# BAYESIAN VARIABLE SELECTION FOR HIGH DIMENSIONAL GENERALIZED LINEAR MODELS: CONVERGENCE RATES OF THE FITTED DENSITIES 

By Wenxin Jiang

Northwestern University


#### Abstract

Bayesian variable selection has gained much empirical success recently in a variety of applications when the number $K$ of explanatory variables $\left(x_{1}, \ldots, x_{K}\right)$ is possibly much larger than the sample size $n$. For generalized linear models, if most of the $x_{j}$ 's have very small effects on the response $y$, we show that it is possible to use Bayesian variable selection to reduce overfitting caused by the curse of dimensionality $K \gg n$. In this approach a suitable prior can be used to choose a few out of the many $x_{j}$ 's to model $y$, so that the posterior will propose probability densities $p$ that are "often close" to the true density $p^{*}$ in some sense. The closeness can be described by a Hellinger distance between $p$ and $p^{*}$ that scales at a power very close to $n^{-1 / 2}$, which is the "finite-dimensional rate" corresponding to a low-dimensional situation. These findings extend some recent work of Jiang [Technical Report 05-02 (2005) Dept. Statistics, Northwestern Univ.] on consistency of Bayesian variable selection for binary classification.


1. Introduction. Bayesian variable selection (BVS) is a fruitful method for studying regression models that relate a response $y$ to a vector of candidate explanatory variables $x=\left(x_{1}, \ldots, x_{K}\right)^{T}$. For example, when generalized linear models (GLM) are considered, the density of $y$ and the mean function of $y$ conditional on $x$ both depend on a linear combination $x^{T} \beta$ through the regression coefficients $\beta=\left(\beta_{1}, \ldots, \beta_{K}\right)^{T}$. The BVS approach uses priors that propose different model $\gamma$ 's and the corresponding sets of regression coefficient $\beta_{\gamma}$ 's, where $\gamma$ indicates the components of $x$ that are included in regression. The posterior distribution $\pi\left[\gamma, \beta_{\gamma} \mid D^{n}\right]$ for the model and the model parameters ( $\gamma, \beta_{\gamma}$ ) can then be obtained based on an observed data set $D^{n}=\left(x^{(i)}, y^{(i)}\right)_{1}^{n}$, which is often assumed to consist of i.i.d. (independent and identically distributed) copies of $(x, y)$. Computational simplification is achievable in the cases of linear regression and probit regression, where the unknown regression coefficients $\beta_{\gamma}$ can often be analytically integrated out in the posterior-based computations (e.g., Kohn, Smith and Chan [17]; Lee et al. [19]).
[^0]The BVS approach has had many successful applications. For example, when applied in a linear regression framework, BVS is used in basis selection for nonparametric regression (e.g., Smith and Kohn [23], Kohn, Smith and Chan [17]) and in construction of financial index tracking portfolios (e.g., George and McCulloch [7]). Other work applying BVS in the GLM framework includes, for example, Clyde and DeSimone-Sasinowska [3], Nott and Leonte [21] and Wang and George [24]. Recently, BVS has been applied to the area of bioinformatics. In order to construct Gaussian graphical models for gene expression pathways, Dobra et al. [5] obtain biologically meaningful results by applying Bayesian variable selection to model how each gene in the graph relates to tens of thousands of other genes. In order to classify binary responses based on microarray data, Lee et al. [19] and Sha et al. [22] (via probit regression) and Zhou, Liu and Wong [27] (via logistic regression) use BVS to achieve excellent cross-validated classification errors. These most recent applications are especially noteworthy since they are all in the situation of $K \gg n$, where the number of candidate variables $K$ can be several thousand and the sample size $n$ is often less than a hundred.

