

BAYESIAN VARIABLE SELECTION WITH SHRINKING AND DIFFUSING PRIORS¹

BY NAVEEN NAIDU NARISSETTY AND XUMING HE

University of Michigan

We consider a Bayesian approach to variable selection in the presence of high dimensional covariates based on a hierarchical model that places prior distributions on the regression coefficients as well as on the model space. We adopt the well-known spike and slab Gaussian priors with a distinct feature, that is, the prior variances depend on the sample size through which appropriate shrinkage can be achieved. We show the strong selection consistency of the proposed method in the sense that the posterior probability of the true model converges to one even when the number of covariates grows nearly exponentially with the sample size. This is arguably the strongest selection consistency result that has been available in the Bayesian variable selection literature; yet the proposed method can be carried out through posterior sampling with a simple Gibbs sampler. Furthermore, we argue that the proposed method is asymptotically similar to model selection with the L_0 penalty. We also demonstrate through empirical work the fine performance of the proposed approach relative to some state of the art alternatives.

1. Introduction. We consider the linear regression setup with high dimensional covariates where the number of covariates p can be large relative to the sample size n . When $p > n$, the estimation problem is ill-posed without performing variable selection. A natural assumption to limit the number of parameters in high dimensional settings is that the regression function (i.e., the conditional mean) is sparse in the sense that only a small number of covariates (called active covariates) have nonzero coefficients. We aim to develop a new Bayesian methodology for selecting the active covariates that is asymptotically consistent and computationally convenient. A large number of methods have been proposed for variable selection in the literature from both frequentist and Bayesian viewpoints. Many frequentist methods based on penalization have been proposed following the well-known least absolute shrinkage and selection operator [LASSO, Tibshirani (1996)]. We mention the smoothly clipped absolute deviation [SCAD, Fan and Li (2001)], adaptive LASSO [Zou (2006)], octagonal shrinkage and clustering algorithm for regression [OSCAR, Bondell and Reich (2008)] and the Dantzig selector [Candes and Tao

Received May 2013; revised January 2014.

¹Supported in part by the NSF Grants DMS-12-37234 and DMS-13-07566, and by the National Natural Science Foundation of China Grant 11129101.

MSC2010 subject classifications. 62J05, 62F12, 62F15.

Key words and phrases. Bayes factor, hierarchical model, high dimensional data, shrinkage, variable selection.

(2007); James, Radchenko and Lv (2009)] just to name a few. Fan and Lv (2010) provided a selective overview of high dimensional variable selection methods. Various authors reported inconsistency of LASSO and its poor performance for variable selection under high dimensional settings; see Zou (2006) and Johnson and Rossell (2012). On the other hand, several penalization based methods were shown to have the oracle property [Fan and Li (2001)] under some restrictions on p . For example, Fan and Peng (2004) and Huang and Xie (2007) showed the oracle property for some nonconcave penalized likelihood methods when $p = O(n^{1/3})$ and $p = o(n)$, respectively. Shen, Pan and Zhu (2012) showed that L_0 penalized likelihood method has the oracle property under exponentially large $p = e^{o(n)}$.

Many Bayesian methods have also been proposed for variable selection including the stochastic search variable selection [George and McCulloch (1993)], empirical Bayes variable selection [George and Foster (2000)], spike and slab selection method [Ishwaran and Rao (2005)], penalized credible regions [Bondell and Reich (2012)], nonlocal prior method [Johnson and Rossell (2012)], among others. We shall describe the typical framework used for Bayesian variable selection methods before discussing their theoretical properties.

We use the standard notation $Y_{n \times 1} = X_{n \times p} \beta_{p \times 1} + \varepsilon_{n \times 1}$ to represent the linear regression model. Bayesian variable selection methods usually introduce latent binary variables for each of the covariates to be denoted by $Z = (Z_1, \dots, Z_p)$. The idea is that each Z_i would indicate whether the i th covariate is active in the model or not. For this reason, the prior distribution on the regression coefficient β_i under $Z_i = 0$ is usually a point mass at zero, but a diffused (noninformative) prior under $Z_i = 1$. The concentrated prior of β_i under $Z_i = 0$ is referred to as the spike prior, and the diffused prior under $Z_i = 1$ is called the slab prior. Further, a prior distribution on the binary random vector Z is assumed, which can be interpreted as a prior distribution on the space of models. A Bayesian variable selection method then selects the model with the highest posterior probability. Various selection procedures with this structure have been proposed; they essentially differ in the form of the spike and slab priors, or in the form of the prior on the model space.

Mitchell and Beauchamp (1988) considered a uniform distribution for the slab prior. George and McCulloch (1993) used the Gaussian distribution with a zero mean and a small but fixed variance as the spike prior, and another Gaussian distribution with a large variance as the slab prior. This allowed the use of a Gibbs sampler to explore the posterior distribution of Z . However, as we argue in Section 3, this prior specification does not guarantee model selection consistency at any fixed prior. Ishwaran and Rao (2005) also used Gaussian spike and slab priors, but with continuous bimodal priors for the variance of β to alleviate the difficulty of choosing specific prior parameters. More recently, Ishwaran and Rao (2011) established the oracle property for the posterior mean as n converges to infinity (but p is fixed) under certain conditions on the prior variances. They noted that in the orthogonal design case, a uniform complexity prior leads to correct complexity recovery (i.e., the expected size of the posterior model size converges to the true

model size) under weaker conditions on the prior variances. In another development, Yang and He (2012) used shrinking priors to explore commonality across quantiles in the context of Bayesian quantile regression, but the use of such priors for achieving model selection consistency has not been explored. In this paper, we continue to work with the framework where both the spike and slab priors are Gaussian, but our prior parameters depend explicitly on the sample size through which appropriate shrinkage is achieved. We shall establish model selection consistency properties for general design matrices while allowing p to grow with n at a nearly exponential rate. In particular, the strong selection consistency property we establish is a stronger result for model selection than complexity recovery.

One of the most commonly used priors on the model space is the independent prior given by $P[Z = z] = \prod_{i=1}^p w_i^{z_i} (1 - w_i)^{1-z_i}$, where the marginal probabilities w_i are usually taken to be the same constant. However, when p is diverging, this implies that the prior probability on models with sizes of order less than p goes to zero, which is against model sparsity. We consider marginal probabilities w_i in the order of p^{-1} , which will impose vanishing prior probability on models of diverging size. Yuan and Lin (2005) used a prior that depends on the Gram matrix to penalize models with unnecessary covariates at the prior level. The vanishing prior probability in our case achieves similar prior penalization.

A common notion of consistency for Bayesian variable selection is defined in terms of pairwise Bayes factors, that is, the Bayes factor of any under- or overfitted model with respect to the true model goes to zero. Moreno, Girón and Casella (2010) proved that intrinsic priors give pairwise consistency when $p = O(n)$, and similar consistency of the Bayesian information criterion [BIC, Schwarz (1978)] when $p = O(n^\alpha)$, $\alpha < 1$. Another notion of consistency for both frequentist and Bayesian methods is that the selected model equals the true model with probability converging to one. We refer to this as selection consistency. Bondell and Reich (2012) proposed a method based on penalized credible regions that is shown to be selection consistent when $\log p = O(n^c)$, $c < 1$. Johnson and Rossell (2012) proposed a stronger consistency for Bayesian methods under which the posterior probability of the true model converges to one, which we shall refer to as strong selection consistency. The authors used nonlocal distributions (distributions with small probability mass close to zero) as slab priors, and proved strong selection consistency when $p < n$. However, apart from the limitation $p < n$, their method involves approximations of the posterior distributions and an application of MCMC methods, which are computationally intensive if at all feasible for modest size problems.

We make the following contributions to variable selection in this article. We introduce shrinking and diffusing priors as spike and slab priors, and establish strong selection consistency of the approach for $p = e^{o(n)}$. This approach is computationally advantageous because a standard Gibbs sampler can be used to sample from the posterior. In addition, we find that the resultant selection on the model space is closely related to the L_0 penalized likelihood function. The merits of the

L_0 penalty for variable selection have been discussed by many authors including Schwarz (1978), Liu and Wu (2007), Dicker, Huang and Lin (2013), Kim, Kwon and Choi (2012) and Shen, Pan and Zhu (2012).

We now outline the remaining sections of the paper as follows. The first part of Section 2 describes the model, conditions on the prior parameters and motivation for these conditions. The latter part describes our proposed methodology for variable selection based on the proposed model. Section 3 motivates the use of sample size dependent prior parameters by considering orthogonal design matrices, and provides insight into the variable selection mechanism using those priors. Section 4 presents our main results on the convergence of the posterior distribution of the latent vector Z , and the strong selection consistency of our model selection methodology. Section 5 provides an asymptotic connection between the proposed method and the L_0 penalization. Section 6 provides a discussion on the conditions assumed for proving the results of Section 4. Some computational aspects of the proposed method are noted in Section 7. We present simulation studies in Section 8 to illustrate how the proposed method compares with some existing methods. Application to a gene expression data set is given in Section 9, followed by a conclusion in Section 10. Section 11 provides proofs of some results not given in the earlier sections.

2. The model. From now on, we use p_n to denote the number of covariates to indicate that it grows with n . Consider the $n \times 1$ response vector Y , and the $n \times p_n$ design matrix X corresponding to the p_n covariates of interest. Let β be the regression vector, that is, the conditional mean of Y given X is given by $X\beta$. We assume that β is sparse in the sense that only a few components of β are nonzero; this sparsity assumption can be relaxed as in Condition 4.3. Our goal is to identify the nonzero coefficients to learn about the active covariates. We describe our working model as follows:

$$\begin{aligned}
 Y|X, \beta, \sigma^2 &\sim N(X\beta, \sigma^2 I), \\
 \beta_i | (\sigma^2, Z_i = 0) &\sim N(0, \sigma^2 \tau_{0,n}^2), \\
 \beta_i | (\sigma^2, Z_i = 1) &\sim N(0, \sigma^2 \tau_{1,n}^2), \\
 P(Z_i = 1) &= 1 - P(Z_i = 0) = q_n, \\
 \sigma^2 &\sim \text{IG}(\alpha_1, \alpha_2),
 \end{aligned}
 \tag{1}$$

where i runs from 1 to p_n , $q_n, \tau_{0,n}, \tau_{1,n}$ are constants that depend on n , and $\text{IG}(\alpha_1, \alpha_2)$ is the Inverse Gamma distribution with shape parameter α_1 and scale parameter α_2 .

The intuition behind this set-up is that the covariates with zero or very small coefficients will be identified with zero Z values, and the active covariates will be classified as $Z = 1$. We use the posterior probabilities of the latent variables Z to identify the active covariates.

Notation: We now introduce the following notation to be used throughout the paper.

Rates: For sequences a_n and b_n , $a_n \sim b_n$ means $\frac{a_n}{b_n} \rightarrow c$ for some constant $c > 0$, $a_n \succeq b_n$ (or $b_n \leq a_n$) means $b_n = O(a_n)$, and $a_n \succ b_n$ (or $b_n \prec a_n$) means $b_n = o(a_n)$.

