

Power-Expected-Posterior Priors as Mixtures of g -Priors in Normal Linear Models

Dimitris Fouskakis* and Ioannis Ntzoufras†

Abstract. One of the main approaches used to construct prior distributions for objective Bayes methods is the concept of random imaginary observations. Under this setup, the expected-posterior prior (EPP) offers several advantages, among which it has a nice and simple interpretation and provides an effective way to establish compatibility of priors among models. In this paper, we study the power-expected-posterior prior as a generalization to the EPP in objective Bayesian model selection under normal linear models. We prove that it can be represented as a mixture of g -prior, like a wide range of prior distributions under normal linear models, and thus posterior distributions and Bayes factors are derived in closed form, keeping therefore its computational tractability. Following this result, we can naturally prove that desiderata (criteria for objective Bayesian model comparison) hold for the PEP prior. Comparisons with other mixtures of g -prior are made and results are presented in simulated and real-life datasets.

Keywords: Bayesian model comparison, expected-posterior priors, imaginary training samples, mixtures of g -priors, objective priors.

1 Introduction

Let $\mathbf{y} = (y_1, \dots, y_n)^T$ denote some available observations. Under the objective Bayesian perspective, suppose we wish to compare the following two models (or hypotheses):

$$\begin{aligned} \text{model } M_0 &: f(\mathbf{y}|\boldsymbol{\theta}_0, M_0), \quad \boldsymbol{\theta}_0 \in \Theta_0, \\ \text{model } M_1 &: f(\mathbf{y}|\boldsymbol{\theta}_1, M_1), \quad \boldsymbol{\theta}_1 \in \Theta_1, \end{aligned} \quad (1)$$

where $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$ are unknown, model specific, parameters. Let further suppose that M_0 is nested in M_1 . By $\pi_\ell^N(\boldsymbol{\theta}_\ell)$, for $\ell \in \{0, 1\}$, we denote the baseline prior of $\boldsymbol{\theta}_\ell$ under model M_ℓ . Here, as a baseline prior we consider any prior that will express low information, for example the reference prior; see Berger et al. (2009). These reference priors are typically improper, resulting in a Bayes factor when comparing M_0 to M_1 which typically cannot be determined due to the unknown normalizing constants of these improper priors.

In order to specify these unknown normalizing constants, Pérez and Berger (2002) developed priors through utilization of the device of “imaginary training samples”. If we denote by \mathbf{y}^* the imaginary training sample, of size n^* , they defined the expected-

*Department of Mathematics, National Technical University of Athens, Greece, fouskakis@math.ntua.gr

†Department of Statistics, Athens University of Economics and Business, Greece, ntzoufras@aueb.gr

posterior prior (EPP) for the parameter vector $\boldsymbol{\theta}_\ell$, of model M_ℓ , as

$$\pi_\ell^{EPP}(\boldsymbol{\theta}_\ell) = \int \pi_\ell^N(\boldsymbol{\theta}_\ell|\mathbf{y}^*) m^*(\mathbf{y}^*) d\mathbf{y}^*, \quad (2)$$

where $\pi_\ell^N(\boldsymbol{\theta}_\ell|\mathbf{y}^*)$ is the posterior of $\boldsymbol{\theta}_\ell$ for model M_ℓ using the baseline prior $\pi_\ell^N(\boldsymbol{\theta}_\ell)$ and data \mathbf{y}^* and $m^*(\mathbf{y}^*)$ is a reference predictive distribution generating sets of imaginary data. A usual choice of m^* is $m^*(\mathbf{y}^*) = m_0^N(\mathbf{y}^*) \equiv f(\mathbf{y}^*|M_0)$, i.e. the marginal likelihood, evaluated at \mathbf{y}^* , for the simplest model M_0 under the baseline prior $\pi_0^N(\boldsymbol{\theta}_0)$. Then model M_0 is called the reference model. EPP offers several advantages, among which it has a nice interpretation and also provides an effective way to establish compatibility of priors among models (Consonni and Veronese, 2008). Furthermore this selection makes the EPP approach essentially equivalent to the arithmetic intrinsic Bayes factor approach of Berger and Pericchi (1996) since it is straightforward to prove that the EPP can be written equivalently as

$$\pi_\ell^{EPP}(\boldsymbol{\theta}_\ell) = \pi_\ell^N(\boldsymbol{\theta}_\ell) E_{\mathbf{y}^*}^{M_\ell} \left[\frac{m_0^N(\mathbf{y}^*)}{m_\ell^N(\mathbf{y}^*)} \right] = \pi_\ell^N(\boldsymbol{\theta}_\ell) \int \frac{m_0^N(\mathbf{y}^*)}{m_\ell^N(\mathbf{y}^*)} f(\mathbf{y}^*|\boldsymbol{\theta}_\ell, M_\ell) d\mathbf{y}^*.$$

When information on covariates is also available, under the EPP methodology, imaginary design matrices X^* with n^* rows should also be introduced. The selection of a *minimal* training sample size n^* has been proposed by Berger and Pericchi (2004), to make the information content of the prior as small as possible, and this is an appealing idea. Then X^* can be extracted from the original design matrix X , by randomly selecting n^* from the n rows.

To diminish the effect of training samples, Fouskakis et al. (2015), generalized the EPP approach, by introducing the *power-expected-posterior* (PEP) priors, combining ideas from the power-prior approach of Ibrahim and Chen (2000) and the unit-information-prior approach of Kass and Wasserman (1995). As a first step, the likelihoods involved in the EPP formula are raised to the power $1/\delta$ ($\delta \geq 1$) and then are density-normalized; for a discussion on the density-normalization and different versions of the PEP priors when this normalization leads to no standard forms see Fouskakis et al. (2018). This power parameter δ could be then set equal to the size of the training sample n^* , to represent information equal to one data point. In Fouskakis et al. (2015) the authors further set $n^* = n$; this choice gives rise to significant advantages, for example when covariates are available it results in the automatic choice $X^* = X$ and therefore the selection of a training sample and its effects on the posterior model comparison is avoided, while still holding the prior information content at one data point.

Specifically, for the model selection problem (1), the PEP prior is defined as

$$\pi_\ell^{PEP}(\boldsymbol{\theta}_\ell|\delta) \equiv \pi_\ell^{PEP}(\boldsymbol{\theta}_\ell) = \int \pi_\ell^N(\boldsymbol{\theta}_\ell|\mathbf{y}^*, \delta) m^*(\mathbf{y}^*|\delta) d\mathbf{y}^*, \quad (3)$$

with

$$\pi_\ell^N(\boldsymbol{\theta}_\ell|\mathbf{y}^*, \delta) \propto f(\mathbf{y}^*|\boldsymbol{\theta}_\ell, \delta, M_\ell) \pi_\ell^N(\boldsymbol{\theta}_\ell), \quad (4)$$

$$f(\mathbf{y}^*|\boldsymbol{\theta}_\ell, \delta, M_\ell) = \frac{f(\mathbf{y}^*|\boldsymbol{\theta}_\ell, M_\ell)^{1/\delta}}{\int f(\mathbf{y}^*|\boldsymbol{\theta}_\ell, M_\ell)^{1/\delta} d\mathbf{y}^*}. \tag{5}$$

As before we choose

$$m^*(\mathbf{y}^*|\delta) = m_0^N(\mathbf{y}^*|\delta) = \int f(\mathbf{y}^*|\boldsymbol{\theta}_0, \delta, M_0)\pi_0^N(\boldsymbol{\theta}_0)d\boldsymbol{\theta}_0, \tag{6}$$

where $f(\mathbf{y}^*|\boldsymbol{\theta}_0, \delta, M_0)$ is given by (5) for $\ell = 0$ (i.e. the null/reference model).

In this work we show, using sufficient statistics (see Fouskakis, 2019 for the equivalent definitions of EPP and PEP prior using sufficient statistics), that the PEP prior (and therefore also the EPP), in normal linear model comparison, can be expressed as a mixture of g -priors, centred around null models. This has the advantage that posterior distributions, as well as, marginal likelihoods are available in closed form and desiderata (criteria for objective Bayesian model comparison), see Bayarri et al. (2012), hold. We compare the PEP prior with other scale normal mixture priors, we present prior summaries for model parameters and we derive posterior distributions as well as marginal likelihoods. Emphasis is given to the Bayesian inference of the shrinkage parameter which is also involved in Bayesian model averaging estimation. Finally we present results from a simulation study, as well as from a real-life dataset.

2 PEP Prior as a Mixture of Normal Distribution in Normal Linear Model Comparison

Let $\mathbf{y} = (y_1, \dots, y_n)^T$ be a random sample. We would like to compare the nested models:

$$H_0 : \text{model } M_0 : \text{Normal}(\mathbf{y}|X_0\boldsymbol{\beta}_0, \sigma_0^2), \pi_0^N(\boldsymbol{\beta}_0, \sigma_0) = c_0\pi_0^U(\boldsymbol{\beta}_0, \sigma_0) = c_0\sigma_0^{-(1+d_0)}$$

$$\text{vs. } H_1 : \text{model } M_1 : \text{Normal}(\mathbf{y}|X_1\boldsymbol{\beta}_1, \sigma_1^2), \pi_1^N(\boldsymbol{\beta}_1, \sigma_1) = c_1\pi_1^U(\boldsymbol{\beta}_1, \sigma_1) = c_1\sigma_1^{-(1+d_1)}$$

where c_0 and c_1 are the unknown normalizing constants of $\pi_0^U(\boldsymbol{\beta}_0, \sigma_0)$ and $\pi_1^U(\boldsymbol{\beta}_1, \sigma_1)$ respectively, X_0 is an $(n \times k_0)$ design matrix under model M_0 , X_1 is an $(n \times k_1)$ design matrix under model M_1 , $k_0 < k_1 < n$ and M_0 is nested in M_1 . Furthermore let $\boldsymbol{\beta}_1 = (\boldsymbol{\beta}_0^T, \boldsymbol{\beta}_{e_1}^T)^T$, $X_1 = [X_0|X_{e_1}]$, $P_0 = X_0(X_0^T X_0)^{-1} X_0^T$ and $P_1 = X_1(X_1^T X_1)^{-1} X_1^T$. All matrices are assumed to be of full rank. Usual choices for d_0 and d_1 are $d_0 = d_1 = 0$ (resulting to the reference prior) or $d_0 = k_0$ and $d_1 = k_1$ (resulting to the dependence Jeffreys prior).

In the above comparison we assume that model M_0 is nested in model M_1 , so that $k_0 < k_1$ and thus we henceforth assume that $\boldsymbol{\beta}_1 = (\boldsymbol{\beta}_0^T, \boldsymbol{\beta}_{e_1}^T)^T$, so that $\boldsymbol{\beta}_0$ is a parameter “common” between the two models, where $\boldsymbol{\beta}_{e_1}$ is model specific. The use of a “common” parameter $\boldsymbol{\beta}_0$ in nested model comparison is often made to justify the employment of the same, potentially improper, prior on $\boldsymbol{\beta}_0$ across models. This usage is becoming standard, see for example Bayarri et al. (2012) and Consonni et al. (2018). It can be justified if, without essential loss of generality, we assume that the model has been parametrized

in an orthogonal fashion, so that $X_0^T X_1 = 0$. In the special case where M_0 is the “null” model, with only the intercept, this assumption can be justified, if we assume, again without loss of generality, that the columns of the design matrix of the full model have been centred on their corresponding means, which makes the covariates orthogonal to the intercept, and gives the intercept an interpretation that is “common” to all models. Regarding the error variance, although it is also standard to be treated as a “common” parameter across models, in this paper we follow the “intrinsic prior methodology” (see for example Moreno and Girón, 2008) and we treat it as a model specific parameter. As we will see later in this Section, this causes no issues about the indeterminacy of Bayes factors due to the “intrinsicification”.