Despite these empirical successes, there has not been a systematic study of the frequentist properties of BVS, such as posterior consistency and convergence rates. It is the aim of this paper to study these convergence properties for BVS, allowing $K$ to be possibly much larger than $n$. The consistency that we will consider is neither the traditional sense (i) of consistency in estimating the true regression parameters, nor the sense (ii) of consistency in identifying the true model (the $x$-components with nonzero regression coefficients). Sense (i) is not feasible since in cases with $K \gg n$, the $\beta$-coefficients are often not identifiable. Sense (ii) is not a totally satisfactory framework when, as in many realistic situations, none of the $K$ regression coefficients is exactly zero, even though many of them may be very small. The consistency we consider is the closeness between the true (conditional) density $p^{*}=p^{*}(y \mid x)$ and the densities $p=p\left(y \mid x ; \gamma, \beta_{\gamma}\right)$ proposed by the posterior $\pi\left(\gamma, \beta_{\gamma} \mid D^{n}\right)$. We do not attempt to identify the "true parameter" or the "true model" (the nonzero coefficients). Rather, we allow all coefficients to be not exactly zero, and attempt to construct the posterior to propose models that include only a few of those nonzero coefficients, but have the corresponding densities $p$ "often close" to the true $p^{*}$ in some sense.

Let $v_{x}(d x)$ be the probability measure for $x$ and $v_{y}(d y)$ be the dominating measure for conditional densities $p$ and $p^{*}$. Define the Hellinger distance between $p$ and $p^{*}$ as $d\left(p, p^{*}\right)=\sqrt{\iint \nu_{y}(d y) v_{x}(d x)\left(\sqrt{p}-\sqrt{p^{*}}\right)^{2}}$. The convergence results we consider describe the "often closeness" between $p$ and $p^{*}$ that can be formulated as, for example,

$$
\begin{equation*}
P^{*}\left[\pi\left[d\left(p, p^{*}\right)<\varepsilon_{n} \mid D^{n}\right]>1-\delta_{n}\right] \geq 1-\lambda_{n}, \tag{1}
\end{equation*}
$$

for all large enough $n$, for some small $\varepsilon_{n}, \delta_{n}, \lambda_{n}$ converging to zero as $n \rightarrow \infty$, where $P^{*}$ is the probability measure for data $D^{n}$, when they are generated as i.i.d. copies from $p^{*} v_{y}(d y) v_{x}(d x)$.

It is noted that BVS is essential for achieving convergence results as above, when $K \gg n$. A usual approach without variable selection, using the full model and putting a prior on all the regression coefficients, can be shown to lead to bad results in the following counterexample.

Example. Suppose $K>n$. Let the random variable $z$ take values from $\{j / K\}_{j=1}^{n}$ with equal probability and let $x=\left(x_{1}, \ldots, x_{K}\right)^{T}$, where $x_{j}=I[z=$ $j / K$ ] for each $j$. Let $\left(z^{(i)}, x^{(i)}, y^{(i)}\right)_{i=1}^{n}$ and $(z, x, y)$ be i.i.d., where $y \mid x \sim$ $N(0,1)$. Suppose the fitted model is $y \mid x \sim N\left(\sum_{j=1}^{K} \beta_{j} x_{j}, 1\right)$, where, without selecting among the $\left(x_{j}\right)_{1}^{K}$, one proposes a prior for $\left(\beta_{j}\right)_{1}^{K}$ as i.i.d. $N(0,1)$.

The Hellinger distance $d\left(p, p^{*}\right)$ between $p^{*}=e^{-y^{2} / 2} / \sqrt{2 \pi}$ and $p=e^{-\left(y-\sum_{j=1}^{K} \beta_{j} x_{j}\right)^{2} / 2} / \sqrt{2 \pi}$ is such that $d\left(p, p^{*}\right)^{2}=(2 / K) \sum_{j=1}^{K}\left(1-e^{-\beta_{j}^{2} / 8}\right)$.