Convergence: Convergence in probability is denoted by \xrightarrow{P} , and equivalence in distribution is denoted by $\stackrel{d}{=}$.

Models: We use k to index an arbitrary model which is viewed as a $p_n \times 1$ binary vector. The i th entry k_i of k indicates whether the i th covariate is active (1) or not (0). We use X_k as the design matrix corresponding to the model k , and β_k to denote the corresponding regression coefficients. In addition, t is used to represent the true model.

Model operations: We use $|k|$ to represent the size of the model k . For two models k and j , the operations $k \vee j$ and $k \wedge j$ denote entry-wise maximum and minimum, respectively. Similarly, $k^c = \mathbf{1} - k$ is entrywise operation, where $\mathbf{1}$ is the vector of 1's. We also use the notation $k \supset j$ (or $k \geq j$) to denote that the model k includes all the covariates in model j , and $k \not\supset j$ otherwise.

Eigenvalues: We use $\phi_{\min}(A)$ and $\phi_{\max}(A)$ to denote the minimum and maximum eigenvalues, respectively, and $\phi_{\min}^\#(A)$ to denote the minimum nonzero eigenvalue (MNEV) of the matrix A . Moreover, we use λ_M^n to be the maximum eigenvalue of the Gram matrix $X'X/n$, and for $\nu > 0$, we define

$$m_n(\nu) = p_n \wedge \frac{n}{(2 + \nu) \log p_n} \quad \text{and} \quad \lambda_m^n(\nu) := \inf_{|k| \leq m_n(\nu)} \phi_{\min}^\# \left(\frac{X'_k X_k}{n} \right).$$

Matrix inequalities: For square matrices A and B of the same order, $A \geq B$ or $(A - B) \geq 0$ means that $(A - B)$ is positive semidefinite.

Residual sum of squares: We define $\tilde{R}_k = Y'(I - X(D_k + X'X)^{-1}X')Y$, where $D_k = \text{Diag}(k\tau_{1n}^{-2} + (\mathbf{1} - k)\tau_{0n}^{-2})$. \tilde{R}_k approximates the usual residual sum of squares $R_k^* = Y'(I - P_k)Y$, where P_k is the projection matrix corresponding to the model k .

Generic constants: We use c' and w' to denote generic positive constants that can take different values each time they appear.

2.1. *Prior parameters.* We consider $\tau_{0,n}^2 \rightarrow 0$ and $\tau_{1,n}^2 \rightarrow \infty$ as n goes to ∞ , where the rates of convergence depend on n and p_n . To be specific, we assume that for some $\nu > 0$, and $\delta > 0$,

$$n\tau_{0n}^2\lambda_M^n = o(1) \quad \text{and} \quad n\tau_{1n}^2\lambda_m^n(\nu) \sim (n \vee p_n^{2+2\delta}).$$

As will be seen later, these rates ensure desired model selection consistency for any $\delta > 0$, where larger values of δ will correspond to higher penalization and vice versa.

Note that the variance τ_{0n}^2 depends on the sample size n and the scale of the Gram matrix. Since the prior distribution of a coefficient under $Z = 0$ is mostly concentrated in

$$\left(-\frac{3\sigma}{\sqrt{n\lambda_M^n}}, \frac{3\sigma}{\sqrt{n\lambda_M^n}}\right),$$

one can view this as the shrinking neighborhood around 0 that is being treated as the region of inactive coefficients. The variance τ_{1n}^2 increases to ∞ , where the rate depends on p_n . However, when $p_n < \sqrt{n}$, τ_{1n}^2 can be of constant order [if $\lambda_m^n(\nu)$ is bounded away from zero].

Now consider the prior probability that a coefficient is nonzero (denoted by q_n). The following calculation gives insight into the choice of q_n . Let K_n be a sequence going to ∞ , then

$$P\left(\sum_{i=1}^{p_n} Z_i > K_n\right) \approx 1 - \Phi\left(\frac{K_n - p_n q_n}{\sqrt{p_n q_n (1 - q_n)}}\right) \rightarrow 0,$$

if $p_n q_n$ is bounded. Therefore, we typically choose q_n such that $q_n \sim p_n^{-1}$. This can be viewed as a priori penalization of the models with large size in the sense that the prior probability on models with diverging number of covariates goes to zero. To this respect, if K is an initial upper bound for the size of the model t , by choosing $q_n = c/p_n$ such that $\Phi((K - c)/\sqrt{c}) \approx 1 - \alpha$, our prior probability on the models with sizes greater than K will be α .

We would like to note that the hierarchical model considered by [George and McCulloch \(1993\)](#) is similar to our model (1), but their prior parameters are fixed and, therefore, do not satisfy our conditions. In Section 3, we give an example illustrating model selection inconsistency under fixed prior parameters.

2.2. Methodology for variable selection. We use the posterior distribution of the latent variables Z_i to select the active covariates. Note that the sample space of Z , denoted by M , has 2^{p_n} points, each of which corresponds to a model. For this reason, we call M the model space. To find the model with the highest posterior probability is computationally challenging for large p_n . In this paper, we use a simpler alternative, that is, we use the p_n marginal posterior probabilities $P(Z_i = 1|Y, X)$, and select the covariates with the corresponding probability more than a fixed threshold $\underline{p} \in (0, 1)$. A threshold probability of 0.5 is a natural choice for \underline{p} . This corresponds to what [Barbieri and Berger \(2004\)](#) call the median probability model. In the orthogonal design case, [Barbieri and Berger \(2004\)](#) showed that the median probability model is an optimal predictive model. The median probability model may not be the same as the maximum a posteriori (MAP) model in general, but the two models are the same with probability converging to one under strong selection consistency.

On the other hand, Dey, Ishwaran and Rao (2008) argued that the median probability model tends to underfit in finite samples. We also consider an alternative by first ranking the variables based on the marginal posterior probabilities and then using BIC to choose among different model sizes. This option avoids the need to specify a threshold. In either case, it is computationally advantageous to use the marginal posterior probabilities, because we need fewer Gibbs iterations to estimate only p_n of them. The proposed methods based on marginal posteriors achieve model selection consistency because the results in Section 4 assure that (i) the posterior probability of the true model converges to 1, and (ii) the marginal posterior based variable selection selects the true model with probability going to 1. We now motivate these results and the necessity of sample size dependent priors in a simple but illustrative case with orthogonal designs.

3. Orthogonal design. In this section, we consider the case where the number of covariates $p_n < n$, and assume that the design matrix X is orthogonal, that is, $X'X = nI$. We also assume σ^2 to be known. Though this may not be a realistic set-up, this simple case provides motivation for the necessity of sample size dependent prior parameters as well as an insight into the mechanism of model selection using these priors. At this moment, we do not impose any assumptions on the prior parameters. *All the probabilities used in the rest of the paper are conditional on X .* Under this simple set-up, the joint posterior of β and Z can be written as

$$\begin{aligned}
 P(\beta, Z | \sigma^2, Y) &\propto \exp\left\{-\frac{1}{2\sigma^2} \|Y - X\beta\|_2^2\right\} \prod_{i=1}^{p_n} ((1 - q_n)\pi_0(\beta_i))^{1-Z_i} (q_n\pi_1(\beta_i))^{Z_i} \\
 &\propto \exp\left\{-\frac{1}{2\sigma^2} (\beta'X'X\beta - 2\beta'X'Y)\right\} \prod_{i=1}^{p_n} ((1 - q_n)\pi_0(\beta_i))^{1-Z_i} (q_n\pi_1(\beta_i))^{Z_i} \\
 &\propto \exp\left\{-\frac{n}{2\sigma^2} \sum_{i=1}^p (\beta_i - \hat{\beta}_i)^2\right\} \prod_{i=1}^{p_n} ((1 - q_n)\pi_0(\beta_i))^{1-Z_i} (q_n\pi_1(\beta_i))^{Z_i},
 \end{aligned}$$

where for $k = 0, 1$, $\pi_k(x) = \phi(x, 0, \sigma^2\tau_{k,n}^2)$ is the probability density function (p.d.f.) of the normal distribution with mean zero and variance $\sigma^2\tau_{k,n}^2$ evaluated at x , and $\hat{\beta}_i$ is the OLS estimator of β_i , that is, $\hat{\beta}_i = X_i'Y/n$.

The product form of the joint posterior of (Z_i, β_i) implies that (Z_i, β_i) and $\{(Z_j, \beta_j), j \neq i\}$ are independent given data. Hence, the marginal posterior of Z_i is given by

$$P(Z_i | \sigma^2, Y) \propto \int \exp\left\{-\frac{n}{2\sigma^2} (b - \hat{\beta}_i)^2\right\} ((1 - q_n)\pi_0(b))^{1-Z_i} (q_n\pi_1(b))^{Z_i} db.$$

Therefore,

$$(2) \quad P(Z_i = 0|\sigma^2, Y) = \frac{(1 - q_n)E_{\hat{\beta}_i}(\pi_0(B))}{(1 - q_n)E_{\hat{\beta}_i}(\pi_0(B)) + q_n E_{\hat{\beta}_i}(\pi_1(B))},$$

where $E_{\hat{\beta}_i}$ is the expectation under B following the normal distribution with mean $\hat{\beta}_i$ and variance σ^2/n . These expectations can be calculated explicitly, that is, for $k = 0$ and 1 ,

$$\begin{aligned} E_{\hat{\beta}_i}(\pi_k(B)) &= \frac{\sqrt{n}}{2\pi\sigma\tau_{k,n}} \int \exp\left\{-\frac{n}{2\sigma^2}(b - \hat{\beta}_i)^2 - \frac{b^2}{2\tau_{k,n}^2}\right\} db \\ &= \frac{1}{\sqrt{2\pi a_{k,n}}} \exp\left\{-\frac{\hat{\beta}_i^2}{2a_{k,n}^2}\right\}, \end{aligned}$$

where $a_{k,n} = \sqrt{\sigma^2/n + \tau_{k,n}^2}$.

This simple calculation gives much insight into the role of our priors and the influence of the prior parameters on variable selection, which we explain in some detail below. In the following subsections, we assume that the i th covariate is identified as active if and only if $P(Z_i = 1|\sigma^2, Y) > 0.5$ for simplicity, and similar arguments can be produced for threshold values other than 0.5.

3.1. *Fixed parameters.* Let us first consider the case of fixed parameters $\tau_{0n}^2 = \tau_0^2 < \tau_{1n}^2 = \tau_1^2$ and $q_n = q = 0.5$. We then have for $k = 0, 1$,

$$(3) \quad E_{\hat{\beta}_i}(\pi_k(B)) \xrightarrow{P} \frac{1}{\tau_k} \exp\left\{-\frac{\beta_i^2}{2\tau_k^2}\right\} \quad \text{as } n \rightarrow \infty \text{ for } \beta_i \neq 0.$$

Now for $\beta_i = \tau_0 \neq 0$, we have $\exp\{-\beta_i^2/2\tau_0^2\}/\tau_0 > \exp\{-\beta_i^2/2\tau_1^2\}/\tau_1$ for any $\tau_1 \neq \tau_0$. Therefore, the limiting value of $P(Z_i = 1|\sigma^2, Y)$ will be less than 0.5 (with high probability) as $n \rightarrow \infty$. This implies that even as $n \rightarrow \infty$, we would not be able to identify the active coefficient in this case.

