how one would compute the ESS of a power prior, consider the beta/binomial model with a beta(1,1) initial prior and D_0 consisting of three successes in 10 historical trials. The power prior is $p(\theta | D_0, a_0) = p(\theta | (3, 10), a_0) \propto \{\theta^3(1 - \theta^7)\}^{a_0} \theta(1 - \theta)$, and it follows easily (case 1) that ESS = $a_0 10 + 2$. More generally, the ESS of $p(\theta | D_0, a_0)$ is $a_0 \times \{L(\theta | D_0)\} + \text{ESS}\{p_0(\theta | c_0)\}$, the weighting parameter times the ESS of the historical data likelihood treated as a function of θ plus the ESS of the initial prior.

Hodges and Sargent (2001) derive a formula for the effective degrees of freedom (EDF) of a richly parameterized model, and illustrate this for a balanced one-way normal linear random effects model for Nn observations $\{Y_{ij}, i = 1, \dots, N, j = 1, \dots, n\}$, given by the likelihood $Y_{i1}, \dots, Y_{in} | \theta_i, \sigma^2 \sim \text{i.i.d. } \text{N}(\theta_i, \sigma^2)$ for each i , and prior $\theta_1, \dots, \theta_N | \tilde{\mu}, \tilde{\sigma}^2 \sim \text{i.i.d. } \text{N}(\tilde{\mu}, \tilde{\sigma}^2)$. They show that the EDF for this model is $\rho = (nN + \phi)/(n + \phi)$, where $\phi = \sigma^2/\tilde{\sigma}^2$, the ratio of the residual variance and the prior variance. Recall from Example 3 that ϕ is the ESS of the simple normal model with known

Table 3

Comparison of ESSs computed using the proposed method and the crude method that matches first and second moments to a beta, for the logistic regression model, $\pi(X_i, \theta) = \text{Pr}(Y_i = 1 | X_i, \theta) = \exp(\mu + \beta X_i) / \{1 + \exp(\mu + \beta X_i)\}$, where the priors are $\mu \sim N(\tilde{\mu}_\mu, \tilde{\sigma}_\mu^2)$ with $\tilde{\mu}_\mu = -0.1313$ and $\beta \sim N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$ with $\tilde{\mu}_\beta = 2.3980$

$\tilde{\sigma}_\mu^2 = \tilde{\sigma}_\beta^2$	Proposed method			Crude method						
	m	m_μ	m_β	\bar{m}^*	$m_{X_{(1)}}$	$m_{X_{(2)}}$	$m_{X_{(3)}}$	$m_{X_{(4)}}$	$m_{X_{(5)}}$	$m_{X_{(6)}}$
0.5 ²	37.1	22.7	101.3	18.2	23.5	18.2	17.0	16.6	16.8	17.3
1.0 ²	9.3	5.7	25.3	4.5	4.1	4.7	4.8	4.6	4.4	4.2
2.0 ²	2.3	1.4	6.3	1.3	1.0	1.4	1.5	1.5	1.3	1.2
3.0 ²	1.0	0.6	2.8	0.7	0.5	0.8	0.8	0.8	0.7	0.7
5.0 ²	0.4	0.2	1.0	0.4	0.3	0.4	0.4	0.4	0.4	0.3

$$\bar{m}^* = 6^{-1} \sum_{z=1}^6 m_{X_{(z)}}.$$

variance. In the limiting case with $\phi \rightarrow \infty$, that is, all θ_i are equal, $\theta_i = \mu$, we find $\rho = 1$. In other words, for large ESS and essentially only one group, Hodges and Sargent report $\rho \approx 1$. At the other extreme, for $\phi \rightarrow 0$, that is, for small ESS and θ_i 's very different from each other, they report $\rho \approx N$. However, such comparisons should not be overinterpreted. EDF and ESS are quite different summaries. Formally, the EDF is a function of the sample size n . In contrast, ESS is not a function of n . Rather, it reports an equivalent sample size for the given model.

Using an information-theoretic argument, Spiegelhalter et al. (2002) also derive a measure for the effective number of parameters in complex models, such as generalized linear (mixed effects) models, p_D , defined as the difference between the posterior mean of the deviance and the deviance evaluated at the posterior means of the parameters of interest. But, similar to the EDF ρ , the nature of p_D is different from the proposed ESS. Formally, p_D is a function of the data, whereas the ESS is not.