Then, in the posterior conditional on $D^{n}=\left(x^{(i)}, y^{(i)}\right)_{i=1}^{n}, \beta_{1}, \ldots, \beta_{K}$ are independent and $\beta_{j} \sim N\left(\sum_{i=1}^{n} y^{(i)} x_{j}^{(i)} /\left(1+\sum_{i=1}^{n} x_{j}^{(i)}\right), 1 /\left(1+\sum_{i=1}^{n} x_{j}^{(i)}\right)\right)$. Note that $\left(x_{j}^{(i)}\right)_{i=1}^{n}$ are zero for at least $K-n$ of the $K j$ 's since $x_{j}^{(i)}=I\left[z^{(i)}=\right.$ $j / K]$, and the $n z^{(i)}$, s can only populate at most $n$ out of $K$ of the $j / K$ locations. Therefore, at least $K-n$ out of the $K \beta_{j}$ 's follow the $N(0,1)$ distribution in the posterior-which is the same as the corresponding prior distribution. Without loss of generality, let $\beta_{1}^{K-n}$ be independent $N(0,1)$ in the posterior. Note that $d\left(p, p^{*}\right)^{2} \geq[2(K-n) / K][1 /(K-n)] \sum_{1}^{K-n}\left(1-e^{-\beta_{j}^{2} / 8}\right)$. For a simple example, let $K=2 n$. An application of Chebyshev's inequality leads to $\pi\left[d\left(p, p^{*}\right) \geq \eta^{1 / 2} \mid D^{n}\right] \geq 1-\left(\eta^{2} n\right)^{-1}$, for $\eta=1 / 2-1 / \sqrt{5}$, which happens with $P^{*}$-probability 1. Therefore, without variable selection, a convergence result such as (1) cannot hold. Such a convergence result, however, can be shown to hold for this example, with $\varepsilon_{n}$ following a near finite-dimensional rate (a power close to $1 / \sqrt{n}$ ), if Bayesian variable selection is used properly, according to later results of this paper (e.g., Remark 1).

There has been considerable interest recently in studying the theoretical properties of high-dimensional regression. Most results are for frequentist methods. For example, Bühlmann [2] considers boosting for high-dimensional regression; Greenshtein and Ritov [12] and Greenshtein [11] consider constrained or $\ell_{1}$-penalized optimization; Meinshausen and Bühlmann [20] apply a similar method of $\ell_{1}$ penalization to high-dimensional graphical models. Recently, Fan and Li [6] have provided a useful overview for methods based on the penalized likelihood for treating high dimensionality, which includes examples of generalized linear models and survival models, among others. In contrast to these frequentist approaches based on optimization, the Bayesian method considered here has the attractive capability of presenting several likely models together with the corresponding posterior probabilities. A theoretical study of Bayesian inference without
variable selection has been carried out by, for example, Ghosal [8, 9]. This work considers $K$ 's growing with $n$ but at a slower rate. On the other hand, in the $K \ll n$ case treated in Ghosal [8, 9], posterior asymptotic normality was established for the whole parameter vector, so the goal there was much higher, and hence, the result there is not comparable with the result in the present paper which focuses on posterior convergence rates.

In contrast to previous work, we consider Bayesian variable selection and allow the cases $K \gg n$. It is noted that it is essential to have the variable selection step in order to obtain good results when $K \gg n$. The counterexample above shows that, without variable selection, it is impossible to have good convergence in general cases with $K>n$, while with variable selection, excellent empirical performance has been reported, for example, in Lee et al. [19] and Sha et al. [22] with $K \gg n$.