3.2. *Shrinking $\tau_{0,n}^2$, fixed $\tau_{1,n}^2$ and q_n .* Now consider the prior parameters such that $\tau_{1,n}^2$ and q_n are fixed, but $\tau_{0,n}^2$ goes to 0 with n . If $\beta_i = 0$, $\sqrt{n}\hat{\beta}_i$ converges in distribution to the standard normal distribution, and we have, for $k = 0, 1$,

$$\exp\left\{-\frac{\hat{\beta}_i^2}{2(\sigma^2/n) + 2\tau_{k,n}^2}\right\} = O_P(1).$$

In this case, (3) will imply that $E_{\hat{\beta}_i}(\pi_1(B)) = O_P(1)$, while $E_{\hat{\beta}_i}(\pi_0(B)) \xrightarrow{P} \infty$.

Therefore, from (2), we have $P(Z_i = 0|\sigma^2, Y) \xrightarrow{P} 1$. For $\beta_i \neq 0$, using $\hat{\beta}_i^2 \xrightarrow{P} \beta_i^2$

and the fact that $x e^{-rx^2} \rightarrow 0$ as $x \rightarrow \infty$ (for fixed $r > 0$), we obtain $E_{\hat{\beta}_i}(\pi_0(B)) \rightarrow 0$. As $E_{\hat{\beta}_i}(\pi_1(B)) \sim c'$, for some $c' > 0$, we have $P(Z_i = 1 | \sigma^2, Y) \xrightarrow{P} 1$.

To summarize, we have argued that $P(Z_i = 0 | \sigma^2, Y) \xrightarrow{P} I(\beta_i = 0)$, where $I(\cdot)$ is the indicator function. That is, for orthogonal design matrices, the marginal posterior probability of including an active covariate or excluding an inactive covariate converges to one under shrinking prior parameter $\tau_{0,n}^2$, with fixed parameters $\tau_{1,n}^2$ and q_n . However, it should be noted that this statement is restricted to the convergence of marginals of Z , and does not assure consistency of overall model selection. To achieve this, we will need to allow $\tau_{1,n}^2, q_n$ to depend on the sample size, too.

3.3. *Shrinking and diffusing priors.* Note that the i th covariate is identified as active if and only if

$$\begin{aligned} P(Z_i = 1 | \sigma^2, Y) &> 0.5 \\ \Leftrightarrow q_n E_{\hat{\beta}_i}(\pi_1(B)) &> (1 - q_n) E_{\hat{\beta}_i}(\pi_0(B)) \\ \Leftrightarrow \hat{\beta}_i^2 (a_{0,n}^{-2} - a_{1,n}^{-2}) &> 2(\log(1 - q_n)a_{1,n} - \log q_n a_{0,n}) \\ \Leftrightarrow \hat{\beta}_i^2 &> 2(\log(1 - q_n)a_{1,n} - \log q_n a_{0,n}) / (a_{0,n}^{-2} - a_{1,n}^{-2}) := \varphi_n. \end{aligned}$$

In particular, when $\tau_{0,n}^2 = o(1/n)$, but the other parameters $\tau_{1,n}^2$ and q_n are fixed, we have $\varphi_n \sim \sigma^2 \log n/n$. Without loss of generality, assume that the first $|t|$ coefficients of β are nonzero. For $i > |t|$, $\beta_i = 0$ which implies that $n \hat{\beta}_i^2 \stackrel{d}{=} \chi_1^2$. Therefore,

$$\begin{aligned} P\left[\hat{\beta}_i^2 > \frac{\sigma^2 \log n}{n}\right] &= P[\chi_1^2 > \log n] \\ &\geq \left(\frac{1}{\sqrt{\log n}} - \frac{1}{\sqrt{\log n^3}}\right) e^{-\log n/2} \\ &\geq n^{-1/2-\varepsilon}, \end{aligned}$$

for $\varepsilon > 0$ and sufficiently large n . Therefore, we have

$$\begin{aligned} P[Z = t | \sigma^2, Y] &\leq P\left[\hat{\beta}_i^2 \leq \frac{\sigma^2 \log n}{n}, \forall i > |t|\right] \\ &\leq (1 - n^{-1/2-\varepsilon})^{p_n - |t|} \\ &\rightarrow 0 \quad \text{if } p_n > n^{1/2+2\varepsilon}. \end{aligned}$$

The above argument shows that having $\tau_{1,n}^2$ and q_n fixed leads to inconsistency of selection if the number of covariates is much greater than \sqrt{n} . In this case, the

threshold φ_n should be larger to bound the magnitude of all the inactive covariates simultaneously. By using the diffusing prior parameters Section 2.1, the threshold will be $(2 + \delta)\sigma^2 \log p_n/n$ in place of $\sigma^2 \log n/n$. Model selection consistency with this threshold can be proved using similar arguments in the orthogonal design case. We will defer the rigorous arguments to the next section.

4. Main results. In this section, we consider our model given by (1) and general design matrices. Because the model selection consistency holds easily with $p_n = O(1)$, we assume throughout the paper that $p_n \rightarrow \infty$ as $n \rightarrow \infty$.

4.1. *Conditions.* We first state the main conditions we use.

CONDITION 4.1 (On dimension p_n). $p_n = e^{nd_n}$ for some $d_n \rightarrow 0$ as $n \rightarrow \infty$, that is, $\log p_n = o(n)$.

CONDITION 4.2 (Prior parameters). $n\tau_{0n}^2 = o(1)$, $n\tau_{1n}^2 \sim (n \vee p_n^{2+3\delta})$, for some $\delta > 0$, and $q_n \sim p_n^{-1}$.

CONDITION 4.3 (On true model). $Y|X \sim N(X_t\beta_t + X_{t^c}\beta_{t^c}, \sigma^2 I)$ where the size of the true model $|t|$ is fixed. The coefficients corresponding to the inactive covariates can be nonzero but satisfy $b_0 := \|X_{t^c}\beta_{t^c}\|_2 = O(1)$.

For any fixed K , define

$$\Delta_n(K) := \inf_{\{k: |k| < K|t|, k \not\supset t\}} \|(I - P_k)X_t\beta_t\|_2^2,$$

where P_k is the projection matrix onto the column space of X_k .

CONDITION 4.4 (Identifiability). There is $K > 1 + 8/\delta$ such that $\Delta_n(K) > \gamma_n := 5\sigma^2|t|(1 + \delta) \log(\sqrt{n} \vee p_n)$.

CONDITION 4.5 (Regularity of the design). For some $\nu < \delta, \kappa < (K - 1)\delta/2$,

$$\lambda_M^n < ((n\tau_{0n}^2)^{-1} \wedge n\tau_{1n}^2) \quad \text{and} \quad \lambda_m^n(\nu) \geq \left(\frac{n \vee p_n^{2+2\delta}}{n\tau_{1n}^2} \vee p_n^{-\kappa} \right).$$

The moderateness of these conditions will be examined in some detail in Section 6.

4.2. *Results for fixed σ^2 .* We suppress ν and K from the notation of $\lambda_m^n(\nu)$, $m_n(\nu)$ and $\Delta_n(K)$ for stating the results for convenience. In addition, we introduce the following notation. The posterior ratio of model k with respect to the true model t is defined as

$$\text{PR}(k, t) := P(Z = k|Y, \sigma^2)/P(Z = t|Y, \sigma^2).$$

The following lemma gives an upper bound on the posterior ratio.

LEMMA 4.1. *Under Conditions 4.2 and 4.5, for any model $k \neq t$, we have*

$$\begin{aligned} \text{PR}(k, t) &= \frac{Q_k}{Q_t} s_n^{|k|-|t|} \exp\left\{-\frac{1}{2\sigma^2}(\tilde{R}_k - \tilde{R}_t)\right\} \\ &\leq w'(n\tau_{1n}^2\lambda_m^n(1 - \phi_n))^{-(1/2)(r_k^* - r_t)} (\lambda_m^n)^{-(1/2)|t \wedge k^c|} s_n^{|k|-|t|} \\ &\quad \times \exp\left\{-\frac{1}{2\sigma^2}(\tilde{R}_k - \tilde{R}_t)\right\}, \end{aligned}$$

where $Q_k = |I + XD_k^{-1}X'|^{-1/2}$, $s_n = q_n/(1 - q_n) \sim p_n^{-1}$, $w' > 0$ is a constant, $r_k = \text{rank}(X_k)$, $r_k^* = r_k \wedge m_n$, $\phi_n = o(1)$, $\tilde{R}_k = Y'(I - X(D_k + X'X)^{-1}X')Y$, and $D_k = \text{Diag}(k\tau_{1n}^{-2} + (\mathbf{1} - k)\tau_{0n}^{-2})$.

The following arguments give some heuristics for the convergence of pairwise posterior ratio. Note that \tilde{R}_k is the residual sum of squares from a shrinkage estimator of β , and the term $LR_n := \exp\{-(\tilde{R}_k - \tilde{R}_t)/2\sigma^2\}$ corresponds to the usual likelihood ratio of the two models k and t . Consider a model k that does not include one or more active covariates, then $(\tilde{R}_k - \tilde{R}_t)$ goes to ∞ at the same rate as n , because it is (approximately) the difference in the residual sums of squares of model k and model t . We then have the posterior ratio converging to zero since $LR_n \sim e^{-cn}$ for some $c > 0$, and due to Conditions 4.1–4.5, $P_n := (n\tau_{1n}^2\lambda_m^n(1 - \phi_n))^{(r_t - r_k^*)/2} (\lambda_m^n)^{-|t \wedge k^c|/2} s_n^{|k|-|t|} (1 - \phi_n)^{-|t|/2} = o(e^{cn})$. On the other hand, if the model k includes all the active covariates and one or more inactive covariates, we have $|k| > |t|$, but $(\tilde{R}_k - \tilde{R}_t)$ is probabilistically bounded. The posterior ratio in this case also converges to zero because P_n goes to zero. Note that when $r_k > r_t$, larger values of τ_{1n}^2 will imply smaller P_n . That is, the posterior ratio for large sized models go to zero faster for larger values of τ_{1n}^2 . A similar observation is made by Ishwaran and Rao (2011). To state our main result, we first consider the posterior distributions of the models Z , assuming the variance parameter σ^2 to be known. We consider the case with the prior on σ^2 in Theorem 4.2.