6. Application to Some Nonstandard Cases

The following examples show how ESS values may be computed in settings where no commonly agreed-upon ESS exists, using the numerical approximations described earlier to obtain $\delta(m, \bar{\theta}, p, q_0)$.

Example 7. Logistic regression. Thall and Lee (2003) use a logistic regression model to determine a maximum tolerable dose in a phase I clinical trial. Each patient receives one of six doses 100, 200, 300, 400, 500, 600 mg/m², denoted by x_1, \dots, x_6 , with standardized doses $X_{(z)} = \log(x_z) - 6^{-1} \sum_{i=1}^6 \log(x_i)$. The outcome variable is the indicator $Y_i = 1$ if a patient i suffers toxicity, 0 if not. A logistic model $\pi(X_i, \theta) = \text{Pr}(Y_i = 1 | X_i, \theta) = \text{logit}^{-1}\{\eta(X_i, \theta)\}$ with $\eta(X_i, \theta) = \mu + \beta X_i$ is assumed, where $\text{logit}^{-1}(x) = e^x / (1 + e^x)$. Hence $d = 2$, $\theta = (\theta_1, \theta_2) = (\mu, \beta)$, and the likelihood for m patients is

$$f_m(\mathbf{Y}_m | \mathbf{X}_m, \theta) = \prod_{i=1}^m \pi(X_i, \theta)^{Y_i} \{1 - \pi(X_i, \theta)\}^{1-Y_i}.$$

Thall and Lee (2003) obtained independent normal priors for μ and β , based on elicited mean $\pi(X, \theta)$ for doses $x_2 = 200$ and $x_5 = 500$, and setting $\tilde{\sigma}_\mu = \tilde{\sigma}_\beta = 2$

based on preliminary sensitivity analyses, which yielded $N(\tilde{\mu}_\mu, \tilde{\sigma}_\mu^2) = N(-0.1313, 2^2)$ and $N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2) = N(2.3980, 2^2)$. For this application, Algorithms 1 and 2 may be applied to compute one ESS of $p(\theta | \tilde{\theta})$ and two ESSs m_μ and m_β of the priors for μ and β , as follows. For step 1, specify $q_0(\theta | \tilde{\theta}_0) = N(\tilde{\mu}_\mu, c\tilde{\sigma}_\mu^2)N(\tilde{\mu}_\beta, c\tilde{\sigma}_\beta^2)$, with $c = 10,000$. Next, compute $D_{p,1}(\theta) = (\tilde{\sigma}_\mu^2)^{-1}$, $D_{p,2}(\theta) = (\tilde{\sigma}_\beta^2)^{-1}$, $D_{q,1}(m, \theta, \mathbf{Y}_m, \mathbf{X}_m) = \sum_{i=1}^m \pi(X_i, \theta)\{1 - \pi(X_i, \theta)\}$, and $D_{q,2}(m, \theta, \mathbf{Y}_m, \mathbf{X}_m) = \sum_{i=1}^m X_i^2 \pi(X_i, \theta)\{1 - \pi(X_i, \theta)\}$. Because $D_{q,1}(m, \theta, \mathbf{Y}_m, \mathbf{X}_m)$ and $D_{q,2}(m, \theta, \mathbf{Y}_m, \mathbf{X}_m)$ depend on \mathbf{X}_m but not on \mathbf{Y}_m , this simplifies the simulation method given in Section 3.2. We assume a uniform distribution on the six doses for the probability model $g(X_i | \xi)$. Draw $X_1^{(t)}, \dots, X_M^{(t)}$ independently from $\{X_{(1)}, \dots, X_{(6)}\}$ with probability $1/6$ each, for $t = 1, \dots, 100,000$. Then, using the plug-in vector $\bar{\theta} = (\bar{\mu}, \bar{\beta}) = (\tilde{\mu}_\mu, \tilde{\mu}_\beta)$, compute $\delta(m, \bar{\theta}, p, q_0)$ for each $m = 0, \dots, M$, $\delta_1(m_\mu, \bar{\theta}, p, q_0)$ for each $m_\mu = 0, \dots, M_1$, and $\delta_2(m_\beta, \bar{\theta}, p, q_0)$ for each $m_\beta = 0, \dots, M_2$. As shown in Table 3, $m = 2.3$, $m_\mu = 1.4$, and $m_\beta = 6.3$.