In variable selection problems we are interested in posterior model probabilities, which, nevertheless, are functions of posterior odds (and the corresponding Bayes factors) of pairwise model comparisons. Under the usual case where the reference model M_0 is the null (with only the intercept), we could perform all pairwise comparisons between any model M_1 and the null model M_0 , in order to derive the corresponding posterior model probabilities under the null based approach; see for example in Liang et al. (2008). Moreover, the posterior odds of any two non-nested models is simply the ratio of the null based posterior odds.

In the following theorem, we prove (see Section 1 of the Appendix (Fouskakis and Ntzoufras, 2021)) that the PEP prior, for comparing models M_0 and M_1 , can be represented as a mixture of g -prior, in a similar way as a variety of prior distributions under normal linear models (see Section 3), and thus posterior distributions and Bayes factors are readily available in closed formed expressions (see Sections 4 and 5), keeping, therefore, its computational tractability.

Theorem 1. *The PEP priors (or EPPs for $\delta = 1$) for comparing models M_0 and M_1 are*

$$\{\pi_0^{PEP}(\boldsymbol{\beta}_0, \sigma_0) = \pi_0^N(\boldsymbol{\beta}_0, \sigma_0), \pi_1^{PEP}(\boldsymbol{\beta}_1, \sigma_1)\},$$

with

$$\begin{aligned} \pi_1^{PEP}(\boldsymbol{\beta}_1, \sigma_1) &= \pi_1^{PEP}(\boldsymbol{\beta}_0, \sigma_1) \int_0^1 \pi_1^{PEP}(\boldsymbol{\beta}_{e_1}, t|\sigma_1, \boldsymbol{\beta}_0) dt \\ &\propto \sigma_1^{-(d_0+1)} \int_0^1 f_N(\boldsymbol{\beta}_{e_1}; \mathbf{0}, \frac{\delta\sigma_1^2}{t} V_{e_1}) f_B\left(t; \frac{n^*+d_0-k_1}{2}, \frac{n^*+d_1-d_0-k_1}{2}\right) dt. \end{aligned}$$

□

From Theorem 1 we see that, conditionally on $(\boldsymbol{\beta}_0, \sigma_1)$, the PEP prior is a beta mixture of a multivariate normal prior and overall can be written using the following hierarchical structure

$$\begin{aligned} \boldsymbol{\beta}_{e_1}|t, \sigma_1, \boldsymbol{\beta}_0 &\sim N_{k_1-k_0}\left(\mathbf{0}, \frac{\delta\sigma_1^2}{t} V_{e_1}\right), \quad t|\sigma_1 \sim \text{Beta}\left(\frac{n^*+d_0-k_1}{2}, \frac{n^*+d_1-d_0-k_1}{2}\right), \quad (7) \\ (\boldsymbol{\beta}_0, \sigma_1) &\sim \pi_1^{PEP}(\boldsymbol{\beta}_0, \sigma_1) \propto \sigma_1^{-(d_0+1)}. \end{aligned}$$

The EPP is directly available for $\delta = 1$.

In the above expression, $\pi_1^{PEP}(\beta_{e_1}, t | \sigma_1, \beta_0) = \pi_1^{PEP}(\beta_{e_1} | t, \sigma_1, \beta_0) \pi_1^{PEP}(t)$ is proper and $\pi_1^{PEP}(\beta_0, \sigma_1) \propto \sigma_1^{-(d_0+1)}$; i.e. the reference prior for the baseline model M_0 . Therefore there are no issues about the indeterminacy of the Bayes factor, when comparing model M_0 to M_1 , since after the “intrinsicification” the unknown constants of the imposed priors will be the same for the two competing models; see Section 2 of the Appendix for a detailed explanation.

Under the usual case where the reference model M_0 is the null model (with only the intercept), we have that the prior variance-covariance matrix of the model coefficients is given by $V_{e_1} = (Z_{e_1}^{*T} Z_{e_1}^*)^{-1}$, where $Z_{e_1}^*$ is the matrix of the centred (at the mean) imaginary covariates. In practice, when using the PEP prior with centred covariates and imaginary design matrices equal to actual ones, as in Fouskakis et al. (2015), then the induced approach results in a mixture of g -priors (Liang et al., 2008) with a different hyper-prior on $g = \delta/t$.

	$\pi_1(\beta_1, \sigma_1) \propto \sigma_1^{-(d_0+1)} \int_0^1 f_N(\beta_{e_1}; \mathbf{0}, a \frac{\sigma_1^2}{t} V_{e_1}) f_B(t; b_1, b_2) dt$ $V(\beta_{e_1} \beta_0, \sigma_1) = a E[t^{-1}] V_{e_1} \sigma_1^2$							
	a		b ₁		b ₂		Approx. $V(\beta_{e_1} \beta_0, \sigma_1)$ (for large n^*)	
	EPP	PEP	EPP/PEP	EPP/PEP	$E[t^{-1}]$ EPP/PEP	EPP	PEP	
For any n^*	1	δ	$\frac{n^*+d_0-k_1}{2}$	$\frac{n^*+d_1-d_0-k_1}{2}$	$\frac{2n^*+d_1-2k_1-2}{n^*+d_0-k_1-2}$	$2V_{e_1} \sigma_1^2$	$2\delta V_{e_1} \sigma_1^2$	
Minimal n^* ($n^* = k_1 + 1$)	1	δ	$\frac{d_0+1}{2}$	$\frac{d_1-d_0+1}{2}$	$\frac{d_1}{d_0-1}$	$\frac{d_1}{d_0-1} V_{e_1} \sigma_1^2$	$\delta \frac{d_1}{d_0-1} V_{e_1} \sigma_1^2$	

Table 1: EPPs and PEP priors, under the alternative hypothesis, for the normal linear case, for any n^* and for $n^* = k_1 + 1$ (minimal training sample size).

Table 1 summarizes the EPP and PEP priors, under the alternative hypothesis, for minimal training sample size ($n^* = k_1 + 1$) as well as for any training sample size $n^* \in [k_1 + 1, n]$. Concerning the prior distribution of $\beta_{e_1} | \beta_0, \sigma_1$ (after integrating out the hyper-parameter t), for large n^* , its corresponding variance will be equivalent to the variance of a g -prior with $g = 2$ and $g = 2\delta$ for the EPP and PEP prior, respectively. Clearly, the PEP prior is more dispersed accounting for information equivalent to $n^*/2\delta$ additional data points, while EPP will account for $n^*/2$ additional data-points. When we consider the EPP, with the minimal training sample, that is $n^* = k_1 + 1$, then $V(\beta_{e_1} | \beta_0, \sigma_1)$ is similar to the variance of a g -prior with $g = d_1/(d_0 - 1)$. This means that it can be defined only for choices of $d_0 > 1$. On the other hand, $V(\beta_{e_1} | \beta_0, \sigma_1)$ can be defined without any problem when we consider any training sample of size $n^* > k_1 - d_0 + 2$. Finally, under the PEP prior, the variance of $\beta_{e_1} | \beta_0, \sigma_1$ is further multiplied by δ making larger the spread of the prior and overall the imposed prior less informative. For this reason, the corresponding posterior summaries will be more robust to the specific choices of d_1 and d_0 , especially when $\delta = n^* = n$ and n is large.

3 A General Framework for Scale Normal Mixtures Priors

Generally, a wide range of prior distributions for variable selection in regression can be written with the following form of a normal scale mixture distribution:

$$\pi_1(\boldsymbol{\beta}_{e_1}, \boldsymbol{\beta}_0, \sigma_1) = \sigma_1^{-(d_0+1)} \int_0^{+\infty} f_{N_{k_1-k_0}}(\boldsymbol{\beta}_{e_1}; \mathbf{0}, g\sigma_1^2 \Sigma_{e_1}) \pi_1(g) dg, \quad (8)$$

where $f_{N_d}(\mathbf{y}; \boldsymbol{\mu}, \Sigma)$ denotes the density of the d -dimensional Normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix Σ , evaluated at \mathbf{y} and $\pi_1(g)$ denotes the prior distribution of the parameter g under model M_1 . Under the PEP prior, the hyper-prior $\pi_1(g)$ for g is given by

$$g \sim \text{SGBP}\left(a = \frac{n^*+d_0-k_1}{2}, b = \frac{n^*+d_1-d_0-k_1}{2}, p = 1, q = \delta, s = \delta\right),$$

where *SGBP* stands for the *shifted generalized beta prime* (SGBP) distribution with density

$$f(g; a, b, p, q, s) = \frac{p \left(\frac{q-s}{q}\right)^{bp-1} \left(1 + \left(\frac{q-s}{q}\right)^p\right)^{-a-b}}{qB(a, b)}, \quad g \geq s. \quad (9)$$

The beta prime distribution is a special case of (9) with $p = q = 1$ and $s = 0$. Furthermore, the generalized beta prime distribution is a special case of (9) with $s = 0$.

The prior expectation of g is given by

$$E(g) = q \frac{\Gamma(a - \frac{1}{p}) \Gamma(b + \frac{1}{p})}{\Gamma(a) \Gamma(b)} + s, \quad \text{if } ap > 1.$$

Under the PEP prior, since $p = 1$ and $q = s = \delta$, the density of the hyper-prior for g simplifies to

$$f(g; a, b, \delta) \propto (g - \delta)^{b-1} g^{-a-b}, \quad g \geq \delta, \quad (10)$$

where $a = \frac{n^*+d_0-k_1}{2}$ and $b = \frac{n^*+d_1-d_0-k_1}{2}$ and the expectation is given by

$$E(g) = \delta \times \frac{a + b - 1}{a - 1} = \delta \times \frac{2n^* + d_1 - 2k_1 - 2}{n^* + d_0 - k_1 - 2}.$$

From the above expression, it is evident that the PEP prior implements an indirect averaging approach across all values of $g \geq \delta$. For the recommended setup; see Fouskakis et al. (2015), of $\delta = n^* = n$, this might look quite dramatic at the first sight. But in practice, it is reasonable, under lack of prior information, to consider at most a value of g that will correspond to one unit of information. Moreover, in such cases, the shrinkage w given by $\frac{g}{g+1} = \frac{\delta}{\delta+t}$ (see Section 5) should approach the value of one, such that most of the posterior information comes from the data. In the case where the likelihood mass supports values of g lower than δ , this means that the data do not have enough information in order to estimate sufficiently the model coefficients. An unrestricted prior for g leads to greater shrinkage towards the prior mean of model coefficients $\boldsymbol{\beta}$. The truncation avoids over-shrinkage and the posterior of g will be concentrated at the value of δ ensuring a minimum value of shrinkage towards the prior.