THEOREM 4.1. *Assume Conditions 4.1–4.5. Under model (1), we have $P(Z = t|Y, \sigma^2) \xrightarrow{P} 1$ as $n \rightarrow \infty$, that is, the posterior probability of the true model goes to 1 as the sample size increases to ∞ .*

REMARK 1. The statement of Theorem 4.1 is equivalent to

$$(4) \quad \frac{1 - P(Z = t|Y, \sigma^2)}{P(Z = t|Y, \sigma^2)} = \sum_{k \neq t} \text{PR}(k, t) \xrightarrow{P} 0.$$

REMARK 2. It is worth noting that for Theorem 4.1 to hold, we do not actually need the true σ^2 to be known. Even for a misspecified $\tilde{\sigma}^2 \neq \sigma^2$, $P(Z =$

$t|Y, \tilde{\sigma}^2) \xrightarrow{P} 1$ under the conditions $\Delta_n > \tilde{\sigma}^2 \gamma_n / \sigma^2$ and $2(1 + \delta)\tilde{\sigma}^2 > (2 + \delta)\sigma^2$. The same proof for Theorem 4.1 works.

To see why (4) holds, we provide specific rates of convergence of individual posterior ratio summed over subsets of the model space. We divide the set of models (excluding the model t) into the following subsets:

1. Unrealistically large models: $M_1 = \{k : r_k > m_n\}$, all the models with dimension (i.e., the rank) greater than m_n .
2. Over-fitted models: $M_2 = \{k : k \supset t, r_k \leq m_n\}$, that is, the models of dimension smaller than m_n which include all the active covariates plus one or more inactive covariates.
3. Large models: $M_3 = \{k : k \not\supset t, K|t| < r_k \leq m_n\}$, the models which do not include one or more active covariates, and dimension greater than $K|t|$ but smaller than m_n .
4. Under-fitted models: $M_4 = \{k : k \not\supset t, r_k \leq K|t|\}$, the models of moderate dimension which miss an active covariate.

The proof of Theorem 4.1 shows the following results.

LEMMA 4.2 (Rates of convergence). *For some constants $c', w' > 0$ (which may depend on δ), we have*

1. *The sum of posterior ratio $\sum_{k \in M_1} \text{PR}(k, t) \leq \exp\{-w'n\}$, with probability at least $1 - 2 \exp\{-c'n\}$.*
2. *The sum $\sum_{k \in M_2} \text{PR}(k, t) \leq v_n := (p_n^{-\delta/2} \wedge \frac{p_n^{1+\delta/2}}{\sqrt{n}})$, with probability greater than $1 - \exp\{-c' \log p_n\}$.*
3. *The sum $\sum_{k \in M_3} \text{PR}(k, t) \leq v_n^{(K-1)|t|/2+1}$, with probability greater than $1 - \exp\{-c' K|t| \log p_n\}$.*
4. *For some $w'' < 1$, we have $\sum_{k \in M_4} \text{PR}(k, t) \leq \exp\{-w'(\Delta_n - w''\gamma_n)\}$, with probability greater than $1 - \exp\{-c'\Delta_n\}$.*

4.3. *Results with prior on σ^2 .* We now consider the case with the inverse Gamma prior on the variance parameter σ^2 . Define the constant w as $w := \delta/8(1 + \delta)^2$ in the rest of the section.

THEOREM 4.2. *Under the same conditions as in Theorem 4.1, if we only consider models of dimension at most $|t| + wn / \log p_n$, we have $P(Z = t|Y) \xrightarrow{P} 1$ as $n \rightarrow \infty$.*

REMARK 3. Note that the dimension of the models that need to be excluded for Theorem 4.2 to hold is in the order of $n / \log p_n$. These are unrealistically large models that are uninteresting to us. From now on, we implicitly assume this restriction when a prior distribution is used for σ^2 .

The following corollary ensures that the variable selection procedure based on the marginal posterior probabilities finds the right model with probability tending to 1. It is a direct consequence of Theorems 4.1 and 4.2, but is particularly useful for computations because it ensures that the marginal posterior probabilities can be used for selecting the active covariates.

COROLLARY 4.1. *Under the conditions of Theorem 4.2, we have for any $0 < \underline{p} < 1$, $P[P(Z_i = t_i|Y) > \underline{p} \text{ for all } i = 1, \dots, p_n] \rightarrow 1$ as $n \rightarrow \infty$.*

PROOF. Let E_i be the event that the marginal posterior probability of i th covariate $P(Z_i = t_i|Y) > \underline{p}$. We shall show that $P[\bigcup_{i=1}^{p_n} E_i^c] \rightarrow 0$ as $n \rightarrow \infty$. For each $i = 1, \dots, p_n$, we have

$$\begin{aligned} P(Z_i \neq t_i|Y) &= \sum_{k:k_i \neq t_i} P(Z = k|Y) \\ &\leq \sum_{k \neq t} P(Z = k|Y) \\ &= 1 - P(Z = t|Y). \end{aligned}$$

Then $P[\bigcup_{i=1}^{p_n} E_i^c] = P[P(Z_i = t_i|Y) \leq \underline{p} \text{ for some } i = 1, \dots, p_n] \leq P[P(Z = t|Y) \leq \underline{p}] \rightarrow 0$, due to Theorem 4.2. \square

5. Connection with penalization methods. Due to Lemma 4.1, the maximum a posteriori (MAP) estimate of the model using our Bayesian set-up is equivalent to minimizing the objective function

$$\begin{aligned} (5) \quad B(k) &:= \tilde{R}_k + 2\sigma^2(-(|k| - |t|) \log s_n - \log(Q_k/Q_t)) \\ &= \tilde{R}_k + (|k| - |t|)\psi_{n,k}, \end{aligned}$$

where

$$\psi_{n,k} = 2\sigma^2 \left(-\log s_n - \frac{\log(Q_k/Q_t)}{(|k| - |t|)} \right).$$

Lemma 4.2 implies that with exponentially small probability, the sum of posterior ratio of the models with dimension greater than m_n goes to zero (exponentially) for the fixed σ case. We therefore focus on all the models with dimension less than m_n in this section. In addition, assume that the maximum and minimum nonzero eigenvalues of models of size $2|t|$ are bounded away from ∞ and 0, respectively. Then, due to Condition 4.5 and the proof of Lemma 11.1(iii), we have

$$(6) \quad c \log(n \vee p_n) \leq -\frac{\log(Q_k/Q_t)}{(r_k - r_t)} \leq C \log(n \vee p_n),$$

for some $0 < c \leq C < \infty$.

In particular, if the models with dimension less than m_n are of full rank, that is, $|k| = r_k$, then due to (6), we have

$$(7) \quad 2\sigma^2 c' \log(n \vee p_n) \leq \psi_{n,k} \leq 2\sigma^2 C' \log(n \vee p_n),$$

where $0 < c' \leq C' < \infty$. As $n\tau_{0n}^2 \lambda_M^n \rightarrow 0$, and $n\tau_{1n}^2 \lambda_m^n \rightarrow \infty$,

$$\tilde{R}_k \sim Y'(I - X(1/\tau_{1n}^2 + X'X)^{-1}X')Y = \|Y - \hat{Y}_k\|^2 + O(1).$$

Therefore, the MAP estimate can be (asymptotically) described as the model corresponding to minimizing the following objective function:

$$(8) \quad m(\beta) := \|Y - X\beta\|_2^2 + \psi_{n,k}(\|\beta\|_0 - |t|).$$

Due to the bounds (7) on $\psi_{n,k}$, any inactive covariate will be penalized in the order of $\log(n \vee p_n)$ irrespective of the size of the coefficient. This is however not the case with the L_1 penalty or SCAD penalty, which are directly proportional to the magnitude of the coefficient in some interval around zero.

The commonly used model selection criteria AIC and BIC are special cases of L_0 penalization. The objective functions of AIC and BIC are similar to $m(\beta)$, which have the quotient of penalty equal to 2 and $\log n$ in place of $\psi_{n,k}$. Due to the results in Section 4 and the above arguments, selection properties of our proposed method are similar to those of the L_0 penalty. In particular, it attempts to find the model with the least possible size that could explain the conditional mean of the response variable. A salient feature of our approach is that the L_0 -type penalization is implied by the hierarchical model. The tuning parameters are more transparent than those in penalization methods. Another feature to note is that our model allows high (or even perfect) correlations among inactive covariates. This is practically very useful in high dimensional problems because the number of inactive covariates is often large and the singularity of the design matrix is a common occurrence. Also, high correlations between active and inactive covariates is not as harmful to the proposed method as they are to the L_1 -type penalties. This point is illustrated in Table 4 of our simulation studies in Section 8.

6. Discussion of the conditions. The purpose of this section is to demonstrate that Conditions 4.1–4.5 that we use in Section 4 are quite mild. Condition 4.1 restricts the number of covariates to be no greater than exponential in n , and Condition 4.2 provides the shrinking and diffusing rates for the spike and slab priors, respectively. We note that Conditions 4.3–4.5 allow β to depend on n . For instance, consider $p_n < n$ and the design matrix X with $X'X/n \rightarrow D$, where D is a positive definite matrix. Ishwaran and Rao (2005), Zou (2006), Bondell and Reich (2012) and Johnson and Rossell (2012) assumed this condition on the design under which Conditions 4.3 and 4.4 only require β to be such that

$$\|\beta_{r^c}\|_2^2 = O\left(\frac{1}{n}\right) \quad \text{and} \quad \|\beta_r\|_2^2 > c' \frac{\log n}{n},$$

for some $c' > 0$. Condition 4.5 is also satisfied in this case, so Conditions 4.3–4.5 allow a wider class of design matrices.

In general, Condition 4.4 is a mild regularity condition that allows us to identify the true model. It serves to restrict the magnitude of the correlation between active and inactive covariates, and also to bound the signal to noise ratio from below. The following two remarks provide some insight into the role of Condition 4.4 in these aspects.

REMARK 4. Consider the case where the active coefficients β_t are fixed. We then have some $w' > 0$, such that

$$\begin{aligned} \Delta_n(K) &\geq \|\beta_t\|_2^2 \inf_{\{k:|k|<K|t|,k\not\supset t\}} \phi_{\min}(X'_t(I - P_k)X_t) \\ &\geq w'n \inf_{\{k:|k|<K|t|,k\not\supset t\}} \phi_{\min}\left(\frac{X'_{k\vee t}X_{k\vee t}}{n}\right), \end{aligned}$$

where we have used the fact that $\phi_{\min}(X'_{k\vee t}X_{k\vee t}) \leq \phi_{\min}(X'_t(I - P_k)X_t)$. To see this, we just need to consider the cases where $X_{k\vee t}$ is of full rank. Then it follows from the observation that $(X'_t(I - P_k)X_t)^{-1}$ is a submatrix of $(X'_{k\vee t}X_{k\vee t})^{-1}$. Therefore, Condition 4.4 is satisfied if the minimum eigenvalues of the submatrices of $X'X/n$ with size smaller than $(K + 1)|t|$ are uniformly larger than $c' \log(n \vee p_n)/n$. In the other end of the spectrum, where the inactive covariates can be perfectly correlated, Condition 4.4 could still hold.