We study the convergence behavior of BVS for generalized linear models, which include linear regression, logistic regression, probit regression, Poisson regression, and so on. We also include a discussion of Gaussian graphical models that uses linear regression for neighborhood selection. Therefore, the current paper forms an extension to Jiang [15], who only considers consistency of BVS for binary logistic and probit regression, without studying the convergence rates. Here we study the convergence rate $\varepsilon_{n}$ as well, and will show that despite the highdimension $K \gg n$, BVS can still lead to a near finite-dimensional rate (with $\varepsilon_{n}$ close to $1 / \sqrt{n}$ in order), if we are in some "sparse" situations when most of the regression coefficients are very small. (For binary regression, this rate $\varepsilon_{n}$ also forms a good convergence rate for the purpose of classification, as shown in Section 5 later.) For such sparse high-dimensional problems, Bayesian variable selection can therefore help to reduce "overfitting" or the "curse of dimensionality." Note that such a conclusion can only be drawn by a careful study of the convergence rates in high dimensions; just proving the consistency, as in Jiang [15], is not enough; for example, it is well known (e.g., Hastie, Tibshirani and Friedman [13], Chapter 13) that the $k$-nearest neighbor rules are consistent for classification, but can suffer considerably from the curse of dimensionality. Also, it is well known (e.g., Devroye, Györfi and Lugosi [4], Chapters 6 and 7) that (at least in finite dimensions) there exist universally consistent classification rules, but any rule can have a very slow convergence rate under some data distribution.

Below we will first specify the notation and the framework of the paper.
2. Notation and framework. The explanatory variable is a $K_{n}$-dimensional random vector $x=\left(x_{1}, x_{2}, \ldots, x_{K_{n}}\right)^{T}$. Following the typical practice of studying high-dimensional problems, we will formally consider the asymptotics when $K_{n}$ increases as $n \rightarrow \infty$.

For simplicity, we will assume that $\left|x_{j}\right| \leq 1$ for all $j$ for most of the later discussion. The results can be easily extended to the case when all $\left|x_{j}\right|$ 's are bounded above by a large constant.

The response is $y$. The true relation between $y$ and $x$ is assumed to follow a parametric generalized linear model (GLM) with true conditional density $p^{*}(y \mid x)$ and the corresponding mean function $\mu^{*}(x)$. Generalized linear models (GLM) are a class of popular regression models relating a response $y$ to a vector of covariates $x$. The GLM with one natural parameter is constructed with a density of the form $p^{*}(y \mid x)=\exp \left\{a\left(h^{*}\right) y+b\left(h^{*}\right)+c(y)\right\} \equiv f\left(y, h^{*}\right)$, where $h^{*}=x^{T} \beta^{*}$ is the linear parameter, $a(h)$ and $b(h)$ are continuously differentiable, and $a(h)$ has nonzero derivative. The mean function $\mu^{*}=E(y \mid x)=-b^{\prime}\left(h^{*}\right) / a^{\prime}\left(h^{*}\right) \equiv$ $\psi\left(x^{T} \beta^{*}\right)$ follows a transformed linear model, where the transform $\psi$ is the inverse of a chosen link function. This formalism includes regression models for responses that are binary, Poisson and Gaussian (with known error variance), and can be easily extended to the cases with a dispersion parameter, which can then include Gaussian models with unknown error variance.

We assume that corresponding to the true model $p^{*}$, there exists a true regression parameter vector $\beta^{*}$, which satisfies some "sparseness" conditions, describing situations when most components of $\beta^{*}$ are very small in magnitude. One such condition states that $\lim _{n \rightarrow \infty} \sum_{j=1}^{K_{n}}\left|\beta_{j}^{*}\right|<\infty$. Other conditions can be formulated to describe how fast the sum of $\left|\beta_{j}^{*}\right|$ 's converges.

The condition limiting the sum of $\left|\beta_{j}^{*}\right|$ 's has been considered by Bühlmann [2] for studying how boosting algorithms handle high-dimensional linear regression. As Bühlmann points out, as a special case, this condition is satisfied when only a finite and fixed number of $x_{j}$ 's are relevant, that is, when the number of nonzero $\beta_{j}^{*}$ 's is independent of $n$. More generally, the sparseness conditions can describe situations when all $x_{j}$ 's are relevant, but most of them have very small effects (| $\beta_{j}^{*} \mid$ 's).