Because the standardized doses X_i were defined to be centered at 0, one may interpret m_μ as the ESS for the prior on the average effect, and m_β as the ESS for the dose effect. The prior indicates greater knowledge about the effects of the doses than about the average response. Because $m = 2.3$, after enrolling 3 patients, the information from the likelihood starts to dominate the prior, as desired.

As a sensitivity analysis, Table 3 summarizes corresponding results for $\tilde{\sigma}_\mu^2 = \tilde{\sigma}_\beta^2 = 0.5^2, 1.0^2, 3.0^2$, and 5.0^2 . As a basis for comparison, we also include the ESS at each dose obtained by the crude method of equating the mean and variance of $\pi(X_{(z)}, \theta)$ at each dose to the corresponding values for a beta, $E(\theta) = \tilde{\alpha} / (\tilde{\alpha} + \tilde{\beta})$, and $\text{Var}(\theta) = \{E(\theta)(1 - E(\theta))\} / (\tilde{\alpha} + \tilde{\beta} + 1)$, and solving for $\tilde{\alpha} + \tilde{\beta}$. We denote by \bar{m} the average of the ESSs $m_{X_{(1)}}, \dots, m_{X_{(6)}}$ at the six doses, obtained in this way. The results indicate that the crude method provides smaller estimates of the ESS for $\tilde{\sigma}^2 < 5.0^2$.

It also is useful to examine how the ESS in this example would vary with a_0 if one wished to reweight the prior by replacing it with a power prior $\{p(\theta | \tilde{\theta})\}^{a_0}$. Identifying $p(\theta | \tilde{\theta})$ with $L(\theta | D_0)$ in the set-up of Ibrahim and Chen (2000), and considering the additional ESS of an initial prior to be negligible, the ESS may be computed by applying Algorithms 1 and 2 and setting the ϵ -information prior to be $\{q_0(\theta | \tilde{\theta}_0)\}^{a_0}$.

Table 4
 ESSs for power priors $\{p(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}})\}^{a_0}$ based on the prior $\{p(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}})\}$ in the logistic regression example, using hyperparameter values $\tilde{\mu}_\mu = -0.1313$ and $\tilde{\mu}_\beta = 2.3980$, as in Table 3, with $\tilde{\sigma}_\mu^2 = \tilde{\sigma}_\beta^2 = 4$

a_0	m	m_μ	m_β
0.5	1.2	0.7	3.2
1	2.3	1.4	6.3
2	4.6	2.8	12.6
4	9.3	5.7	25.3

This yields the values summarized in Table 4. These values illustrate, as in Example 6 given earlier, that the power a_0 acts essentially as a multiplier in the ESS domain, aside from the additive ESS of an initial prior.

Example 8. Two-agent dose-response model. The next example is the one described earlier in Section 2—a design to find acceptable dose combinations of two cytotoxic agents used together in a phase I trial. Recall the definition of $\pi(X, \boldsymbol{\theta})$ given in equation (1). The likelihood for m patients with toxicity indicators $\mathbf{Y}_m = (Y_1, \dots, Y_m)$ and dose pairs $\mathbf{X}_m = (X_1, \dots, X_m)$ is

$$f(\mathbf{Y}_m | \mathbf{X}_m, \boldsymbol{\theta}) = \prod_{i=1}^m \pi(X_i, \boldsymbol{\theta})^{Y_i} \{1 - \pi(X_i, \boldsymbol{\theta})\}^{1-Y_i}. \quad (5)$$

Based on (5) and the gamma priors given in Section 2, for this case, Algorithm 1 is used to compute one ESS, m , of $p(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}})$. The three ESSs m_1 , m_2 , and m_3 for $\boldsymbol{\theta}_1$, $\boldsymbol{\theta}_2$, and $\boldsymbol{\theta}_3$ can be computed using Algorithm 2. In step 1, with $c = 10,000$, $q_0(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}}_0) = \prod_{k=1}^3 \text{Ga}(\tilde{a}_{k,1}/c, \tilde{a}_{k,2}/c) \times \text{Ga}(\tilde{b}_{k,1}/c, \tilde{b}_{k,2}/c)$. In step 2, we computed $D_{p,1}(\boldsymbol{\theta}) = (\tilde{a}_{1,1} - 1)\alpha_1^{-2}, \dots, D_{p,6}(\boldsymbol{\theta}) = (\tilde{b}_{3,1} - 1)\beta_3^{-2}$ analytically. The numerical methods given in Section 3 give $\delta_k(m_k, \tilde{\boldsymbol{\theta}}, p, q_0)$ for $k = 1, 2, 3$, yielding the values $m = 1.5$, $m_1 = 547.3$, $m_2 = 756.8$, and $m_3 = 0.01$, as reported earlier.