REMARK 5. If the infimum of $\phi_{\min}(X'_t(I - P_k)X_t/n)$ is uniformly bounded away from zero, then $\Delta_n(K) \geq w'n \|\beta_t\|_2^2$. Then Condition 4.4 is satisfied if

$$\left\| \frac{\beta_t}{\sigma} \right\|_2^2 \geq \frac{c' \log(n \vee p_n)}{n}.$$

Condition 4.5 provides conditions on the eigenvalues of the Gram matrix in terms of the prior parameters. The condition is weaker than the assumption that the maximum and minimum nonzero eigenvalues of the Gram matrix are bounded away from infinity and zero, respectively. In Condition 4.5, $\lambda_M^n < (n\tau_{0n}^2)^{-1}$ will be satisfied if τ_{0n}^2 is small enough. However, the assumption on $\lambda_m^n(\nu)$ is nontrivial as it needs to be greater than $p_n^{-\kappa}$. We now show that this requirement is satisfied with high probability if the design matrix consists of independent sub-Gaussian rows.

LEMMA 6.1 (MNEV for sub-Gaussian random matrices). *Suppose that the rows of $X_{n \times p_n}$ are independent isotropic sub-Gaussian random vectors in R^{p_n} . Then there exists a $\nu > 0$ such that, with probability greater than $1 - \exp(-w'n)$,*

$$\inf_{|k| \leq m_n(\nu)} \phi_{\min}\left(\frac{X'_k X_k}{n}\right) > 0.$$

A proof of Lemma 6.1 is provided in Section 11. Lemma 6.1 implies that the Gram matrix of a sub-Gaussian design matrix has the minimum eigenvalues of all the $m_n(v)$ dimensional submatrices to be uniformly bounded away from zero. This clearly is stronger than Condition 4.5, which only requires the minimum nonzero eigenvalues to be uniformly greater than $p_n^{-\kappa}$. In particular, unlike the restricted isometry conditions which control the minimum eigenvalue, Condition 4.5 allows the minimum eigenvalue to be exactly zero to allow even perfect correlation among inactive (or active) covariates.

7. Computation. The implementation of our proposed method involves using the Gibbs sampler to draw samples from the posterior of Z . The full conditionals are standard distributions due to the use of conjugate priors. The conditional distribution of β is given by

$$f(\beta|Z, \sigma^2, Y) \propto \exp\left\{-\frac{1}{2\sigma^2}\|Y - X\beta\|_2^2\right\} \prod_{i=1}^{p_n} \phi(\beta_i, 0, \sigma^2\tau_{Z_i,n}^2),$$

where $\phi(x, 0, \tau^2)$ is the p.d.f. of the normal distribution with mean zero, and variance τ^2 evaluated at x . This can be rewritten as

$$f(\beta|Z = k, \sigma^2, Y) \propto \exp\left\{-\frac{1}{2\sigma^2}(\beta'X'X\beta - 2\beta'X'Y)\right\} \exp\left\{-\frac{1}{2\sigma^2}\beta'D_k\beta\right\},$$

where $D_k = \text{Diag}(\tau_{k_i,n}^{-2})$. Hence, the conditional distribution of β is given by $\beta \sim N(m, \sigma^2V)$, where $V = (X'X + D_k)^{-1}$, and $m = VX'Y$. Furthermore, the conditional distribution of Z_i is

$$P(Z_i = 1|\beta, \sigma^2) = \frac{q_n\phi(\beta_i, 0, \sigma^2\tau_{1,n}^2)}{q_n\phi(\beta_i, 0, \sigma^2\tau_{1,n}^2) + (1 - q_n)\phi(\beta_i, 0, \sigma^2\tau_{0,n}^2)}.$$

The conditional of σ^2 is the inverse Gamma distribution $IG(a, b)$ with $a = \alpha_1 + n/2 + p_n/2$, and $b = \alpha_2 + \beta'D_k\beta/2 + (Y - X\beta)'(Y - X\beta)/2$.

The only possible computational difficulty in the Gibbs sampling algorithm is the step of drawing from the conditional distribution of β , which is a high dimensional normal distribution for large values of p_n . However, due to the structure of the covariance matrix $(X'X + D_k)^{-1}$, it can be efficiently sampled using block updating that only requires drawing from smaller dimensional normal distributions. Details of the block updating can be found in [Ishwaran and Rao \(2005\)](#).

8. Simulation study. In this section, we study performance of the proposed method in several experimental settings, and compare it with some existing variable selection methods. We will refer to the proposed method as BASAD for BAYesian Shrinking And Diffusing priors.

The proposed BASAD method has three tuning parameters. In all our empirical work, we use

$$\tau_{0n}^2 = \frac{\hat{\sigma}^2}{10n}, \quad \tau_{1n}^2 = \hat{\sigma}^2 \max\left(\frac{p_n^{2.1}}{100n}, \log n\right),$$

where $\hat{\sigma}^2$ is the sample variance of Y , and we choose $q_n = P[Z_i = 1]$ such that $P[\sum_{i=1}^{p_n} Z_i = 1 > K] = 0.1$, for a prespecified value of K . Our default value is $K = \max(10, \log(n))$, unless otherwise specified in anticipation of a less sparse model. The purpose of using $\hat{\sigma}^2$ is to provide appropriate scaling. If a preliminary model is available, it is better to use as $\hat{\sigma}^2$ the residual variance from such a model. It is clear that those choices are not optimized for any given problem, but they provide a reasonable assessment on how well BASAD can do. In the simulations, we use 1000 burn-in iterations for the Gibbs sampler followed by 5000 updates for estimating the posterior probabilities. As mentioned in Section 2, we consider both the median probability model (denoted by BASAD) and the BIC-based model (denoted by BASAD.BIC) where the threshold for marginal posterior probability is chosen by the BIC. The R function used for obtaining the results in this section is publicly available on the authors' website.

In this paper, we report our simulation results for six cases under several (n, p) combinations, varied correlations, signal strengths and sparsity levels.

- *Case 1:* In the first case, we use the set-up of [Johnson and Rossell \(2012\)](#) with $p = n$. Two sample sizes, $n = 100$ and $n = 200$, are considered, and the covariates are generated from the multivariate normal distributions with zero mean and unit variance. The compound symmetric covariance with pairwise covariance of $\rho = 0.25$ is used to represent correlation between covariates. Five covariates are taken active with coefficients $\beta_t = (0.6, 1.2, 1.8, 2.4, 3.0)$. This is a simple setting with moderate correlation between covariates and strong signal strength.
- *Case 2:* We consider the $p > n$ scenario with $(n, p) = (100, 500)$ and $(n, p) = (200, 1000)$, but the other parameters are same as in case 1.

For the next three cases (cases 3–5), we keep $(n, p) = (100, 500)$ but vary model sparsity, signal strength and correlation among covariates.

- *Case 3:* We keep $\rho = 0.25$ and $|t| = 5$ but have low signals $\beta_t = (0.6, 0.6, 0.6, 0.6, 0.6)$.
- *Case 4:* We consider a block covariance setting where the active covariates have common correlation (ρ_1) equal to 0.25, the inactive covariates have common correlation (ρ_3) equal to 0.75 and each pair of active and inactive covariate has correlation (ρ_2) 0.50. The other aspects of the model are the same as in case 1.
- *Case 5:* We consider a less sparse true model with $|t| = 25$ and β_t is the vector containing 25 equally spaced values between 1 and 3 (inclusive of 1 and 3).
- *Case 6:* We consider the more classical case of $n > p$ with $(n, p) = (100, 50)$ and $(n, p) = (200, 50)$. Following [Bondell and Reich \(2012\)](#), the covariates are drawn from a normal distribution with the covariance matrix distributed as the

TABLE 1

Performance of BASAD for case 1: $n = p$. The methods under comparison are piMOM with nonlocal priors of Johnson and Rossell (2012), BCR.Joint of Bondell and Reich (2012), SpikeSlab of Ishwaran and Rao (2005) and three penalization methods Lasso, elastic net (EN), and SCAD tuned by the BIC. The other columns of the table are as follows: pp_0 and pp_1 (when applicable) are the average posterior probabilities of inactive and active variables, respectively; $Z = t$ is the proportion that the exact models is selected; $Z \supset t$ is the proportion that the selected model contains all the active covariates; FDR is the false discovery rate, and MSPE is the mean squared prediction error of the selected models

	pp_0	pp_1	$Z = t$	$Z \supset t$	FDR	MSPE
$(n, p) = (100, 100); \rho = 0.25; t = 5$						
BASAD	0.016	0.985	0.866	0.954	0.015	1.092
BASAD.BIC	0.016	0.985	0.066	0.996	0.256	1.203
piMOM	0.012	0.991	0.836	0.982	0.030	1.083
BCR.Joint			0.442	0.940	0.157	1.165
SpikeSlab			0.005	0.216	0.502	1.660
Lasso.BIC			0.010	0.992	0.430	1.195
EN.BIC			0.398	0.982	0.154	1.134
SCAD.BIC			0.356	0.990	0.160	1.157
$(n, p) = (200, 200); \rho = 0.25; t = 5$						
BASAD	0.002	1.000	0.944	1.000	0.009	1.037
BASAD.BIC	0.002	1.000	0.090	1.000	0.187	1.087
piMOM	0.003	1.000	0.900	1.000	0.018	1.038
BCR.Joint			0.594	0.994	0.102	1.064
SpikeSlab			0.008	0.236	0.501	1.530
Lasso.BIC			0.014	1.000	0.422	1.101
EN.BIC			0.492	1.000	0.113	1.056
SCAD.BIC			0.844	1.000	0.029	1.040

Wishart distribution centered at the identity matrix with p degrees of freedom. Three of the 50 covariates are taken to be active with their coefficients drawn from the uniform distribution $U(0, 3)$ to imply a mix of weak and strong signals.

The summary of our results are presented in Tables 1–6. In those tables, BASAD denotes the median probability model, BASAD.BIC denotes the model obtained by using the threshold probability chosen by the BIC. Three competing Bayesian model selection methods are: (1) piMOM, the nonlocal prior method proposed by Johnson and Rossell (2012) but only when $p \leq n$; (2) BCR.Joint, the Bayesian joint credible region method of Bondell and Reich (2012) (using the default priors followed by an application of BIC); (3) SpikeSlab, the generalized elastic net model obtained using the R package spikeslab [Ishwaran, Kogalur and Rao (2010)] for the spike and slab method of Ishwaran and Rao (2005). Three penalization methods under consideration are: (1) LASSO; (2) Elastic Net (EN); and (3) SCAD, all tuned by the BIC. Our simulation experiment used 500 data sets

TABLE 2
Performance of BASAD for case 2: $p > n$

	pp_0	pp_1	$Z = t$	$Z \supset t$	FDR	MSPE
$(n, p) = (100, 500); \rho = 0.25; t = 5$						
BASAD	0.001	0.948	0.730	0.775	0.011	1.130
BASAD.BIC	0.001	0.948	0.190	0.915	0.146	1.168
BCR.Joint			0.070	0.305	0.268	1.592
SpikeSlab			0.000	0.040	0.626	3.351
Lasso.BIC			0.005	0.845	0.466	1.280
EN.BIC			0.135	0.835	0.283	1.223
SCAD.BIC			0.045	0.980	0.328	1.260
$(n, p) = (200, 1000); \rho = 0.25; t = 5$						
BASAD	0.000	0.986	0.930	0.950	0.000	1.054
BASAD.BIC	0.000	0.986	0.720	0.990	0.046	1.060
BCR.Joint			0.090	0.250	0.176	1.324
SpikeSlab			0.000	0.050	0.574	1.933
Lasso.BIC			0.020	1.000	0.430	1.127
EN.BIC			0.325	1.000	0.177	1.077
SCAD.BIC			0.650	1.000	0.091	1.063

from each model when $n \geq p$, but used 200 data sets when $p > n$ to aggregate the results.