Note that $p^{*}$ is the conditional density of $y \mid x$, which is also the joint density of $(x, y)$ if the dominating measure $v_{x}(d x) v_{y}(d y)$ is the product of the probability measure of $x$ and the dominating measure of $y$. We will always use this kind of dominating measure.

The data for $n$ subjects are assumed to be independent and identically distributed (i.i.d.) based on $p^{*} v_{x}(d x) v_{y}(d y)$. Therefore, showing the subject index $i$, the data set is of the form $D^{n}=\left(x_{1}^{(i)}, \ldots, x_{K_{n}}^{(i)}, y^{(i)}\right)_{i=1}^{n}$.

The prior selects a subset of the $K_{n} x$-variables in the data set to model $y$, using density $f\left(y, x_{\gamma}^{T} \beta_{\gamma}\right)$, where $\gamma=\left(\gamma_{1}, \ldots, \gamma_{K_{n}}\right)$ has $0 / 1$ valued components which are 1 only when the corresponding $x$ component is included in the model, that is, $\gamma_{j}=I\left[\left|\beta_{j}\right|>0\right]$. Sometimes we will also use $\gamma$ to denote the corresponding set of index $j$ 's for which $\left|\beta_{j}\right|>0$. The notation $v_{\gamma}$ denotes the subvector of a vector $v$ with components $\left\{v_{j}\right\}$, for all $j$ 's with $\gamma_{j}=1$ (or for all $j \in \gamma$, if $\gamma$ is understood as the corresponding index set).

We use the probability measure $\pi_{n}\left(\gamma, d \beta_{\gamma}\right)$ to denote the prior distribution of the subset model $\gamma$ and the corresponding regression coefficients $\beta_{\gamma}$. (The prior
depends on the sample size $n$, but we will often drop the subscript $n$ for simpler notation.) This induces a posterior measure conditional on the data set $D^{n}$,

$$
\begin{aligned}
& \pi\left(\gamma, d \beta_{\gamma} \mid D^{n}\right) \\
& \quad=\prod_{i=1}^{n} p\left(y_{i}, x_{i} \mid \gamma, \beta_{\gamma}\right) \pi\left(\gamma, d \beta_{\gamma}\right) / \sum_{\gamma^{\prime}} \int_{\beta_{\gamma^{\prime}}} \prod_{i=1}^{n} p\left(y_{i}, x_{i} \mid \gamma^{\prime}, \beta_{\gamma^{\prime}}^{\prime}\right) \pi\left(\gamma^{\prime}, d \beta_{\gamma^{\prime}}^{\prime}\right)
\end{aligned}
$$

where $p\left(y, x \mid \gamma, \beta_{\gamma}\right)=f\left(y, x_{\gamma}^{T} \beta_{\gamma}\right)$. The prior and posterior distributions for ( $\gamma, \beta_{\gamma}$ ) induce distributions for the corresponding parameterized densities.

For notational simplification, we will use $|v|$ to denote the sum of the absolute values of the components for any vector $v$. For two positive sequences $a_{n}$ and $b_{n}$, $a_{n} \prec b_{n}\left(\right.$ or $b_{n} \succ a_{n}$ ) means $\lim _{n \rightarrow \infty} a_{n} / b_{n}=0$.
3. A prior specification. General conditions on the prior will be given later in Section 7. Here, for being specific, we first consider the following prior for $\left(\gamma, \beta_{\gamma}\right)$. Conditional on $\gamma, \beta_{\gamma}$ follows $N\left(0, V_{\gamma}\right)$, where $V_{\gamma}$ is a $|\gamma| \times|\gamma|$ covariance matrix.