Prior	hyper-prior	Parameters of the SGBP distribution				
		a	b	p	q	s
PEP						
(General)	SGBP	$\frac{n^* - k_1}{2}$	$\frac{n^* - k_1}{2}$	1	δ	δ
(Recommended)	SGBP	$\frac{n - k_1}{2}$	$\frac{n - k_1}{2}$	1	n	n
EPP						
(General)	SGBP	$\frac{n^* - k_1}{2}$	$\frac{n^* - k_1}{2}$	1	1	1
(Recommended)	SGBP	$1/2$	$1/2$	1	1	1
Intrinsic	SGBP	$1/2$	$1/2$	1	$\frac{n}{k_1 + 1}$	$\frac{n}{k_1 + 1}$
Robust						
(General)	SGBP	a_r	1	1	$\frac{b_r + n}{-1}$	$\frac{b_r + n}{-1} - b_r$
(Recommended)	SGBP	$1/2$	1	1	$\frac{\rho_{1,r}}{n+1}$	$\frac{\rho_{1,r}}{n+1} - 1$
MG*						
(General)	Beta'	$a_{mg} + 1$	$b_{mg} + 1$	1	1	0
(Recommended)	Beta'	$1/4$	$\frac{n - q_{mg} - 5}{2} + \frac{3}{4}$	1	1	0
Hyper-g						
(General)	Beta'	$\frac{a_h}{2} - 1$	1	1	1	0
(Recommended)	Beta'	$1/2$	1	1	1	0
Hyper-g/n						
(General)	GBP	$\frac{a_h}{2} - 1$	1	1	n	0
(Recommended)	GBP	$1/2$	1	1	n	0
Benchmark**						
(General)	Beta'	c_b	$c_b \max(n, p^2)$	1	1	0
(Recommended)	Beta'	0.01	$0.01 \max(n, p^2)$	1	1	0

* Maruyama and George (2011) prior but only for the case where $q_{mg} < n - 1$; where q_{mg} is the dimension of an orthogonal matrix which diagonalizes $\mathbf{X}^T \mathbf{X}$.

** Under the Benchmark prior, $p = k - 1$ denotes the total number of regressors.

SGBP: Shifted generalized beta prime distribution.

GBP: Generalized beta prime distribution.

Beta': Beta prime distribution.

Table 2: Mixing distributions of g under different prior setups ($g \geq s$).

3.1 Comparisons

Most of the known priors used for variable selection assume that $\Sigma_{e_1}^{-1} = \mathbf{X}_{e_1}^T (\mathbf{I}_n - \mathbf{P}_0) \mathbf{X}_{e_1}$ in (8). This is also the case for the PEP prior if we consider $\mathbf{X}_{e_1}^* = \mathbf{X}_{e_1}$ as in Fouskakis et al. (2015). Similarly, the benchmark prior (Ley and Steel, 2012), the robust prior (Bayarri et al., 2012), the hyper-g and hyper-g/n priors (Liang et al., 2008) can be written as in (8) with the hyper-prior for g to be as in (9); details are provided in Table 2, under the usual choice of $d_1 = d_0 = 0$ for simplicity reasons. Additionally, the EPP, as shown above (and also in Womack et al., 2014) can be written as in (8), but using imaginary design matrices in $\Sigma_{e_1}^{-1}$, with number of rows usually equal to the minimal training sample ($n^* = k_1 + 1$). Also, the intrinsic prior of Casella and Moreno (2006) can be viewed as an EPP. In their approach, as an approximation, using ideas from the arithmetic intrinsic Bayes factor approach, they used the original design matrix in $\Sigma_{e_1}^{-1}$, with all n rows, using an additional multiplier in the covariance matrix of the normal component in (8) given by $\frac{n}{k_1 + 1}$; see for example Womack et al. (2014). Therefore, this intrinsic prior can be viewed as a PEP prior, with (a) $\mathbf{X}_{e_1}^* = \mathbf{X}_{e_1}$; (b) $n^* = k_1 + 1$ (minimal training sample) and (c) a model dependent power parameter $\delta = \frac{n}{k_1 + 1}$; in the rest of the paper will call this prior intrinsic. Finally the prior by

Maruyama and George (2011) is also closely related, where in the normal component in (8) the rotated coordinates are used, while the Zellner and Siow prior (Zellner and Siow, 1980) is as in (8) with the hyper-prior for g to be an inverted Gamma distribution with parameters $1/2$ and $n/2$.

3.2 Criteria

Bayarri et al. (2012) developed criteria (*desiderata*) to be satisfied by objective prior distributions for Bayesian model choice. Obviously PEP prior satisfies the basic criterion (C1). Furthermore Fouskakis and Ntzoufras (2016) proved that the PEP prior leads to a consistent model selection procedure (criterion C2). Fouskakis and Ntzoufras (2017) showed that the PEP prior satisfies the information consistency criterion (C3). Additionally, as shown here, for $d_0 = 0$, the PEP prior belongs to a more general class of conditional priors

$$\pi_1(\boldsymbol{\beta}_{e_1}, \boldsymbol{\beta}_0, \sigma_1) \propto \sigma_1^{-1-(k_1-k_0)} h_1\left(\frac{\boldsymbol{\beta}_{e_1}}{\sigma_1}\right), \quad (11)$$

where $h_1(\cdot)$ is a proper density with support $\mathbb{R}^{k_1-k_0}$. Bayarri et al. (2012) prove that the group invariance criterion (C7) hold if and only if $\pi_1(\boldsymbol{\beta}_{e_1}, \boldsymbol{\beta}_0, \sigma_1)$ has the form of (11). Additionally, if $h_1(\cdot)$ is symmetric around zero, which is the case under the PEP prior, predictive matching criterion (C5) also holds. When, finally $\mathbf{X}_{e_1}^* = \mathbf{X}_{e_1}$, the conditional scale matrix has the form $\Sigma_{e_1}^{-1} = \mathbf{X}_{e_1}^T (\mathbf{I}_n - \mathbf{P}_0) \mathbf{X}_{e_1}$ and then null predictive matching, dimensional predictive matching and the measurement invariance criterion (C6) hold, according to Bayarri et al. (2012).

3.3 Effective Sample Size

The effective sample size of the prior distribution of a mixture prior of type (8), with $\Sigma_{e_1} = (\mathbf{Z}_{e_1}^T \mathbf{Z}_{e_1})^{-1}$, is given by $ESS = n/E(g)$ or alternatively by the expected $ESS = nE(g^{-1})$; see Section 3 of the Appendix for details.

For the PEP prior we have

$$ESS = \frac{n}{\delta E(t^{-1})} = \frac{n}{\delta} \times \frac{n^* + d_0 - k_1 - 2}{2n^* + d_1 - 2k_1 - 2}.$$

For the default choices, $\delta = n$, $n^* = n$ and $d_1 = d_0 = 0$

$$ESS = \frac{n - k_1 - 2}{2n - 2k_1 - 2} \approx 1/2, \text{ for large } n/k_1.$$

Hence, the PEP prior's effective sample size will be approximately equal to $1/2$ data point, as n goes to infinity and for fixed k_1 . This means that the PEP prior is on average less informative than the unit information prior and this justifies why PEP supports slightly more parsimonious solutions than other competitive methods (see for example Section 7). Similar are the results if we consider the expression for the expected effective sample size.

Table 3 presents the asymptotic values of the effective sample size, under the PEP prior (with $n^* = n$ and $d_1 = d_0 = 0$), for various choices of the parameter δ and for $k_1 = \mathcal{O}(1)$ or $k_1 = \mathcal{O}(n)$ (such that $r = \lim_{n, k_1 \rightarrow \infty} \frac{n}{k_1} > 1$).

Prior parameter	$ESS(r)$	Asymptotic behavior of ESS , as $n \rightarrow \infty$	
		$k_1 = \mathcal{O}(1)$ $r \rightarrow \infty$	$k_1 = \mathcal{O}(n)$ $1 < r < \infty$
$\delta = n$	$N_0(r, k_1)$	$\frac{1}{2}$	$\frac{1}{2}$
$\delta = n - k_1$	$\frac{r}{r-1} N_0(r, k_1)$	$\frac{1}{2}$	$\frac{r}{2(r-2)}$
$\delta = k_1$	$r N_0(r, k_1)$	∞	$\frac{r}{2}$
$\delta = \delta_0$ (fixed)	$\frac{n}{\delta_0} N_0(r, k_1)$	∞	∞
$r = \frac{n}{k_1}$ and $N_0(r, k_1) = \frac{r-1-2/k_1}{2(r-1)-2/k_1}$			

Table 3: Asymptotic values of the effective sample size (ESS), under the PEP prior, for various choices of the parameter δ .

4 Posterior Distributions of Model Parameters

In this Section we the present posterior distributions of model parameters, under the PEP approach. For compatibility with the mixtures of g-prior, we work with the hyperparameter $g = \delta/t$.

4.1 Full Conditional Posteriors and Gibbs Sampling

Under the PEP approach, the full conditional posterior distribution of β_{e_1} is a multivariate normal distribution of the form

$$\beta_{e_1} | g, \sigma_1, \beta_0, \mathbf{y}, M_1 \sim N_{k_{e_1}} \left(W_{e_1} \tilde{\beta}_{e_1}, W_{e_1} (X_{e_1}^T X_{e_1})^{-1} \sigma_1^2 \right), \tag{12}$$

where $k_{e_1} = k_1 - k_0$ and

$$\begin{aligned} \tilde{\beta}_{e_1} &= (X_{e_1}^T X_{e_1})^{-1} X_{e_1}^T (\mathbf{y} - X_0^T \beta_0) = \hat{\beta}_{e_1} - (X_{e_1}^T X_{e_1})^{-1} X_{e_1}^T X_0^T \beta_0, \\ W_{e_1} &= (w X_{e_1}^T X_{e_1} + (1 - w) V_{e_1}^{-1})^{-1} (w X_{e_1}^T X_{e_1}), \\ w &= \frac{g}{g + 1} = \frac{\delta}{\delta + t}; \\ &\text{for } \delta = 1 \Rightarrow \text{EPP}; \quad \text{for } \delta > 1 \Rightarrow \text{PEP}. \end{aligned}$$

The matrix W_{e_1} plays the role of a multivariate shrinkage factor which penalizes each coefficient locally, while w is a global shrinkage factor which affects uniformly the posterior mean and posterior variance-covariance matrix. For example, if $w \rightarrow 0$ all the conditional posterior information is taken from the prior, while for $w \rightarrow 1$ the conditional posterior information will be derived from the data.