The columns of the tables show the average marginal posterior probability assigned to inactive covariates and active covariates (pp_0 and pp_1 , resp.), proportion of choosing the true model ($Z = t$), proportion of including the true model ($Z \supset t$) and false discovery rate (FDR). The last column (MSPE) gives the average test mean squared prediction error based on n new observations as testing data. From our simulation experiment, we have the following findings:

TABLE 3
Performance of BASAD for case 3: $(n, p) = (100, 500)$

	pp_0	pp_1	$Z = t$	$Z \supset t$	FDR	MSPE
$\rho = 0.25; t = 5; \beta_t = (0.6, 0.6, 0.6, 0.6, 0.6)$						
BASAD	0.002	0.622	0.185	0.195	0.066	2.319
BASAD.BIC	0.002	0.622	0.160	0.375	0.193	1.521
BCR.Joint			0.030	0.315	0.447	1.501
SpikeSlab			0.000	0.000	0.857	2.466
Lasso.BIC			0.000	0.520	0.561	1.555
EN.BIC			0.040	0.345	0.478	1.552
SCAD.BIC			0.045	0.340	0.464	1.561

TABLE 4
Performance of BASAD for case 4: $(n, p) = (100, 500)$

	<i>pp</i> ₀	<i>pp</i> ₁	$Z = t$	$Z \supset t$	FDR	MSPE
$\rho_1 = 0.25, \rho_2 = 0.50, \rho_3 = 0.75$						
BASAD	0.002	0.908	0.505	0.530	0.012	1.199
BASAD.BIC	0.002	0.908	0.165	0.815	0.179	1.210
BCR.Joint			0.000	0.000	0.515	2.212
SpikeSlab			0.000	0.000	0.995	10.297
Lasso.BIC			0.000	0.015	0.869	8.579
EN.BIC			0.000	0.000	0.898	8.360
SCAD.BIC			0.000	0.000	0.899	8.739

(i) The Bayesian model selection methods BASAD and piMOM (whenever available) tend to perform better than the other methods in terms of selecting the true model and controlling the false discovery rate in variable selection, and our proposed BASAD stands out in this regard. The penalization methods often have

TABLE 5
Performance of BASAD for case 5: $(n, p) = (100, 500)$. In this case, two versions of BASAD are included, where BASAD.K10 uses our default value of $K = 10$, and BASAD.K50 uses a less sparse specification of $K = 50$

	<i>pp</i> ₀	<i>pp</i> ₁	$Z = t$	$Z \supset t$	FDR	MSPE
$\rho = 0.25; t = 25$						
BASAD.K50	0.020	0.988	0.650	0.950	0.036	3.397
BASAD.BIC.K50	0.020	0.988	0.005	0.960	0.283	4.019
BASAD.K10	0.003	0.548	0.405	0.420	0.011	170.862
BASAD.BIC.K10	0.003	0.548	0.035	0.430	0.076	88.881
BCR.Joint			0.000	0.000	0.622	49.299
SpikeSlab			0.000	0.000	0.816	111.911
Lasso.BIC			0.000	0.005	0.685	58.664
EN.BIC			0.000	0.000	0.693	59.058
SCAD.BIC			0.000	0.000	0.666	72.122
$\rho = 0.75; t = 25$						
BASAD.K50	0.048	0.914	0.005	0.355	0.289	6.103
BASAD.BIC.K50	0.048	0.914	0.000	0.445	0.498	6.611
BASAD.K10	0.003	0.298	0.025	0.030	0.018	349.992
BASAD.BIC.K10	0.003	0.298	0.000	0.060	0.087	61.709
BCR.Joint			0.000	0.000	0.772	34.113
SpikeSlab			0.000	0.000	0.899	48.880
Lasso.BIC			0.000	0.000	0.734	24.310
EN.BIC			0.000	0.000	0.754	29.171
SCAD.BIC			0.000	0.000	0.736	27.236

TABLE 6
Performance of BASAD for case 6: $n > p$

	pp_0	pp_1	$Z = t$	$Z \supset t$	FDR	MSPE
$(n, p) = (100, 50)$						
BASAD	0.037	0.899	0.654	0.714	0.026	1.086
BASAD.BIC	0.037	0.899	0.208	0.778	0.267	1.151
piMOM	0.011	0.892	0.656	0.708	0.021	1.066
SpikeSlab			0.064	0.846	0.567	1.226
BCR.Joint			0.336	0.650	0.216	1.124
Lasso.BIC			0.076	0.744	0.397	1.152
EN.BIC			0.378	0.742	0.194	1.110
SCAD.BIC			0.186	0.772	0.284	1.147
$(n, p) = (200, 50)$						
BASAD	0.026	0.926	0.738	0.784	0.017	1.029
BASAD.BIC	0.026	0.926	0.338	0.842	0.193	1.055
piMOM	0.005	0.908	0.694	0.740	0.020	1.036
BCR.Joint			0.484	0.770	0.133	1.045
SpikeSlab			0.038	0.900	0.629	1.121
Lasso.BIC			0.082	0.752	0.378	1.059
EN.BIC			0.428	0.748	0.165	1.039
SCAD.BIC			0.358	0.812	0.193	1.046

higher probabilities of selecting all the active covariates at the cost of overfitting and false discoveries. In terms of the prediction error, however, BASAD does not always outperform its competitors, but remains competitive.

(ii) When the signals are low (case 3), all the methods under consideration have trouble finding the right model, and BASAD.BIC results in lower prediction error than BASAD with 0.5 as the threshold for posterior probabilities. In most cases, BASAD.BIC leads to slightly higher false positive rates than BASAD with similar prediction errors.

(iii) In case 4, there is a moderate level of correlation among inactive covariates and some level of correlation between active and inactive covariates. This is where BASAD outperforms the other methods under consideration because BASAD is similar to the L_0 penalty and is able to accommodate such correlations well. Please refer to our discussion in Sections 5 and 6.

(iv) When the true model is not so sparse and has $|t| = 25$ active covariates (case 5), our default choice of $K = 10$ in BASAD did not perform well, which is not surprising. In fact, no other methods under consideration did well in this case, highlighting the difficulty of finding a nonsparse model with a limited sample size. On the other hand, there is some promising news. If we anticipate a less sparse model with $K = 50$, the proposed method BASAD improved the performance considerably. Our empirical experience suggests that if we are uncertain about the level of sparsity of our model, we may use a generous choice of K or use BIC to choose between different values of K .

9. Real data example. In this section, we apply our variable selection method to a real data set to examine how it works in practice. We consider the data from an experiment conducted by [Lan et al. \(2006\)](#) to study the genetics of two inbred mouse populations (B6 and BTBR). The data include expression levels of 22,575 genes of 31 female and 29 male mice resulting in a total of 60 arrays. Some physiological phenotypes, including the numbers of phosphoenopyruvate carboxykinase (PEPCK) and glycerol-3-phosphate acyltransferase (GPAT) were also measured by quantitative real-time PCR. The gene expression data and the phenotypic data are available at GEO (<http://www.ncbi.nlm.nih.gov/geo>; accession number GSE3330). [Zhang, Lin and Zhang \(2009\)](#) used orthogonal components regression to predict each phenotype based on the gene expression data. [Bondell and Reich \(2012\)](#) used the Bayesian credible region method for variable selection on the same data.

Because this is an ultra-high dimensional problem with $p_n = 22,575$, we prefer to perform simple screenings of the genes first based on the magnitude of marginal correlations with the response. The power of marginal screening has been recognized by [Fan and Lv \(2008\)](#). After the screening, the dataset for each of the responses consisted of $p = 200$ and 400 predictors (including the intercept and gender) by taking 198 and 398 genes based on marginal screening. We performed variable selection with BASAD along with LASSO, SCAD and the BCR method. Following [Bondell and Reich \(2012\)](#), we randomly split the sample into a training set of 55 observations and a test set with the remaining five observations. The fitted models using the training set were used to predict the response in the test set. This process was repeated 100 times to estimate the prediction error.

In [Figure 1](#), we plot the average mean square prediction error (MSPE) for models of various sizes chosen by BASAD, BCR and SCAD methods for the two responses PEPCK and GPAT. We find that the MSPE of BASAD is mostly smaller than that for other methods across different model sizes. In particular, BASAD chooses less correlated variables and achieves low MSPE with fewer predictive genes than the other methods. We also note that the 10-covariate models chosen by BASAD is very different (with the overlap of just one covariate for PEPCK and three covariates for GPAT) from those of SCAD which chose mostly the same covariates as LASSO. There are four common covariates identified by both BASAD and BCR methods. When we perform a linear regression by including the covariates chosen by BASAD and SCAD, we noticed that majority of the covariates chosen by BASAD are significant, which indicates that those genes chosen by BASAD are significant in explaining the response even in the presence of those chosen using SCAD. Most of the genes selected by SCAD, however, are not significant in the presence of those chosen by BASAD. Despite the evidence in favor of the genes selected by BASAD in this example, we must add that the ultimate assessment of a chosen model would need to be made by additional information from the subject matter science and/or additional experiment.