To complete this prior specification, we let the model indicators $\gamma=\left(\gamma_{1}, \ldots\right.$, $\gamma_{K_{n}}$ ) be generated by first proposing i.i.d. binary random variables $\tilde{\gamma}_{1}^{n}$, with $\pi\left(\tilde{\gamma}_{j}=1\right)=\lambda_{n}=r_{n} / K_{n}$, where we assume, for convenience, that $r_{n}$ is some integer smaller than $K_{n}$. We then keep only the $\tilde{\gamma}$ 's satisfying a size restriction $\sum\left|\tilde{\gamma}_{j}\right| \leq \bar{r}_{n}$, and let the prior proposed model $\gamma=\tilde{\gamma}$. Here $r_{n}$ is the prior expectation of model size $|\tilde{\gamma}|$ before applying the size restriction; $\bar{r}_{n}$ is the maximal possible model size. We assume $1 \leq r_{n} \leq \bar{r}_{n}<K_{n}$.

Therefore, $\pi(\gamma) \propto \prod_{j=1}^{K_{n}} \lambda_{n}^{\gamma_{j}}\left(1-\lambda_{n}\right)^{1-\gamma_{j}} I\left[\sum_{l=1}^{K_{n}} \gamma_{l} \leq \bar{r}_{n}\right]$. Although the size restriction is not necessary (see more general conditions in Section 7), it helps to keep the model from becoming too complicated and gives a convenient starting point for proving the theoretical properties. Also, without this kind of restriction, the design matrix $\sum_{i=1}^{n} x_{i \gamma} x_{i \gamma}^{T}$ would become singular when the proposed model size $|\gamma|>n$; such a design matrix is often used in the popular algorithms for generating the posterior distributions in Gaussian regression (e.g., Smith and Kohn [23]), probit regression (e.g., Lee et al. [19]) and logistic regression (e.g., Zhou et al. [27]).

Under this specification of prior, we will present conditions on $V_{\gamma}, r_{n}$ and $\bar{r}_{n}$ for proving results on posterior consistency and convergence rates. The condition on $V_{\gamma}$ will depend on how the largest eigenvalues $\left(\mathrm{ch}_{1}\right)$ of $V_{\gamma}$ and $V_{\gamma}^{-1}$ grow with the size of $|\gamma|$.

Let $H(\gamma)=\max \left\{\mathrm{ch}_{1}\left(V_{\gamma}\right), \mathrm{ch}_{1}\left(V_{\gamma}^{-1}\right)\right\}$. In many typical cases $H(\gamma)$ grows at most polynomially in model size, that is, $H(\gamma) \leq B|\gamma|^{v}$ for some constant $B>0$ and some power $v>0$, for all large $|\gamma|$. For example, when $V_{\gamma}=c I_{\gamma}$ (proportional to the identity matrix; see, e.g., Dobra et al. [5]), $H(\gamma)$ is a constant and does not grow with $|\gamma|$. For another example, $V_{\gamma} \approx$ constant $\times\left(E x_{\gamma} x_{\gamma}^{T}\right)^{-1}$; see, for example, Smith and Kohn [23] and Lee et al. [19], who use a sample approximation
of this choice. Then the largest eigenvalues of $V_{\gamma}$ and $V_{\gamma}^{-1}$ are both bounded linearly for large $|\gamma|$, when $x_{\gamma}$ has components standardized to have mean zero and common variance, and have all pairwise correlations being $\rho \in(0,1)$. In addition, for $V_{\gamma}$ following the covariance matrix of a finite-order AR or MA process, when the lag polynomials have no zeros on the unit circle, the eigenvalues of $V_{\gamma}$ and $V_{\gamma}^{-1}$ are also bounded such that $\max \left\{\operatorname{ch}_{1}\left(V_{\gamma}\right), \operatorname{ch}_{1}\left(V_{\gamma}^{-1}\right)\right\}$ grows like $|\gamma|^{0}$. For a detailed discussion of these eigenvalues, see, for example, Section 3 of Bickel and Levina [1].
4. Convergence results for GLM. Here, for simplicity we will assume that all explanatory variables are bounded and standardized such that $\left|x_{j}\right| \leq 1$ for all $j$. Assume $\lim _{n \rightarrow \infty} \sum_{1}^{K_{n}}\left|\beta_{j}^{*}\right|<\infty$ for a regression parameter $\beta^{*}$ corresponding to the true density $p^{*}$, where $K_{n}$ is a nondecreasing sequence in $n$.