Similarly we can obtain the full conditional posterior distributions of β_0 , σ_1^2 and g by

$$\beta_0 | \beta_{e_1}, \sigma_1, g, \mathbf{y}, M_1 \sim N_{k_0} \left(\widehat{\beta}_0 - (X_0^T X_0)^{-1} X_0^T X_{e_1}^T \beta_{e_1}, (X_0^T X_0)^{-1} \sigma_1^2 \right), \quad (13)$$

$$\sigma_1^2 | \beta_0, \beta_{e_1}, g, \mathbf{y}, M_1 \sim IG \left(\frac{n + k_{e_1} + d_0}{2}, \frac{RSS_1 + \beta_{e_1}^T V_{e_1}^{-1} \beta_{e_1}}{2} \right), \quad (14)$$

$$u | \beta_0, \beta_{e_1}, \sigma_1, \mathbf{y}, M_1 \sim CH \left(b, a + \frac{k_1 - k_0}{2} + 2, -\frac{1}{2\delta\sigma_1^2} \beta_{e_1}^T V_{e_1}^{-1} \beta_{e_1} \right),$$

$$g = \delta / (1 - u),$$

where $CH(p, q, s)$ is the confluent hypergeometric distribution with density function

$$f_{CH}(x; p, q, s) \propto x^{p-1} (1-x)^{q-1} e^{-sx} \quad \text{for } 0 \leq x \leq 1.$$

The above conditional distributions can be easily used to implement a full Gibbs sampler in order to obtain any posterior estimates of interest for any specific model. Similarly, it can be used to build a Gibbs based variable selection sampler; see for example in Dellaportas et al. (2002), to obtain estimates of the posterior model weights. We can further simplify the Gibbs sampler by combining the posterior distributions of β_0 and β_{e_1} given in (12) and (13). Finally, the full conditional posterior distribution of $\beta_1^T = (\beta_0^T, \beta_{e_1}^T)$ is given by

$$\beta_1 | \sigma_1, g, \mathbf{y}, M_1 \sim N_{k_1} \left(W_1 \widehat{\beta}_1, W_1 (X_1^T X_1)^{-1} \sigma_1^2 \right),$$

$$W_1 = (w X_1^T X_1 + (1-w) T_1)^{-1} w X_1^T X_1,$$

$$T_1 = \begin{pmatrix} 0_{k_0 \times k_0} & 0_{k_0 \times k_{e_1}} \\ 0_{k_{e_1} \times k_0} & V_{e_1}^{-1} \end{pmatrix},$$

where $0_{\ell_1 \times \ell_2}$ is a matrix of dimension $\ell_1 \times \ell_2$ with zeros, $w = g/(g+1)$ is the shrinkage parameter while $\widehat{\beta}_1$ is the MLE for β_1 of model M_1 given by $\widehat{\beta}_1 = (X_1^T X_1)^{-1} X_1^T \mathbf{y}$.

4.2 Marginal Likelihoods

The marginal likelihood conditionally on a value of g is given by the usual marginal likelihood of the normal inverse gamma prior. Thus

$$f(\mathbf{y} | g, M_1) = C_1 \times (g+1)^{\frac{n+d_0-k_1}{2}} (1+gR_{10})^{-\frac{n+d_0-k_0}{2}}, \quad (15)$$

with $R_{10} = \frac{1-R_1^2}{1-R_0^2}$; where R_ℓ^2 is the coefficient of determination of model M_ℓ ($\ell \in \{0, 1\}$), and C_1 being constant for all models (assuming that the covariates of X_0 are included in all models) given by

$$C_1 = 2^{\frac{d_0}{2}-1} \pi^{\frac{k_0-n}{2}} |X_0^T X_0|^{-1/2} \Gamma \left(\frac{n+d_0-k_0}{2} \right) (1-R_0^2)^{-\frac{n+d_0-k_0}{2}} \|\mathbf{y} - \bar{y} \mathbf{1}_n\|^{-\frac{n+d_0-k_0}{2}}.$$

The full marginal likelihood $f(\mathbf{y}|M_1)$ is given by

$$\begin{aligned}
 f(\mathbf{y}|M_1) &= \frac{C_1}{\delta B\left(\frac{n^*+d_0-k_1}{2}, \frac{n^*+d_1-d_0-k_1}{2}\right)} \\
 &\times \delta^{1-b} \delta^{a+b} \int_{\delta}^{\infty} (1+g)^{\frac{n+d_0-k_1}{2}} (g-\delta)^{b-1} g^{-a-b} (1+gR_{10})^{-\frac{n+d_0-k_0}{2}} dg \\
 &= \frac{C_1}{\delta B\left(\frac{n^*+d_0-k_1}{2}, \frac{n^*+d_1-d_0-k_1}{2}\right)} \times \delta^{1-b} \delta^{a+b} \delta^{-a} (\delta+1)^{\frac{n+d_0-k_1}{2}} [1+\delta R_{10}]^{-\frac{n+d_0-k_0}{2}} \\
 &\times \int_0^1 u^{b-1} (1-u)^{\frac{k_1-k_0}{2}+a-1} \left(1-\frac{u}{\delta+1}\right)^{\frac{n+d_0-k_1}{2}} \left(1-u\frac{1-R_0^2}{1-R_0^2+\delta(1-R_1^2)}\right)^{-\frac{n+d_0-k_0}{2}} du \\
 &= C_1 \times \frac{B\left(\frac{k_1-k_0}{2}+a, b\right)}{B(a, b)} \times (\delta+1)^{\frac{n+d_0-k_1}{2}} (1+\delta R_{10})^{-\frac{n+d_0-k_0}{2}} \times \tilde{F}_1(0) \tag{16}
 \end{aligned}$$

with $\tilde{F}_1(0) = F_1\left(b, \frac{n+d_0-k_0}{2}, -\frac{n+d_0-k_1}{2}, \frac{k_{e_1}}{2}+a+b; \frac{1}{1+\delta R_{10}}, \frac{1}{\delta+1}\right)$;

where in the above $k_{e_1} = k_1 - k_0$, $a = \frac{n^*+d_0-k_1}{2}$, $b = \frac{n^*+d_1-d_0-k_1}{2}$ (reminder from 10) and $F_1(a', b'_1, b'_2, c'; x, y)$ is the hypergeometric function of two variables or Appell hypergeometric function given by

$$F_1(a', b'_1, b'_2, c'; x, y) = \frac{1}{B(a', c'-a')} \int_0^1 t^{a'-1} (1-t)^{c'-a'-1} (1-xt)^{-b'_1} (1-yt)^{-b'_2} dt.$$

Note that the marginal likelihood of the reference model M_0 is $f(\mathbf{y}|M_0) = C_1$.

4.3 Marginal Posterior Distribution of g

The marginal posterior distribution of g , under model M_1 , is given by

$$\pi_1(g|\mathbf{y}) \equiv \pi(g|\mathbf{y}, M_1) = C_2 \times (1+g)^{\frac{n+d_0-k_1}{2}} (1+gR_{10})^{-\frac{n+d_0-k_0}{2}} (g-\delta)^{b-1} g^{-a-b} \tag{17}$$

for $g \geq \delta$ with the normalizing constant C_2 given by

$$C_2 = \frac{\delta^a (\delta+1)^{-\frac{n+d_0-k_1}{2}} (1+\delta R_{10})^{\frac{n+d_0-k_0}{2}}}{B\left(b, \frac{k_{e_1}}{2}+a\right) F_1\left(b, \frac{n+d_0-k_0}{2}, -\frac{n+d_0-k_1}{2}, \frac{k_{e_1}}{2}+a+b; \frac{1}{1+\delta R_{10}}, \frac{1}{\delta+1}\right)}. \tag{18}$$

The κ posterior moment of g is given by

$$E(g^\kappa|\mathbf{y}, M_1) = \delta^\kappa \frac{B\left(b, \frac{k_{e_1}}{2}+a-\kappa\right) \check{F}_1(\kappa)}{B\left(b, \frac{k_{e_1}}{2}+a\right) \check{F}_1(0)},$$

where

$$\check{F}_1(\kappa) = F_1\left(b, \frac{n+d_0-k_0}{2}, -\frac{n+d_0-k_1}{2}, \frac{k_{e_1}}{2}+a+b-\kappa; \frac{1}{1+\delta R_{10}}, \frac{1}{\delta+1}\right)$$

for $\kappa \in \{0, 1, 2, \dots\}$. Note that $\check{F}_1(0) = \tilde{F}_1(0)$.

The posterior expectation and variance of g are now given by

$$\begin{aligned} E(g|\mathbf{y}, M_1) &= \delta \frac{\tilde{a} + b}{\tilde{a}} \frac{\check{F}_1(1)}{\check{F}_1(0)}, \\ V(g|\mathbf{y}, M_1) &= \delta^2 \frac{\tilde{a} + b}{\tilde{a}} \times \frac{1}{\check{F}_1(0)} \times \left(\frac{\tilde{a} + b - 1}{\tilde{a} - 1} \check{F}_1(2) - \frac{\tilde{a} + b}{\tilde{a}} \frac{\check{F}_1(1)^2}{\check{F}_1(0)} \right), \end{aligned}$$

where $\tilde{a} = \frac{k_{e1}}{2} + a - 1$.

5 Bayesian Inference of the Shrinkage Parameter

5.1 Prior distribution of w

Under model M_1 , the imposed hyper-prior (or mixing) distribution for $w = \delta/(\delta + t) = g/(g + 1)$ is induced via the beta hyper-prior for t (see 7) with parameters given in Table 1. Hence, the resulted prior for w is

$$w \sim BTPD\left(a = \frac{n^* + d_0 - k_1}{2}, b = \frac{n^* + d_1 - d_0 - k_1}{2}, \theta = \frac{\delta}{\delta + 1}, \lambda = 1, \kappa = 1\right),$$

where $BTPD(a, b, \theta, \lambda, \kappa)$ is the *Beta truncated Pareto distribution* (Lourenzutti et al., 2014) with parameters $a, b, \theta, \lambda, \kappa$ and density function

$$f(w; a, b, \theta, \lambda, \kappa) = \frac{1}{B(a, b)} \frac{\kappa \theta^\kappa w^{-\kappa-1}}{1 - (\frac{\theta}{\lambda})^\kappa} \left[\frac{1 - (\frac{\theta}{w})^\kappa}{1 - (\frac{\theta}{\lambda})^\kappa} \right]^{a-1} \left[1 - \frac{1 - (\frac{\theta}{w})^\kappa}{1 - (\frac{\theta}{\lambda})^\kappa} \right]^{b-1},$$

for $\theta < w < \lambda$. The prior mean and variance of w are now given by

$$\begin{aligned} E(w) &= {}_2F_1(1, a, a + b; -1/\delta), \text{ and} \\ Var(w) &= {}_2F_1(2, a, a + b; -1/\delta) - {}_2F_1(1, a, a + b; -1/\delta)^2, \end{aligned}$$

where ${}_2F_1(a_0, b_0, c_0; z)$ is the Gauss hyper-geometric function (Abramowitz and Stegun, 1970) given by

$${}_2F_1(a_0, b_0, c_0; z) = \frac{1}{B(a_0, c_0 - b_0)} \int_0^1 x^{b_0-1} (1-x)^{c_0-b_0-1} (1-zx)^{-a_0} dx.$$

Equivalently we can show that the complementary shrinkage factor $u = 1 - w = 1/(1 + g)$ follows the truncated Compound Confluent Hypergeometric distribution; i.e.

$$\begin{aligned} u \sim tCCH\left(t = a = \frac{n^* + d_0 - k_1}{2}, q = b = \frac{n^* + d_1 - d_0 - k_1}{2}, s = 0, r = a + b, s = 0, \right. \\ \left. v = \delta + 1, \kappa = \frac{\delta + 1}{\delta}\right), \end{aligned}$$

with density expressed as

$$f(u; t, q, r, s, v, \kappa) = \frac{v^t \exp(s/v)}{B(t, q) \Phi_1(q, r, t+q, s/v, 1-\kappa)} \frac{u^{t-1} (1-vu)^{q-1} \exp(-su)}{[\kappa + (1-\kappa)vu]^r} \mathbf{1}_{\{0 < u < \frac{1}{v}\}},$$

where $\Phi_1(\cdot)$ is the Humbert series (Humbert, 1920). Thus the PEP prior is a type of a “Confluent Hypergeometric Information Criterion” (CHIC) g -prior introduced by Li and Clyde (2018). For a comparison with other CHIC g -priors, that have been appeared in literature, see Li and Clyde (2018, Table 1).