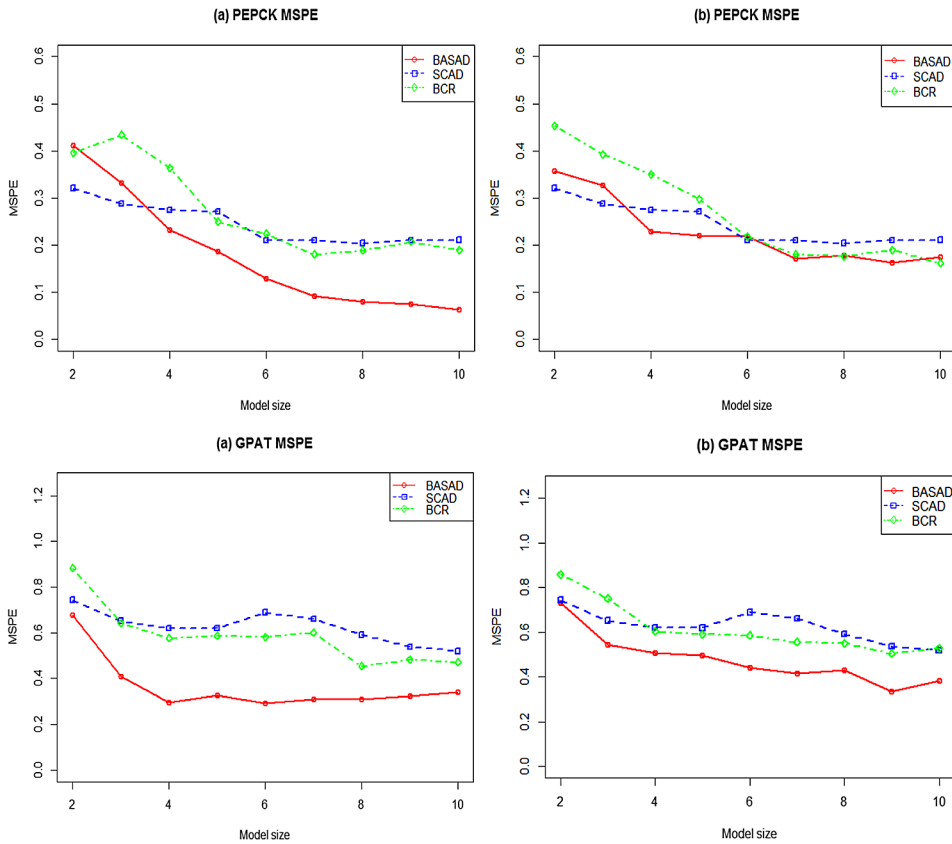


FIG. 1. Mean squared prediction error (MSPE) versus model size for analyzing PEPCK and GPAT in the upper and lower panel, respectively, (a) $p = 200$ and (b) $p = 400$.

10. Conclusion. In this paper, we consider a Bayesian variable selection method for high dimensional data based on the spike and slab priors with shrinking and diffusing priors. We show under mild conditions that this approach achieves strong selection consistency in the sense that the posterior probability of the true model converges to one. The tuning parameters needed for the prior specifications are transparent, and a standard Gibbs sampler can be used for posterior sampling. We also provide the asymptotic relationship between the proposed approach and the L_0 penalty for model selection. Simulation studies in Section 8 and real data example in Section 9 show evidence that the method performs well in a variety of settings even though we do not attempt to optimize the tuning parameters in the proposed method.

The strong selection consistency of Bayesian methods has not been established in the cases of $p > n$ until very recently. For higher dimensional cases, we just became aware of Liang, Song and Yu (2013), which provided the strong selection consistency for Bayesian subset selection based on the theory developed by Jiang

(2007) for posterior density consistency. However, to translate density consistency into selection consistency, Liang, Song and Yu (2013) imposed a condition on the posterior distribution itself, which is not verifiable directly. The techniques we use in this paper might also be used to complete the development of their theory on strong selection consistency.

Throughout the paper, we assume Gaussian errors in the regression model, but this assumption is not necessary to obtain selection consistency. For proving Lemma 4.1, we did not need assumptions on the error distribution, and to prove Theorem 4.2, we just need deviation inequalities of the quadratic forms $\varepsilon' P_k \varepsilon$, which follow the chi-squared distribution for normal errors. Similar proofs with an application of deviation inequalities for other error distributions would work. For instance, Hsu, Kakade and Zhang (2012) provide deviation inequalities for quadratic forms of sub-Gaussian random variables.

The primary focus of our paper is model selection consistency. The model is selected by averaging over the latent indicator variables drawn from the posterior distributions. The strengths of different model selection methods need to be evaluated differently if prediction accuracy is the goal. In our empirical work, we have included comparisons of the mean squared prediction errors, and found that our proposed method based on default tuning parameters is highly competitive in terms of prediction. However, improvements are possible, mainly in the cases of low signals, if the parameters are tuned by BIC or cross-validation, or if model-averaging is used instead of the predictions from a single model.

11. Proofs. In this section, we prove Lemmas 4.1 and 6.1. Please refer to Narisetty and He (2014) for proofs of the remaining results.

PROOF OF LEMMA 4.1. The joint posterior of β, σ^2, Z under model (1) is given by

$$\begin{aligned}
 & P(\beta, Z = k, \sigma^2 | Y) \\
 (9) \quad & \propto \exp \left\{ -\frac{1}{2\sigma^2} (\|Y - X\beta\|_2^2 - \beta' D_k \beta - 2\alpha_2) \right\} \sigma^{-2(n/2 + p_n/2 + \alpha_1 + 1)} \\
 & \quad \times |D_k|^{1/2} s_n^{|k|},
 \end{aligned}$$

where $D_k = \text{Diag}(k\tau_{1n}^{-2} + (\mathbf{1} - k)\tau_{0n}^{-2})$, $s_n = q_n / (1 - q_n)$, α_1, α_2 are the parameters of IG prior, and $|k|$ is the size of the model k . By a simple rearrangement of terms in the above expression, we obtain

$$\begin{aligned}
 & P(\beta, Z = k | Y, \sigma^2) \\
 & \propto \exp \left\{ -\frac{1}{2\sigma^2} ((\beta - \tilde{\beta})'(D_k + X'X)(\beta - \tilde{\beta}) - \tilde{\beta}'(D_k + X'X)\tilde{\beta}) \right\} |D_k|^{1/2} s_n^{|k|},
 \end{aligned}$$

where $\tilde{\beta} = (D_k + X'X)^{-1} X'Y$. Note that $\tilde{\beta}$ is a shrinkage estimator of the regression vector β . Shrinkage of $\tilde{\beta}$ depends on D_k , which is the precision matrix of β

given $Z = k$. The components of $\tilde{\beta}_i$ corresponding to $k_i = 0$ are shrunk towards zero while the shrinkage of coefficients corresponding to $k_i = 1$ is negligible (as τ_{1n}^{-2} is small).

$$\begin{aligned}
 P(Z = k|Y, \sigma^2) &\propto Q_k s_n^{|k|} \exp\left\{-\frac{1}{2\sigma^2}(Y'Y - \tilde{\beta}'(D_k + X'X)\tilde{\beta})\right\} \\
 (10) \qquad &= Q_k s_n^{|k|} \exp\left\{-\frac{1}{2\sigma^2}(Y'Y - Y'X(D_k + X'X)^{-1}X'Y)\right\} \\
 &= Q_k s_n^{|k|} \exp\left\{-\frac{1}{2\sigma^2}\tilde{R}_k\right\},
 \end{aligned}$$

where $Q_k = |D_k + X'X|^{-1/2}|D_k|^{1/2}$. Next, we obtain bounds on Q_k .

LEMMA 11.1. *Let A be an invertible matrix, and B be any matrix with appropriate dimension. Further, let k and j be any pair of models. Then,*

- (i) $|(A + B'B)^{-1}A| = |I + BA^{-1}B'|^{-1}$,
- (ii) $(I + \tau_{1n}^2 X_k X_k' + \tau_{0n}^2 X_j X_j')^{-1} \geq (I + \tau_{1n}^2 X_k X_k')^{-1}(1 - \xi_n)$,
 where $\xi_n = n\tau_{0n}^2 \lambda_M^n = o(1)$, and
- (iii) $Q_k \leq w'(n\tau_{1n}^2 \lambda_m^n (1 - \phi_n))^{-(1/2)(r_k^* - r_l)} (\lambda_m^n)^{-(1/2)|t \wedge k^c|} Q_l$, where $w' > 0$, $r_k = \text{rank}(X_k)$, $r_k^* = r_k \wedge m_n$, and $\phi_n = o(1)$.

PROOF. (i) We use the Sylvester’s determinant theorem, and the multiplicative property of the determinant to obtain

$$\begin{aligned}
 |(A + B'B)^{-1}A| &= |I + A^{-1/2}B'BA^{-1/2}|^{-1} \\
 &= |I + BA^{-1}B'|^{-1}.
 \end{aligned}$$

(ii) By the Sherman–Morrison–Woodbury (SMW) identity, assuming A , C and $(C^{-1} + DA^{-1}B)$ to be nonsingular,

$$(11) \quad (A + BCD)^{-1} = A^{-1} - A^{-1}B(C^{-1} + DA^{-1}B)^{-1}DA^{-1},$$

we have, for any vector a ,

$$a'(I + \tau_{1n}^2 X_k X_k' + \tau_{0n}^2 X_j X_j')^{-1}a = a'G^{-1}a - \tau_{0n}^2 H,$$

where $G = I + \tau_{1n}^2 X_k X_k'$ and $H = a'G^{-1}X_j(I + \tau_{0n}^2 X_j'G^{-1}X_j)^{-1}X_j'G^{-1}a$. Note that

$$\begin{aligned}
 (12) \quad 0 \leq \tau_{0n}^2 H &\leq \tau_{0n}^2 a'G^{-1}X_j X_j'G^{-1}a \\
 &\leq n\tau_{0n}^2 \lambda_M^n a'G^{-1}a,
 \end{aligned}$$

where λ_M^n is the maximum eigenvalue of the Gram matrix $X'X/n$. Therefore,

$$a'(I + \tau_{1n}^2 X_k X_k')^{-1}a(1 - n\tau_{0n}^2 \lambda_M^n) \leq a'(I + \tau_{1n}^2 X_k X_k' + \tau_{0n}^2 X_j X_j')^{-1}a,$$

and hence (ii) is proved.

(iii) From part (i) of the lemma, we have

$$\begin{aligned}
 (13) \quad Q_k &= |I + XD_k^{-1}X'|^{-1/2} \\
 &= |I + \tau_{1n}^2 X_k X'_k + \tau_{0n}^2 X_{k^c} X'_{k^c}|^{-1/2}.
 \end{aligned}$$

Define $A = I + \tau_{1n}^2 X_{k \wedge t} X'_{k \wedge t} + \tau_{0n}^2 X_{k^c \vee t^c} X'_{k^c \vee t^c}$. Then, by (ii) we have

$$(1 - \xi_n)(I + \tau_{1n}^2 X_{k \wedge t} X'^{-1}_{k \wedge t}) \leq A^{-1} \leq (I + \tau_{1n}^2 X_{k \wedge t} X'_{k \wedge t})^{-1}.$$