Example 9: Linear regression. The last example is a linear regression model used to analyze a small data set $(Y_1, X_1), \dots, (Y_{10}, X_{10})$, where Y_i is December rainfall and X_i is November rainfall for 10 consecutive years $i = 1, \dots, 10$ (Congdon, 2001). The sampling model is $Y_i | X_i, \boldsymbol{\theta} \sim N(\mu_i, 1/\tau)$ with $\mu_i = \alpha + \beta(X_i - \bar{X})$ and τ denoting the precision where \bar{X} is the sample average of the original predictor, so $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3) = (\alpha, \beta, \tau)$. Let $N(x; m, s)$ indicate that the random variable x is normally distributed with moments (m, s) . In Congdon (2001), an independent prior $p(\boldsymbol{\theta}) = p_1(\theta_1, \theta_2 | \tilde{\boldsymbol{\theta}}_1, \tilde{\boldsymbol{\theta}}_2) \cdot p_2(\theta_3 | \tilde{\boldsymbol{\theta}}_3)$ is assumed, with $p_1(\theta_1, \theta_2) = N(\theta_1; \tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2) \cdot N(\theta_2; \tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$ and $p_2 = \text{Ga}(\tilde{a}, \tilde{b})$. Congdon (2001) uses $\tilde{\mu}_\alpha = \tilde{\mu}_\beta = 0$, $\tilde{\sigma}_\alpha^2 = \tilde{\sigma}_\beta^2 = 1000$, $\tilde{a} = \tilde{b} = 0.001$. Algorithm 2 was used to compute two ESSs: m_1 for $p_1(\theta_1, \theta_2 | \tilde{\boldsymbol{\theta}}_1, \tilde{\boldsymbol{\theta}}_2)$ and m_2 of $p_2(\theta_3 | \tilde{\boldsymbol{\theta}}_3)$. The plug-in vector is $\tilde{\boldsymbol{\theta}} = \mathbb{E}_p(\boldsymbol{\theta}) = (\tilde{\mu}_\alpha, \tilde{\mu}_\beta, \tilde{a}/\tilde{b})$. In step 1, specify $q_0(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}}_0) = q_{0,1}(\theta_1 | \tilde{\boldsymbol{\theta}}_{0,1})q_{0,1}(\theta_2 | \tilde{\boldsymbol{\theta}}_{0,2})q_{0,2}(\theta_3 | \tilde{\boldsymbol{\theta}}_{0,3}) = N(\tilde{\mu}_\alpha, c\tilde{\sigma}_\alpha^2)N(\tilde{\mu}_\beta, c\tilde{\sigma}_\beta^2) \times \text{Ga}(\tilde{a}/c, \tilde{b}/c)$, with $c = 10,000$. In step 2, compute analytically $D_{p,1}(\boldsymbol{\theta}) = (\tilde{\sigma}_\alpha^2)^{-1}$, $D_{p,2}(\boldsymbol{\theta}) = (\tilde{\sigma}_\beta^2)^{-1}$, $D_{p,3}(\boldsymbol{\theta}) = (\tilde{a} - 1)\tau^{-2}$,