In order to get more insight about the behavior of the prior distribution of w , we can obtain approximations of the prior mean and variance by using the first terms of a Taylor expansion given by

$$\mathbb{E}[w(t)] \approx w(\mu_t) + \frac{1}{2} \frac{d^2 w(\mu_t)}{dt^2} \sigma_t^2 \text{ and } \mathbb{V}[w(t)] \approx \left[\frac{dw(\mu_t)}{dt} \right]^2 \sigma_t^2, \tag{19}$$

where μ_t and σ_t^2 are the prior mean and variance of the hyper-parameter t . By implementing the above approach, we obtain the approximations summarized in Table 4; for the PEP prior we restrict attention on the choice of $\delta = n^*$. Note that for small training samples, the dimensions k_0 and k_1 may influence the imposed prior, making it sometimes more informative than intended.

	Prior Mean of w		Prior St. Deviation of w	
	EPP	PEP ($\delta = n^*$)	EPP	PEP ($\delta = n^*$)
For large n^*	$\frac{2n^* - 2k_1 + d_1}{3n^* - 3k_1 + d_1 + d_0}$	$1 - \frac{1}{2n^*}$	$\frac{2}{9n^{*1/2}}$	$\frac{1}{2n^{*3/2}}$
Minimal n^* ($n^* = k_1 + 1$) Reference ($d_0 = d_1 = 0$)	0.704 (for all k_1)	1 (for large k_1)	0.158 (for all k_1)	$1/\sqrt{8}k_1$ (for large k_1)
Jeffreys ($d_0 = k_0, d_1 = k_1$) $k_0 = 1, k_1 = 2$	0.691	0.86	0.128	0.071
Large k_0 , fixed $k = k_1 - k_0$	1/2	1	$\sqrt{\frac{k+1}{8k_0^2}}$	$\sqrt{\frac{2(k+1)}{k_0^4}}$
Fixed k_0 , large $k = k_1 - k_0$	1	1	$\sqrt{\frac{2(k_0+1)}{k^2}}$	$\sqrt{\frac{2(k_0+1)}{k^4}}$

Table 4: Approximate prior means and variances of the shrinkage parameter w .

From Table 4 it is evident that when considering the usual EPP setup with the minimal training sample, then the prior mean of the shrinkage is far away from the value of one for specific cases (e.g. for the reference prior or for the Jeffreys’ dependence prior when $k_0 = 1$ and $k_1 = 2$). This is not the case for the PEP prior for which the prior mean of the shrinkage is close to one even for models of small dimension; for example, under the reference prior and for $k_0 = 1$ and $k_1 = 2$ we obtain a prior mean of the shrinkage equal to 0.86 and a prior standard deviation of the shrinkage equal to 0.071. Generally the global shrinkage parameter w under the PEP prior is close to the value of one implying that the prior is generally non-informative since most of the information is taken from the data.

5.2 Marginal Posterior Distribution of w

Under model M_1 , the marginal posterior distribution of the shrinkage parameter w can directly derived by (17) resulting in

$$\begin{aligned}\pi_1(w|\mathbf{y}) &\equiv \pi(w|\mathbf{y}, M_1) \\ &= C_2 \times (1 + \delta)^{b-1} \\ &\quad \times (1 - w)^{\frac{k_{e1}}{2} + a - 1} w^{-a-b} \left(w - \frac{\delta}{\delta+1}\right)^{b-1} \{1 - w(1 - R_{10})\}^{-\frac{n+d_0-k_0}{2}}, \\ &\quad \text{for } \frac{\delta}{\delta+1} \leq w \leq 1,\end{aligned}$$

where the constant C_2 given by (18).

The posterior κ moment is given by

$$\begin{aligned}E(w^\kappa|\mathbf{y}, M_1) &= C_2 \times (1 + \delta R_{10})^{-\frac{n+d_0-k_0}{2}} \delta^{-a+\kappa} (\delta + 1)^{\frac{n+d_0-k_1}{2} - \kappa} B\left(b, \frac{k_{e1}}{2} + a\right) \tilde{F}_1(\kappa) \\ &= \left(\frac{\delta}{\delta + 1}\right)^\kappa \times \frac{\tilde{F}_1(\kappa)}{\tilde{F}_1(0)},\end{aligned}\quad (20)$$

where

$$\tilde{F}_1(\kappa) = F_1\left(b, \frac{n+d_0-k_0}{2}, -\frac{n+d_0-k_1}{2} + \kappa, \frac{k_{e1}}{2} + a + b; \frac{1}{1+\delta R_{10}}, \frac{1}{\delta+1}\right),$$

for $\kappa \in \{0, 1, 2, \dots\}$. Therefore the posterior expectation and variance of w are directly derived as

$$\begin{aligned}E(w|\mathbf{y}, M_1) &= \frac{\delta}{\delta + 1} \times \frac{\tilde{F}_1(1)}{\tilde{F}_1(0)}, \\ V(w|\mathbf{y}, M_1) &= \left(\frac{\delta}{\delta + 1}\right)^2 \frac{\tilde{F}_1(2)\tilde{F}_1(0) - \tilde{F}_1(1)^2}{\tilde{F}_1(0)^2}.\end{aligned}\quad (21)$$

6 Bayesian Model Averaging, Computation and Model Search

Details about the implementation of Bayesian model averaging (BMA) for PEP priors using the mixture representation are provided in Section 4 of the Appendix. Moreover, in Section 5 of the Appendix we present three alternative MCMC schemes for implementing model search, model averaging and computation of the parameters of interest under the PEP prior using the mixture representation. Specifically, we provide details about: (a) a vanilla MC^3 algorithm (Madigan and York, 1995); (b) an MC^3 algorithm conditional on g ; and (c) an MCMC variable selection scheme based on the Gibbs variable selection of Dellaportas et al. (2002). All these three schemes are summarized in Algorithms 2–4 at Section 5 of the Appendix.

7 Simulation Study

In this Section we illustrate the proposed methodology in simulated data. We compare the performance of PEP prior and the intrinsic prior, the latest as presented in Section 3. We consider 100 data sets of $n = 50$ observations with $p = 15$ covariates. We run two different scenarios. Under Scenario 1 (independence) all covariates are generated from a multivariate Normal distribution with mean vector $\mathbf{0}$ and covariance matrix I_{15} , while the response is generated from

$$Y_i \sim N(4 + 2X_{i1} - X_{i5} + 1.5X_{i7} + X_{i,11} + 0.5X_{i,13}, 2.5^2), \quad (22)$$

for $i = 1, \dots, 50$. Under Scenario 2 (collinearity), the response is generated again from (22), but this time only the first 10 covariates are generated from a multivariate Normal distribution with mean vector $\mathbf{0}$ and covariance matrix I_{10} , while

$$X_{ij} \sim N(0.3X_{i1} + 0.5X_{i2} + 0.7X_{i3} + 0.9X_{i4} + 1.1X_{i5}, 1), \quad (23)$$

for $j = 11, \dots, 15$; $i = 1, \dots, 50$.

With $p = 15$ covariates there are only 32,768 models to compare; we were able to conduct a full enumeration of the model space, obviating the need for a model-search algorithm in this example.

Regarding the prior on model space we consider the uniform prior on model space (uni), as well as the uniform prior on model size (BB), as a special case of the beta-binomial prior (Scott and Berger, 2010); thus in what follows we compare the following methods: PEP-BB, PEP-Uni, I-BB and I-Uni; the first two names denote the PEP prior under the uniform prior on model space and the uniform prior on model size respectively and the last two names the intrinsic prior under the uniform prior on model space and the uniform prior on model size respectively.

Under Scenario 1, the size of the posterior covariate inclusion probabilities for the non-zero effects (see Figure 1), for each method, is in agreement with the size each covariate's effect as expected. Hence the posterior inclusion probabilities for X_1 ($\beta_1 = 2$) are equal to one, with almost zero sampling variability, followed by X_7 ($\beta_7 = 1.5$) with posterior inclusion probabilities close to one, but with almost all values over 80%. For covariates X_5 and X_{11} , the picture for their posterior inclusion probabilities is almost identical due to the same magnitude of the effects in absolute values (equal to one). Moreover, we observe large sampling variability within and across methods. Finally, all methods fail to identify X_{13} ($\beta_{13} = 0.5$) as an important covariate of the model, with the intrinsic approaches giving higher inclusion probabilities on average (around 40%). Nevertheless, the posterior inclusion probabilities for X_{13} are slightly higher on average and more dispersed across different samples, than the zero effects (see Figure 2 for a representative example). Due to the independence of the covariates, we get similar results as the ones presented in Figure 2 for the remaining zero effect covariates, and therefore plots are omitted for brevity reasons. Concerning the comparison of the different methods we observe that: (a) PEP is systematically more parsimonious than

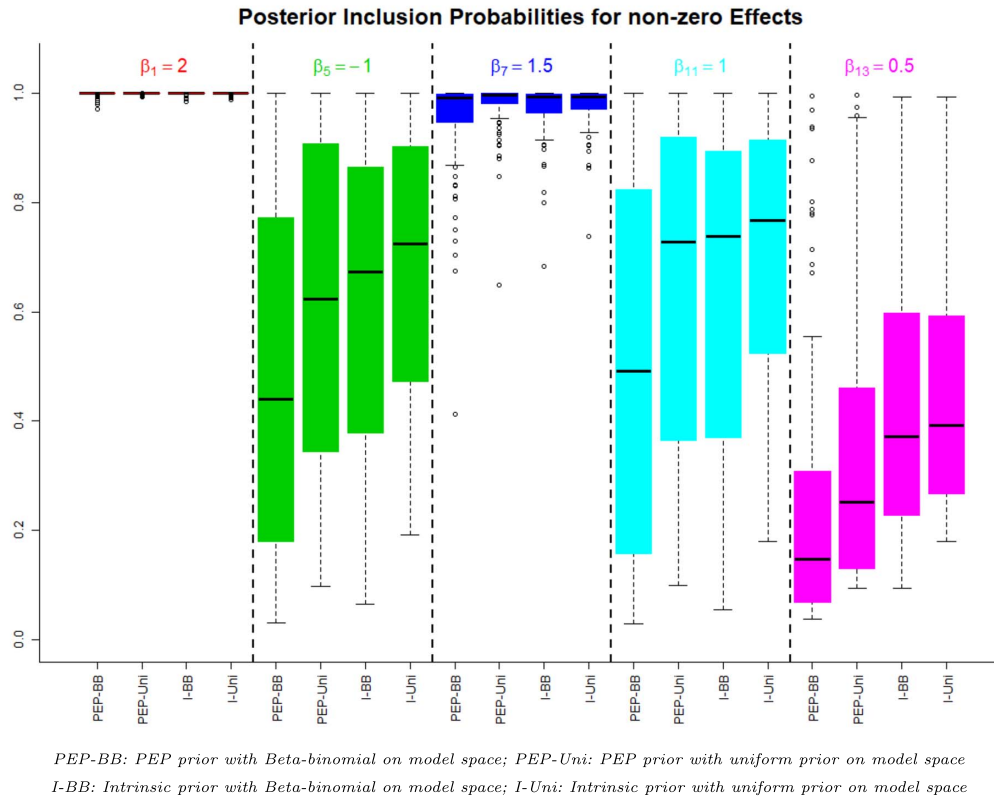


Figure 1: Simulation Scenario 1: Marginal Inclusion Probabilities for Non-Zero Effects.