This, along with Condition 4.5 implies

$$\begin{aligned}
 \frac{Q_k}{Q_{k \wedge t}} &= |I + \tau_{1n}^2 X_k X'_k + \tau_{0n}^2 X_{k^c} X'_{k^c}|^{-1/2} |A|^{1/2} \\
 &= |A + (\tau_{1n}^2 - \tau_{0n}^2) X_{k \wedge t^c} X'_{k \wedge t^c}|^{-1/2} |A|^{1/2} \\
 &= |I + (\tau_{1n}^2 - \tau_{0n}^2) X'_{k \wedge t^c} A^{-1} X_{k \wedge t^c}|^{-1/2} \\
 &\leq |I + (\tau_{1n}^2 - \tau_{0n}^2)(1 - \xi_n) X'_{k \wedge t^c} (I + \tau_{1n}^2 X_{k \wedge t} X'_{k \wedge t})^{-1} X_{k \wedge t^c}|^{-1/2} \\
 &= |I + \tau_{1n}^2 X_t X'_t + \tau_{1n}^2 (1 - \phi_n) X_{k \wedge t^c} X'_{k \wedge t^c}|^{-1/2} |I + \tau_{1n}^2 X_{k \wedge t} X'_{k \wedge t}|^{1/2} \\
 &\leq |I + \tau_{1n}^2 (1 - \phi_n) X_k X'_k|^{-1/2} |I + \tau_{1n}^2 X_{k \wedge t} X'_{k \wedge t}|^{1/2} \\
 &\leq (n \tau_{1n}^2 \lambda_m^n (1 - \phi_n))^{-(r_k^* - r_{t \wedge k})/2} (1 - \phi_n)^{-|t \wedge k|/2},
 \end{aligned}$$

where $(1 - \phi_n) = (\tau_{1n}^2 - \tau_{0n}^2)(1 - \xi_n)/\tau_{1n}^2 \rightarrow 1$. Similarly, let $A = I + \tau_{1n}^2 X_t X'_t + \tau_{0n}^2 X_{t^c} X'_{t^c}$ to obtain

$$\begin{aligned}
 \frac{Q_{k \wedge t}}{Q_t} &= |A - (\tau_{1n}^2 - \tau_{0n}^2) X_{k \wedge t^c} X'_{k \wedge t^c}|^{-1/2} |A|^{1/2} \\
 &\leq |I + \tau_{1n}^2 X_{k \wedge t} X'_{k \wedge t}|^{-1/2} |I + \tau_{1n}^2 X_t X'_t|^{1/2} \\
 &\leq |I + \tau_{1n}^2 X_{t \wedge k^c} X'_{t \wedge k^c}|^{1/2} \\
 &\leq (n \tau_{1n}^2 c')^{|t \wedge k^c|/2}.
 \end{aligned}$$

The above two inequalities give

$$\frac{Q_k}{Q_t} \leq w'(n \tau_{1n}^2 \lambda_m^n (1 - \phi_n))^{-(r_k^* - r_t)/2} (\lambda_m^n)^{-|t \wedge k^c|/2}.$$

□

Due to (10), we have

$$PR(k, t) = \frac{Q_k}{Q_t} s_n^{|k| - |t|} \exp\left\{-\frac{1}{2\sigma^2} (\tilde{R}_k - \tilde{R}_t)\right\}.$$

Therefore, Lemma 11.1(iii) implies Lemma 4.1. □

PROOF OF LEMMA 6.1. The rows of X_k are n independent sub-Gaussian random isotropic vectors in $R^{|k|}$. Note that $|k| \leq m_n$ implies $|k| = o(n)$. Due to Theorem 5.39 of Vershynin (2012), with probability at least $1 - 2 \exp(-cs)$, we have

$$(14) \quad \phi_{\min}\left(\frac{X'_k X_k}{n}\right) > \left(1 - C\sqrt{\frac{|k|}{n}} - \sqrt{\frac{s}{n}}\right)^2,$$

where c and C are absolute constants that depend only on the sub-Gaussian norms of the rows of the matrix X_k .

Let us fix $s = n(1 - \phi)$ for some $\phi > 0$, and define the event given by equation (14) as A_k . We then have $P[A_k^c] < 2 \exp(-c(1 - \phi)n)$ for all k . By taking an union bound over $\{k : |k| \leq m_n\}$, we obtain

$$\begin{aligned} P\left[\bigcup_{|k| \leq m_n} A_k^c\right] &\leq p_n^{m_n} \exp(-c(1 - \phi)n) \\ &= \exp\left\{\frac{n}{2 + \nu} - c(1 - \phi)n\right\} \rightarrow 0, \end{aligned}$$

if $\nu > (\frac{1}{c(1-\phi)} - 2)$. Therefore, in the event $\bigcap_{|k| \leq m_n} A_k$, whose probability goes to 1, we have $\phi_{\min}(X'_k X_k/n) \geq \phi^2/4 - O(\sqrt{m_n/n}) > 0$, for all k . \square

Acknowledgments. The authors are grateful to anonymous referees and an Associate Editor for their encouraging and helpful comments on an earlier version of the paper. The authors would also like to thank Professors Howard Bondell, Val Johnson and Faming Liang for sharing with us their code to perform Bayesian model selection.

SUPPLEMENTARY MATERIAL

Supplement to “Bayesian variable selection with shrinking and diffusing priors” (DOI: [10.1214/14-AOS1207SUPP](https://doi.org/10.1214/14-AOS1207SUPP); .pdf). This supplement contains the proofs of Theorems 4.1, 4.2 and Lemma 4.2.

REFERENCES

BARBIERI, M. M. and BERGER, J. O. (2004). Optimal predictive model selection. *Ann. Statist.* **32** 870–897. [MR2065192](#)

BONDELL, H. D. and REICH, B. J. (2008). Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with OSCAR. *Biometrics* **64** 115–123, 322–323. [MR2422825](#)

BONDELL, H. D. and REICH, B. J. (2012). Consistent high-dimensional Bayesian variable selection via penalized credible regions. *J. Amer. Statist. Assoc.* **107** 1610–1624. [MR3036420](#)

CANDES, E. and TAO, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n . *Ann. Statist.* **35** 2313–2351. [MR2382644](#)

- DEY, T., ISHWARAN, H. and RAO, J. S. (2008). An in-depth look at highest posterior model selection. *Econometric Theory* **24** 377–403. [MR2391616](#)
- DICKER, L., HUANG, B. and LIN, X. (2013). Variable selection and estimation with the seamless- L_0 penalty. *Statist. Sinica* **23** 929–962. [MR3086661](#)
- FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. [MR1946581](#)
- FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **70** 849–911. [MR2530322](#)
- FAN, J. and LV, J. (2010). A selective overview of variable selection in high dimensional feature space. *Statist. Sinica* **20** 101–148. [MR2640659](#)
- FAN, J. and PENG, H. (2004). Nonconcave penalized likelihood with a diverging number of parameters. *Ann. Statist.* **32** 928–961. [MR2065194](#)
- GEORGE, E. I. and FOSTER, D. P. (2000). Calibration and empirical Bayes variable selection. *Biometrika* **87** 731–747. [MR1813972](#)
- GEORGE, E. I. and MCCULLOCH, R. E. (1993). Variable selection via Gibbs sampling. *J. Amer. Statist. Assoc.* **88** 881–889.
- HSU, D., KAKADE, S. M. and ZHANG, T. (2012). A tail inequality for quadratic forms of subgaussian random vectors. *Electron. Commun. Probab.* **17** 6. [MR2994877](#)
- HUANG, J. and XIE, H. (2007). Asymptotic oracle properties of SCAD-penalized least squares estimators. In *Asymptotics: Particles, Processes and Inverse Problems. Institute of Mathematical Statistics Lecture Notes—Monograph Series* **55** 149–166. IMS, Beachwood, OH. [MR2459937](#)
- ISHWARAN, H., KOGALUR, U. B. and RAO, J. S. (2010). spikeslab: Prediction and variable selection using spike and slab regression. *The R Journal* **2** 68–73.
- ISHWARAN, H. and RAO, J. S. (2005). Spike and slab variable selection: Frequentist and Bayesian strategies. *Ann. Statist.* **33** 730–773. [MR2163158](#)
- ISHWARAN, H. and RAO, J. S. (2011). Consistency of spike and slab regression. *Statist. Probab. Lett.* **81** 1920–1928. [MR2845909](#)
- JAMES, G. M., RADCHENKO, P. and LV, J. (2009). DASSO: Connections between the Dantzig selector and lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **71** 127–142. [MR2655526](#)
- JIANG, W. (2007). Bayesian variable selection for high dimensional generalized linear models: Convergence rates of the fitted densities. *Ann. Statist.* **35** 1487–1511. [MR2351094](#)
- JOHNSON, V. E. and ROSSELL, D. (2012). Bayesian model selection in high-dimensional settings. *J. Amer. Statist. Assoc.* **107** 649–660. [MR2980074](#)
- KIM, Y., KWON, S. and CHOI, H. (2012). Consistent model selection criteria on high dimensions. *J. Mach. Learn. Res.* **13** 1037–1057. [MR2930632](#)
- LAN, H., CHEN, M., FLOWERS, J. B., YANDELL, B. S., STAPLETON, D. S., MATA, C. M., MUI, E. T., FLOWERS, M. T., SCHUELER, K. L., MANLY, K. F., WILLIAMS, R. W., KENDZIORSKI, K. and ATTIE, A. D. (2006). Combined expression trait correlations and expression quantitative trait locus mapping. *PLoS Genetics* **2** e6.
- LIANG, F., SONG, Q. and YU, K. (2013). Bayesian subset modeling for high-dimensional generalized linear models. *J. Amer. Statist. Assoc.* **108** 589–606. [MR3174644](#)
- LIU, Y. and WU, Y. (2007). Variable selection via a combination of the L_0 and L_1 penalties. *J. Comput. Graph. Statist.* **16** 782–798. [MR2412482](#)
- MITCHELL, T. J. and BEAUCHAMP, J. J. (1988). Bayesian variable selection in linear regression. *J. Amer. Statist. Assoc.* **83** 1023–1036. [MR0997578](#)
- MORENO, E., GIRÓN, F. J. and CASELLA, G. (2010). Consistency of objective Bayes factors as the model dimension grows. *Ann. Statist.* **38** 1937–1952. [MR2676879](#)
- NARISSETTY, N. N. and HE, X. (2014). Supplement to “Bayesian variable selection with shrinking and diffusing priors.” DOI:10.1214/14-AOS1207SUPP.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Ann. Statist.* **6** 461–464. [MR0468014](#)

- SHEN, X., PAN, W. and ZHU, Y. (2012). Likelihood-based selection and sharp parameter estimation. *J. Amer. Statist. Assoc.* **107** 223–232. [MR2949354](#)
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **58** 267–288. [MR1379242](#)
- VERSHYNIN, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing* 210–268. Cambridge Univ. Press, Cambridge. [MR2963170](#)
- YANG, Y. and HE, X. (2012). Bayesian empirical likelihood for quantile regression. *Ann. Statist.* **40** 1102–1131. [MR2985945](#)
- YUAN, M. and LIN, Y. (2005). Efficient empirical Bayes variable selection and estimation in linear models. *J. Amer. Statist. Assoc.* **100** 1215–1225. [MR2236436](#)
- ZHANG, D., LIN, Y. and ZHANG, M. (2009). Penalized orthogonal-components regression for large p small n data. *Electron. J. Stat.* **3** 781–796. [MR2534201](#)
- ZOU, H. (2006). The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.* **101** 1418–1429. [MR2279469](#)

DEPARTMENT OF STATISTICS
UNIVERSITY OF MICHIGAN
ANN ARBOR, MICHIGAN 48109
USA
E-MAIL: naveennn@umich.edu
xmhe@umich.edu