$D_{q,1}(m_1, \boldsymbol{\theta}, \mathbf{Y}_{m_1}, \mathbf{X}_{m_1}) = (c\tilde{\sigma}_\alpha^2)^{-1} + \tau m_1$, and $D_{q,3}(m_2, \boldsymbol{\theta}, \mathbf{Y}_{m_2}, \mathbf{X}_{m_2}) = (\tilde{a}/c - 1)\tau^{-2} + m_2 \tau^{-2}/2$. For this case, only $D_{q,2}(m_1, \boldsymbol{\theta}, \mathbf{Y}_{m_1}, \mathbf{X}_{m_1}) = (c\tilde{\sigma}_\beta^2)^{-1} + \tau \sum_{i=1}^{10} X_i^2$ depends on \mathbf{X} . Following the methods in Section 3, we simulated $X_1^{(t)}, \dots, X_{M_1}^{(t)} \sim$ i.i.d. $N(0, 1)$ for $t = 1, \dots, 100,000$ to obtain $m_1 = 0.001$ and $m_2 = 0.002$. We interpret the reported ESSs as evidence of very vague priors. As a sensitivity analysis, we also computed the ESSs of two alternative priors $p'(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}}) = N(0, 100) N(0, 10) \text{Ga}(1, 1)$ and $p''(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}}) = N(0, 1) N(0, 1) \text{Ga}(2, 2)$, which gave $m_1 = 0.06$ and $m_2 = 2.0$ for $p'(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}})$, and $m_1 = 1.0$ and $m_2 = 4.0$ for $p''(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}})$.

7. Discussion

The methods proposed in this article are useful in Bayesian analysis, particularly in settings with elicited priors or where the data consist of a relatively small number of observations. By computing ESSs, one may avoid the use of an overly informative prior in the sense that inference is dominated by the prior rather than the data. As noted in our guidelines for application, other uses of ESS values include interpreting or reviewing others' Bayesian analyses or designs, using the ESS values themselves to perform sensitivity analyses in the prior's informativeness, and calibrating the parameters of outcome-adaptive Bayesian designs.

Extension of our methods to accommodate hierarchical models is not straightforward. This is a potentially important area for future research, because it would be useful to compute ESS values in such settings. Other potential applications involving more complicated problems include mixture priors synthesizing multiple component priors, or the class of ε -contaminated priors, where ε reflects the amount of uncertainty in the prior information (Greenhouse and Wasserman, 1995).

8. Supplementary Materials

The R program to compute ESS values and instructions for using this program are available under the Paper Information link at the *Biometrics* website <http://www.biometrics.tibs.org>.

ACKNOWLEDGEMENTS

Satoshi Morita's work was supported in part by grant H16-TRANS-003 from the Ministry of Health, Labour, and Welfare in Japan. Peter Thall's work was partially supported by NCI grant RO1 CA 83932. Peter Müller's work was partially supported by NCI grant R01 CA 075981. We thank the associate editor and the referee for their thoughtful and constructive comments and suggestions.

REFERENCES

Berger, J. and Pericchi, L. (2001). Objective Bayesian methods for model selection: Introduction and comparison (with discussion). In *Model Selection*, P. Lahiri (ed), Institute of Mathematical Statistics Lecture Notes, Monograph Series, vol. 38, Bethesda, Maryland.

Chaloner, K. and Rhame, F. S. (2001). Quantifying and documenting prior beliefs in clinical trials. *Statistics in Medicine* **20**, 581–600.

- Congdon, P. (2001). *Bayesian Statistical Modelling*. Chichester, U.K.: John Wiley and Sons.
- Garthwaite, P. H., Kadane, J. B., and O'Hagan, A. (2005). Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association* **100**, 680–701.
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2004). *Bayesian Data Analysis*, 2nd edition. New York: Chapman and Hall/CRC.
- Greenhouse, J. and Wasserman, L. (1995). Robust Bayesian methods for monitoring clinical trials. *Statistics in Medicine* **14**, 1379–1391.
- Hodges J. S. and Sargent, D. J. (2001). Counting degrees of freedom in hierarchical and other richly-parameterized models. *Biometrika* **88**, 367–379.
- Ibrahim, J. G. and Chen, M.-H. (2000). Power prior distributions for regression models. *Statistical Science* **15**, 46–60.
- Spiegelhalter, D. J., Freedman, L. S., and Parmar, M. K. B. (1994). Bayesian approaches to randomized trials. *Journal of the Royal Statistical Society, Series A* **157**, 357–416.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002). Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society, Series B* **64**, 583–639.
- Thall, P. F. and Lee, S. J. (2003). Practical model-based dose-finding in phase I clinical trials: Methods based on toxicity. *International Journal of Gynecological Cancer* **13**, 251–261.
- Thall, P. F., Millikan, R. E., Mueller, P., and Lee, S. J. (2003). Dose-finding with two agents in phase I oncology trials. *Biometrics* **59**, 487–496.

Received April 2006. Revised May 2007.

Accepted May 2007.