intrinsic, as previously reported in bibliography (see for example Fouskakis et al., 2020); (b) PEP-BB is more parsimonious than PEP-Uni; (c) I-BB supports slightly more parsimonious solutions than I-Uni. The last finding seems, at a first glance, surprising since the BB prior on model space is promoted in bibliography as a multiplicity adjustment prior. Nevertheless, in this example, the mean of the inclusion probabilities, under the uniform prior, in each data set is around 0.46, which is slightly reduced after the BB implementation to 0.44, leaving the results virtually unchanged. This is in accordance with what is expected by this prior, since it places a U shaped distribution on the prior probabilities of each model depending on its dimensionality. This results in: (a) shrinkage of the inclusion probabilities when the observed proportion of variables and the average of posterior inclusion probabilities under the uniform prior is small (resulting in good sparsity properties in large p problems); (b) inflation of the inclusion probabilities when the observed proportion of variables and the average of posterior inclusion probabilities under the uniform prior is high (leading to posterior support of over-fitted models; a case which is largely neglected in the bibliography); and finally (c) leaving virtually unchanged the posterior inclusion probabilities when the observed proportion

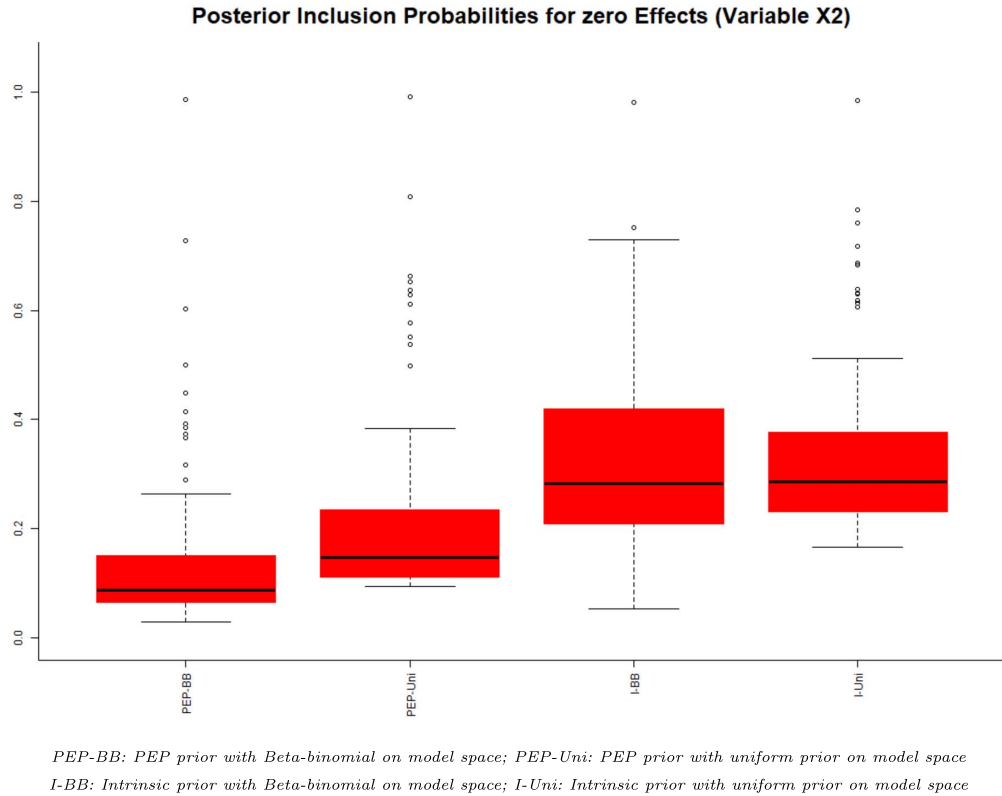


Figure 2: Simulation Scenario 1: Marginal Inclusion Probabilities for X_2 representing Covariates with Zero Effects.

of variables and the average of posterior inclusion probabilities under the uniform prior are close to 0.5. The latter is the case here, where the number of true effects is 5 out of 15 (33%) and the average of the posterior inclusion probabilities under the uniform prior on the model space is equal to 0.46.

Under Scenario 2, the posterior inclusion probabilities for X_1 and X_{13} (see Figure 3), have similar picture as the ones in Scenario 1. For variable X_7 we observe again high inclusion probabilities, but this time with higher uncertainty. Due to collinearity, the posterior inclusion probabilities for covariates X_5 and X_{11} are not longer similar, with the later being close to one (similar picture with the posterior inclusion probabilities for covariate X_7 with true effect equal to 1.5), while the first one (X_5) demonstrates posterior inclusion probabilities around 0.4, or lower, depending on the method. In Figure 4, the posterior inclusion probabilities for all the zero effects are presented. In all cases those are lower under the PEP prior, with the PEP-BB to behave the best. Moreover, they are differences across covariates, depending on collinearities, mainly in the variability across samples.

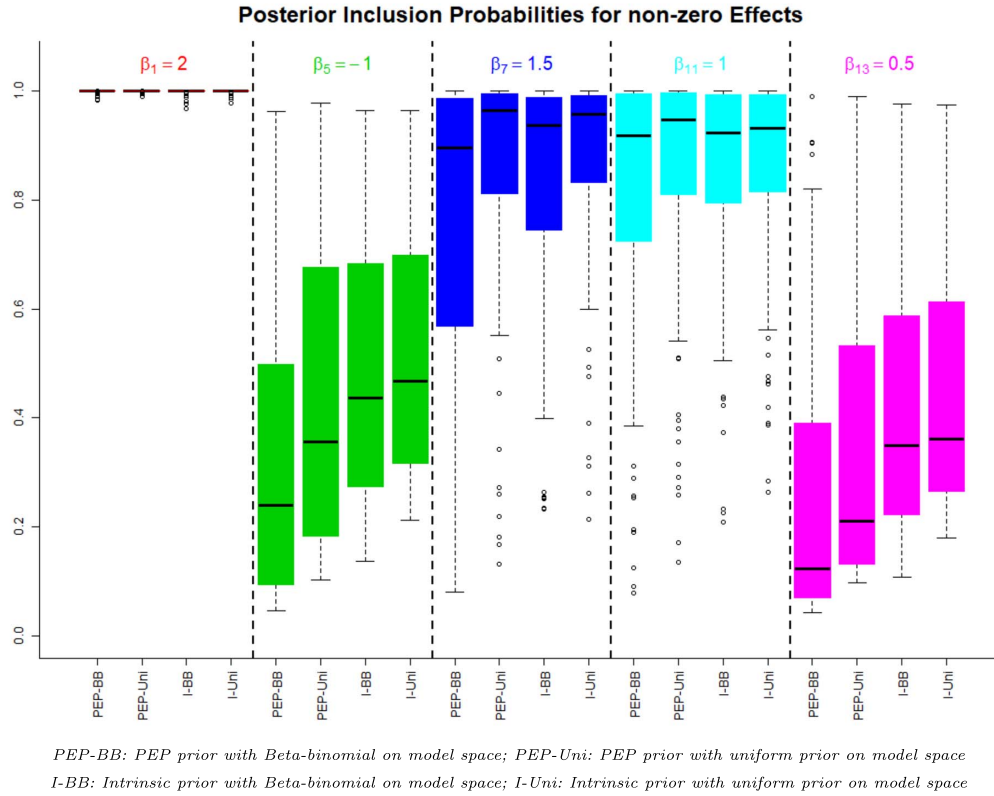


Figure 3: Simulation Scenario 2: Marginal Inclusion Probabilities for non-zero Effects.

8 SDM Dataset

In this Section we consider the SDM data that contains $p = 67$ potential regressors for modelling annual GDP growth per capita between 1960 and 1996 for $n = 88$ countries. More details about this dataset can be found in Ley and Steel (2012).

We compare the performance of the PEP prior (with $\delta = n^* = n$) with that of other scale normal mixtures priors, as presented in Table 2. Specifically, we implement the methods based on: (a) the PEP prior with $\delta = n^* = n$; (b) the intrinsic prior (Womack et al., 2014); (c) the hyper- g and the hyper- g/n priors (Liang et al., 2008) with $a_h = 3$; (d) the robust prior (Bayarri et al., 2012) with $a_r = 0.5$; (e) the benchmark prior (Ley and Steel, 2012) with $c_b = 0.01$. For comparative reasons we have also included the g -prior (Zellner, 1976) with $g = n$.

All methods have been implemented in `MultiBUGS` and `R2MultiBUGS` (Goudie et al., 2020), using the Gibbs variable selection sampler (Dellaportas et al., 2002); see Section 5 of the Appendix for details (Algorithm 4). The obtained results have been generated using 100K MCMC iterations and a 10K burnin period and were additionally compared

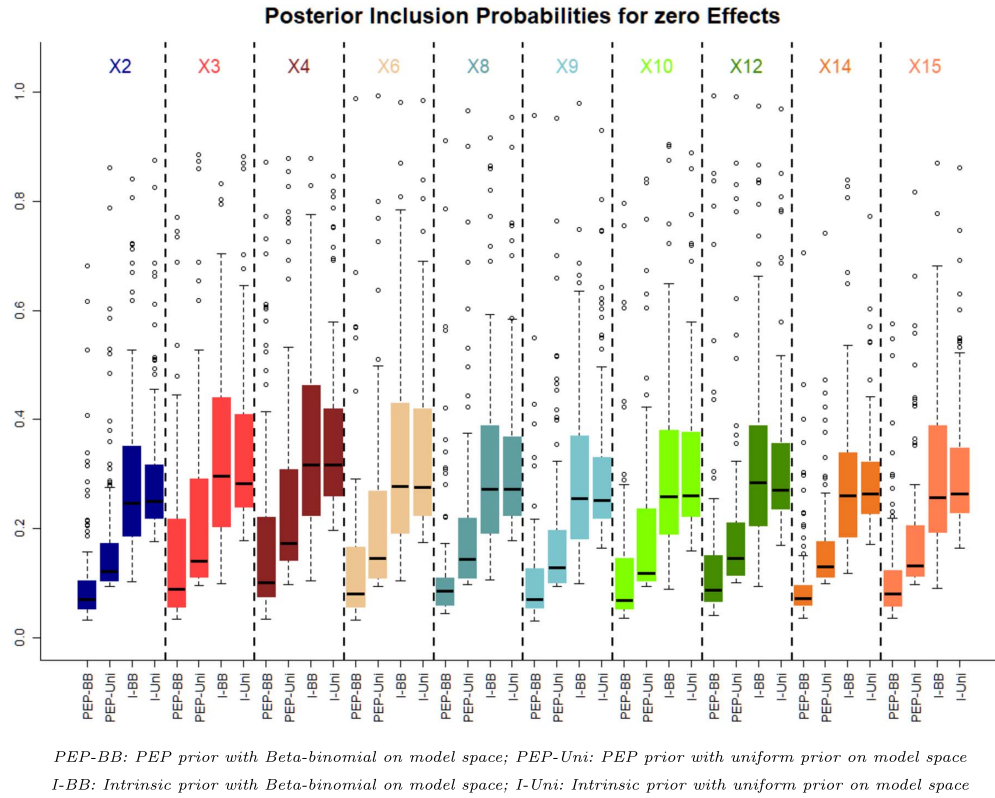


Figure 4: Simulation Scenario 2: Marginal Inclusion Probabilities for Covariates with Zero Effects.

and validated using an MC^3 based algorithm, as described in Section 5 of the Appendix, for the PEP and the intrinsic priors, the `BAS` package in R (Clyde, 2020) for the g , hyper- g and hyper- g/n priors and the `BayesVarSel` package in R (Garcia-Donato and Forte, 2018) for the robust prior.

Regarding the prior on model space we consider, as before, the uniform prior on model space (uni), as well as the uniform prior on model size (BB), as a special case of the beta-binomial prior (Scott and Berger, 2010).

Posterior variable inclusion probabilities and posterior distributions of the model dimension across all visited models are presented in Figures 1–2, respectively, at Section 6.1 of the Appendix, along with detailed discussion and comparison of the related results obtained by all competing methods under consideration using the two prior distributions on the model space.

An interesting point of discussion is the fact that the lower bound imposed on g seems to drive the final results, under the PEP prior. For the recommended PEP prior

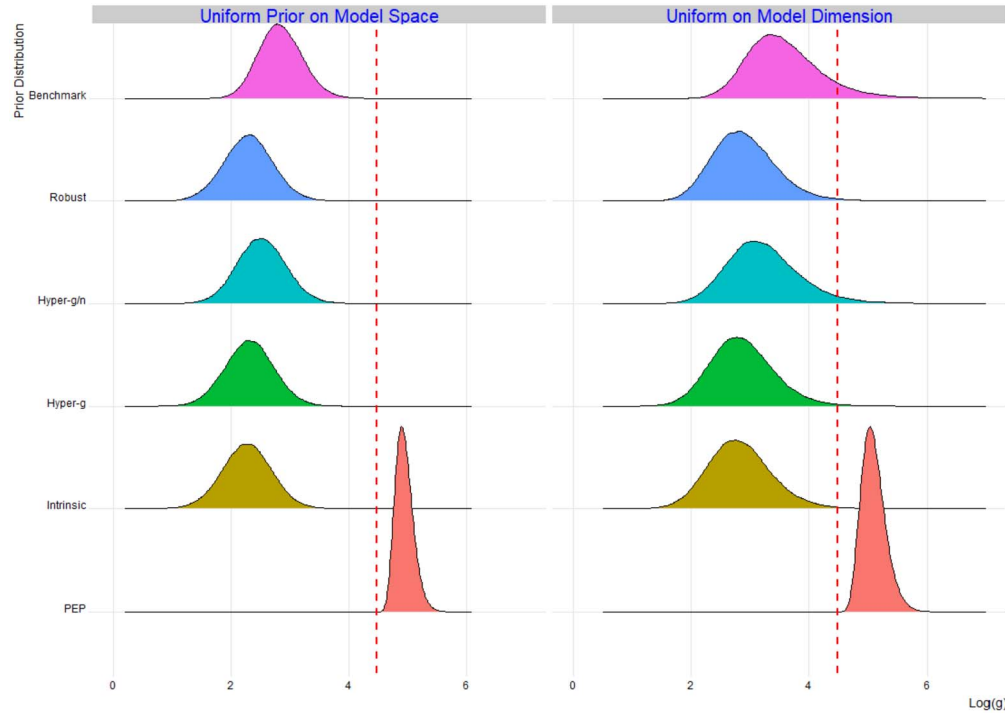


Figure 5: SDM Dataset: Posterior densities of $\log(g)$ for each method (100K MCMC iterations).

specification ($\delta = n$), the posterior density function of g is zero if $g < \delta = n$. Of course we can still specify the PEP prior with smaller values of δ in order to consider different weighting of the imaginary data. By this way, the bound (via the choice of δ) can be lower and thus we might leave g to take values in a wider range. In Figure 5 we present the posterior densities of $\log(g)$ for each competing method (except the g -prior, where g is fixed and equal to n), under both priors on model space, based on the MCMC runs. The vertical dashed lines are referring to the vertical line of $x = \log(n)$. We do not observe any noticeable differences when we move from the uniform prior on the model space to the uniform prior on model size; in the latter case the posterior densities, under each competing method, have a slightly larger variance. On the other hand, there are differences on the posterior distributions of $\log(g)$ when applying different priors on the model parameters. The posterior distributions of $\log(g)$ for all priors, except for the ones under the PEP prior, look similar, with the ones under the hyper- g/n and the benchmark priors to be slightly shifted to the right. Under the PEP prior the posterior distributions of $\log(g)$ are centred on larger values and have smaller standard deviations. Still, the mode of the posterior distribution of $\log(g)$, under the PEP prior, is away from the lower bound, while the posterior standard deviations are high enough to allow for a satisfactorily posterior uncertainty for g . On the other hand, using other hyper-priors for

g without restricting the range of values, like the hyper- g and hyper- g/n , results on high posterior standard deviations of g , which in combination with low posterior modes, may result in a “waste” of valuable posterior probability in informative prior choices (within each model) and to the inflation of the posterior probability of irrelevant models with low practical usefulness. This behaviour has two side effects: (a) the posterior probability of the MAP model is considerably lower than the one obtained by methods with fixed prior choices for g , and (b) the posterior inclusion probabilities for the non-important covariates will be inflated towards 0.5; see Dellaportas et al. (2012) for an empirical illustration within the hyper- g setup.




Figure 6: SDM Dataset: BMA posterior boxplots of $R^{2,(t)}$, $t = 1, \dots, T$, for each method ($T = 100K$ MCMC iterations).

In the following, we further focus on the implementation of a selection of measures concerning the fit and the predictive ability of the models. Figure 6 presents boxplots of the Bayesian version of the coefficient of determination, for each model M_ℓ , given by $R_\ell^2 = 1 - \sigma_\ell^2/S_y^2$, where S_y^2 is the unbiased sample variance, while Figure 7 presents boxplots of the root mean square error of the predictive values, of each model M_ℓ , given by $RMSE_\ell = \frac{1}{n} \sum_{i=1}^n (Y_i - Y_{i,M_\ell}^{pred})$, where Y_{i,M_ℓ}^{pred} were generated from the predictive distribution of each model M_ℓ . These quantities have been calculated for each iteration t of our Gibbs based variable selection algorithm. Therefore, Figures 6 and 7 present the BMA posterior distribution of these quantities; see also Tables 1–2, at Section 6.2




Figure 7: SDM Dataset: BMA posterior boxplots of $RMSE^{(t)}$, $t = 1, \dots, T$, for each method ($T = 100K$ MCMC iterations).

of the Appendix, respectively, for the BMA point estimate of these quantities based on the posterior mean. Both of these quantities evaluate the in-sample overall fit of the models supported by each method. Overall we observe that the PEP and the g -prior under the uniform prior on model space achieve higher posterior R^2 values. On the other hand, all methods when combined with the beta-binomial (uniform on model dimension) prior support models with lower R^2 values. It is notable, that under the BB setup, the PEP prior supports models with lower R^2 values compared to other competing methods, which suggests that the BB prior in combination with PEP overpenalize the inclusion of covariates in sparse problems. Similar is the picture for the RMSE with the inverted relations between methods since lower RMSE values indicate better fitted models. Moreover, the variances of RMSE are now higher since the results are based on the predictive distribution rather than the posterior distribution of the error variance parameter.

Additionally we calculate the *BMA-log predictive score* (BMA-LPS); see for example Fernandez et al. (2001). Specifically, we perform κ -fold cross-validation with $\kappa = 8$, placing 11 randomly selected observations in each fold. We select $\kappa - 1 = 7$ of the folds to form the modelling subsample and the remaining fold to form the validation subsample. We denote by $\mathbb{M} = \{\mathbf{y}^{\mathbb{M}}, \mathbf{X}^{\mathbb{M}}\}$ the modelling subsample, of size $n^{\mathbb{M}} = 77$,

and by $\mathbb{V} = \{\mathbf{y}^{\mathbb{V}}, \mathbf{X}^{\mathbb{V}}\}$ the validation subsample, of size $n^{\mathbb{V}} = 11$, where $n = n^{\mathbb{M}} + n^{\mathbb{V}}$. The BMA-LPS then is given by

$$\text{BMA-LPS} = -\frac{1}{n^{\mathbb{V}}} \sum_{i=1}^{n^{\mathbb{V}}} \log f(y_i^{\mathbb{V}} | \mathbf{y}^{\mathbb{M}}, \mathbf{X}^{\mathbb{V}}), \tag{24}$$

where

$$f(y_i^{\mathbb{V}} | \mathbf{y}^{\mathbb{M}}, \mathbf{X}^{\mathbb{V}}) = \frac{\sum_{M_\ell \in \mathcal{M}} f(y_i^{\mathbb{V}}, \mathbf{y}^{\mathbb{M}} | M_\ell) \pi(M_\ell)}{\sum_{M_\ell \in \mathcal{M}} f(\mathbf{y}^{\mathbb{M}} | M_\ell) \pi(M_\ell)}, \tag{25}$$

with $\pi(M_\ell)$ denoting the prior probabilities of model M_ℓ . Smaller values of BMA-LPS indicate better performance.

In this section we estimate $f(y_i^{\mathbb{V}} | \mathbf{y}^{\mathbb{M}}, \mathbf{X}^{\mathbb{V}})$ from the output of the Gibbs variable selection sampler (see Algorithm 4 at Section 5 of the Appendix) using as an estimate the posterior mean of the normal density evaluated at the model parameter values over the visited model at iteration t ; for $t = 1, \dots, T$. This automatically averages across all visited models.

Prior on models	Prior on model parameters						
	Uniform	PEP	Intrinsic	g -prior	Hyper- g	Hyper- g/n	Robust Benchmark
on models	1.89 (0.31)	1.95 (0.71)	1.75 (0.40)	1.69 (0.25)	1.70 (0.27)	1.70 (0.14)	1.94 (0.35)
on dimension	1.75 (0.37)	1.70 (0.31)	1.84 (0.45)	1.72 (0.31)	1.82 (0.31)	1.75 (0.28)	1.80 (0.40)

Table 5: SDM Dataset: Mean (Standard Deviations) over 8-fold CV of BMA estimates of log-predictive scores obtained using Gibbs based variable selection samplers.

Table 5 presents the means and the standard deviations of the BMA log-predictive score over the eight different modelling and validation combinations of subsamples (folds); see also Figure 3 at Section 6.3 of the Appendix for boxplots of these BMA predictive log-scores. BMA based on hyper- g , hyper- g/n and the robust priors outperform the rest of the methods under the uniform prior on model space for this example. The rest of the methods (g -prior, PEP, intrinsic and the benchmark prior) have higher means and medians (see Figure 3 at Section 6.3 of the Appendix) of log-predictive scores under the uniform prior on model space. Nevertheless, standard deviations do not support important differences between the log-predictive scores of all methods under comparison; note that the log-predictive score for the robust prior has much smaller variability than the corresponding scores for the rest of the methods while for the intrinsic prior the corresponding standard deviation is much higher due to the presence of one outlier. The picture is different when the uniform prior on model dimension is used. In this case, the intrinsic prior seems to perform better than the rest of the methods, followed by the hyper- g prior while the PEP and the robust priors are tied performing slightly worse on average than the two previously mentioned methods. Note that, in terms of medians, the PEP prior performs better than all methods under consideration when the uniform prior on the model dimension is used; see Figure 3 at Section 6.3 of the Appendix. The rest of the methods (g -prior, hyper- g/n and robust) perform slightly worse in terms of mean log-predictive score. The variability of the log-predictive scores is similar for all methods and does not clearly suggest one of the methods as the best.

9 Discussion

In this article we have shown that the power-expected posterior prior, a generalization of the expected-posterior prior in objective Bayesian model selection, under normal linear models can be represented as a mixture of g -prior. This has the great advantage of being able to derive posterior distributions and marginal likelihoods in closed form, permitting fast calculations even when exploring high-dimensional model spaces.

Our results imply that the PEP prior is more parsimonious than its competitors. We do not claim that this property is always the best practice in variable selection problems. The choice of parsimony or sparsity depends on the problem at hand. When we have a sparse dataset, where the important covariates are very few, then the PEP prior will act probably in a better way than other competitors, which may spend a big portion of the posterior probability to models that are impractical in terms of dimension and sparsity.

Additional future extensions of our method include the introduction of two different power parameters in order to derive a family of prior distributions, with members all the prior distributions for variable selection in regression that are written as mixtures of g -priors, that can be derived using either fixed or random imaginary data. Furthermore, we plan to extend the applicability of PEP prior in cases where $k > n$. This can be done by (a) using shrinkage type of baseline priors, such as Lasso or Ridge; (b) assigning zero prior probability to models with dimension larger than n ; and (c) mimicking formal approaches to use g -priors in situations where $k > n$, such as Maruyama and George (2011), based on different ways of generalizing the notion of inverse matrices.

Supplementary Material

Electronic Appendix of the “Power-Expected-Posterior Priors as Mixtures of g -Priors in Normal Linear Models”. (DOI: [10.1214/21-BA1288SUPP](https://doi.org/10.1214/21-BA1288SUPP); .pdf).

References

- Abramowitz, M. and Stegun, I. (1970). *Handbook of Mathematical Functions*. New York: Dover. [MR0208797](#). 1084
- Bayarri, M. J., Berger, J., Forte, A., and Garcia-Donato, G. (2012). “Criteria for Bayesian Model Choice with Application to Variable Selection.” *Annals of Statistics*, 40: 1550–1577. [MR3015035](#). doi: <https://doi.org/10.1214/12-AOS1013>. 1075, 1079, 1080, 1090
- Berger, J., Bernardo, J., and Sun, D. (2009). “The Formal Definition of Reference Priors.” *The Annals of Statistics*, 37: 905–938. [MR2502655](#). doi: <https://doi.org/10.1214/07-AOS587>. 1073
- Berger, J. and Pericchi, L. (1996). “The Intrinsic Bayes Factor for Model Selection and Prediction.” *Journal of the American Statistical Association*, 91: 109–122. [MR1394065](#). doi: <https://doi.org/10.2307/2291387>. 1074

- Berger, J. and Pericchi, L. (2004). “Training Samples in Objective Model Selection.” *Annals of Statistics*, 32: 841–869. MR2065191. doi: <https://doi.org/10.1214/009053604000000238>. 1074
- Casella, G. and Moreno, E. (2006). “Objective Bayesian Variable Selection.” *Journal of the American Statistical Association*, 101: 157–167. MR2268035. doi: <https://doi.org/10.1198/016214505000000646>. 1079
- Clyde, M. (2020). *BAS: Bayesian Variable Selection and Model Averaging using Bayesian Adaptive Sampling*. R package version 1.5.5. 1091
- Consonni, G., Fouskakis, D., Liseo, B., and Ntzoufras, I. (2018). “Prior Distributions for Objective Bayesian Analysis.” *Bayesian Analysis*, 13: 627–679. MR3807861. doi: <https://doi.org/10.1214/18-BA1103>. 1075
- Consonni, G. and Veronese, P. (2008). “Compatibility of Prior Specifications Across Linear Models.” *Statistical Science*, 23: 332–353. MR2483907. doi: <https://doi.org/10.1214/08-STS258>. 1074
- Dellaportas, P., Forster, J., and Ntzoufras, I. (2002). “On Bayesian Model and Variable Selection Using MCMC.” *Statistics and Computing*, 12: 27–36. MR1877577. doi: <https://doi.org/10.1023/A:1013164120801>. 1082, 1086, 1090
- Dellaportas, P., Forster, J., and Ntzoufras, I. (2012). “Joint Specification of Model Space and Parameter Space Prior Distributions.” *Statistical Science*, 27: 232–246. MR2963994. doi: <https://doi.org/10.1214/11-STS369>. 1093
- Fernandez, C., Ley, E., and Steel, M. (2001). “Benchmark Priors For Bayesian Model Averaging.” *Journal of Econometrics*, 100: 381–427. MR1820410. doi: [https://doi.org/10.1016/S0304-4076\(00\)00076-2](https://doi.org/10.1016/S0304-4076(00)00076-2). 1094
- Fouskakis, D. (2019). “Priors via Imaginary Training Samples of Sufficient Statistics for Objective Bayesian Hypothesis Testing.” *Metron*, 77: 179–199. MR4033581. doi: <https://doi.org/10.1007/s40300-019-00159-0>. 1075
- Fouskakis, D. and Ntzoufras, I. (2016). “Limiting Behavior of the Jeffreys Power-Expected-Posterior Bayes Factor in Gaussian Linear Models.” *Brazilian Journal of Probability and Statistics*, 30: 299–320. MR3481105. doi: <https://doi.org/10.1214/15-BJPS281>. 1080
- Fouskakis, D. and Ntzoufras, I. (2017). “Information Consistency of the Jeffreys Power-Expected-Posterior Prior in Gaussian Linear Models.” *Metron*, 75: 371–380. MR3736668. doi: <https://doi.org/10.1007/s40300-017-0110-6>. 1080
- Fouskakis, D. and Ntzoufras, I. (2021). “Supplementary material for: Power-Expected-Posterior Priors as Mixtures of g -Priors in Normal Linear Models.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/21-BA1288SUPP>. 1076
- Fouskakis, D., Ntzoufras, I., and Draper, D. (2015). “Power-Expected-Posterior Priors for variable selection in Gaussian Linear Models.” *Bayesian Analysis*, 10: 75–107. MR3420898. doi: <https://doi.org/10.1214/14-BA887>. 1074, 1077, 1078, 1079

- Fouskakis, D., Ntzoufras, I., and Perrakis, K. (2018). “Power-Expected-Posterior Priors for Generalized Linear Models.” *Bayesian Analysis*, 13: 721–748. MR3807864. doi: <https://doi.org/10.1214/17-BA1066>. 1074
- Fouskakis, D., Ntzoufras, I., and Perrakis, K. (2020). “Variations of Power-Expected-Posterior Priors in Normal Regression Models.” *Computational Statistics and Data Analysis*, 143: 106836; <https://doi.org/10.1016/j.csda.2019.106836>. 1088
- Garcia-Donato, G. and Forte, A. (2018). “Bayesian Testing, Variable Selection and Model Averaging in Linear Models using R with BayesVarSel.” *The R Journal*, 10: 155–174. 1091
- Goudie, R. J. B., Turner, R. M., De Angelis, D., and Thomas, A. (2020). “MultiBUGS: A parallel implementation of the BUGS modelling framework for faster Bayesian inference.” [arXiv:1704.03216](https://arxiv.org/abs/1704.03216). 1090
- Humbert, P. (1920). “Some Extensions of Pincherles Polynomials.” volume 39 of *Proceedings of the Edinburgh Mathematical Society*, 21–24. 1085
- Ibrahim, J. and Chen, M. (2000). “Power Prior Distributions for Regression Models.” *Statistical Science*, 15: 46–60. MR1842236. doi: <https://doi.org/10.1214/ss/1009212673>. 1074
- Kass, R. and Wasserman, L. (1995). “A Reference Bayesian Test for Nested Hypotheses and its Relationship to the Schwarz Criterion.” *Journal of the American Statistical Association*, 90: 928–934. MR1354008. 1074
- Ley, E. and Steel, M. (2012). “Mixtures of g -Priors for Bayesian Model Averaging with Economic Applications.” *Journal of Econometrics*, 171: 251–266. MR2991863. doi: <https://doi.org/10.1016/j.jeconom.2012.06.009>. 1079, 1090
- Li, Y. and Clyde, M. (2018). “Mixtures of g -Priors in Generalized Linear Models.” *Journal of the American Statistical Association*, 113: 1828–1845. MR3902249. doi: <https://doi.org/10.1080/01621459.2018.1469992>. 1085
- Liang, F., Paulo, R., Molina, G., Clyde, M., and Berger, J. (2008). “Mixtures of g Priors for Bayesian Variable Selection.” *Journal of the American Statistical Association*, 103: 410–423. MR2420243. doi: <https://doi.org/10.1198/016214507000001337>. 1076, 1077, 1079, 1090
- Lourenzutti, R., Duarte, D., and Azevedo, M. (2014). *The Beta Truncated Pareto Distribution*. Technical Report. Universidade Federal de Minas Gerais, Belo Horizonte, MG, Brazil. 1084
- Madigan, D. and York, J. (1995). “Bayesian Graphical Models for Discrete Data.” *International Statistical Review*, 63: 215–232. 1086
- Maruyama, Y. and George, E. (2011). “Fully Bayes Factors with a Generalized g -Prior.” *The Annals of Statistics*, 39: 2740–2765. MR2906885. doi: <https://doi.org/10.1214/11-AOS917>. 1079, 1080, 1096
- Moreno, E. and Girón, F. (2008). “Comparison of Bayesian Objective Procedures

- for Variable Selection in Linear Regression.” *Test*, 17: 472–490. MR2470092. doi: <https://doi.org/10.1007/s11749-006-0039-1>. 1076
- Pérez, J. and Berger, J. (2002). “Expected-posterior Prior Distributions for Model Selection.” *Biometrika*, 89: 491–511. MR1929158. doi: <https://doi.org/10.1093/biomet/89.3.491>. 1073
- Scott, J. and Berger, J. (2010). “Bayes and Empirical-Bayes Multiplicity Adjustment in the Variable-Selection Problem.” *The Annals of Statistics*, 38: 2587–2619. MR2722450. doi: <https://doi.org/10.1214/10-AOS792>. 1087, 1091
- Womack, A., Leon-Novelo, L., and Casella, G. (2014). “Inference from Intrinsic Bayes Procedures Under Model Selection and Uncertainty.” *Journal of the American Statistical Association*, 109: 1040–1053. MR3265679. doi: <https://doi.org/10.1080/01621459.2014.880348>. 1079, 1090
- Zellner, A. (1976). “Bayesian and Non-Bayesian Analysis of the Regression Model with Multivariate Student-t Error Terms.” *Journal of the American Statistical Association*, 71: 400–405. MR0405699. 1090
- Zellner, A. and Siow, A. (1980). “Posterior Odds Ratios for Selected Regression Hypothesis (with discussion).” In J. M. Bernardo, M. H. DeGroot, D. V. Lindley and A.F.M. Smith, eds., *Bayesian Statistics*, Vol. 1, 585–606 & 618–647 (discussion). Oxford University Press. MR0862503. 1080