We also assume that the prior specification in Section 3 is used. Define $\Delta\left(r_{n}\right)=\inf _{\gamma:|\gamma|=r_{n}} \sum_{j: j \notin \gamma}\left|\beta_{j}^{*}\right|, \quad B\left(r_{n}\right)=\sup _{\gamma:|\gamma|=r_{n}} \operatorname{ch}_{1}\left(V_{\gamma}^{-1}\right)$ and $\bar{B}\left(r_{n}\right)=$ $\sup _{\gamma:|\gamma|=r_{n}} \operatorname{ch}_{1}\left(V_{\gamma}\right)$. Let $\tilde{B}_{n}=\sup _{\gamma:|\gamma| \leq \bar{x}_{n}} \operatorname{ch}_{1}\left(V_{\gamma}\right)$. Let $D(R)=1+R \times$ $\sup _{|h| \leq R}\left|a^{\prime}(h)\right| \cdot \sup _{|h| \leq R}|\psi(h)|$ for any $R>0$.

THEOREM 1. Assume that the prior specification in Section 3 is used, $\left|x_{j}\right| \leq 1$ for all $j$ and $\lim _{n \rightarrow \infty} \sum_{1}^{K_{n}}\left|\beta_{j}^{*}\right|<\infty$, where $K_{n}$ is a nondecreasing sequence in $n$.

Let $\varepsilon_{n}$ be a sequence such that $\varepsilon_{n} \in(0,1]$ for each $n$ and $n \varepsilon_{n}^{2} \succ 1$ and assume that the following conditions also hold:

$$
\begin{equation*}
1 \leq r_{n} \leq \bar{r}_{n}<K_{n} \tag{5}
\end{equation*}
$$

$$
\begin{equation*}
1 \prec r_{n} \prec K_{n} \tag{6}
\end{equation*}
$$

$$
\begin{align*}
\bar{r}_{n} \ln \left(1 / \varepsilon_{n}^{2}\right) & \prec n \varepsilon_{n}^{2},  \tag{2}\\
\bar{r}_{n} \ln \left(K_{n}\right) & \prec n \varepsilon_{n}^{2},  \tag{3}\\
\bar{r}_{n} \ln D\left(\bar{r}_{n} \sqrt{n \varepsilon_{n}^{2} \tilde{B}_{n}}\right) & \prec n \varepsilon_{n}^{2}, \tag{4}
\end{align*}
$$

$$
\begin{equation*}
\Delta\left(r_{n}\right) \prec \varepsilon_{n}^{2} \tag{7}
\end{equation*}
$$

$$
\begin{equation*}
B\left(r_{n}\right) \prec n \varepsilon_{n}^{2}, \tag{8}
\end{equation*}
$$

$$
\begin{equation*}
r_{n} \ln \bar{B}\left(r_{n}\right) \prec n \varepsilon_{n}^{2} . \tag{9}
\end{equation*}
$$

Denote $d\left(p, p^{*}\right)^{2}=\iint\left|p\left(y, x \mid \gamma, \beta_{\gamma}\right)^{1 / 2}-p^{*}(y, x)^{1 / 2}\right|^{2} \nu_{y}(d y) v_{x}(d x)$. Then we have the following successively stronger results:
(i) for some $r_{0}>0$,

$$
\lim _{n \rightarrow \infty} P^{*}\left\{\pi\left[d\left(p, p^{*}\right) \leq \varepsilon_{n} \mid D^{n}\right] \geq 1-e^{-r_{0} n \varepsilon_{n}^{2}}\right\}=1
$$

(ii) for some $c_{1}>0$, and for all sufficiently large $n$,

$$
P^{*}\left\{\pi\left[d\left(p, p^{*}\right)>\varepsilon_{n} \mid D^{n}\right] \geq e^{-0.5 c_{1} n \varepsilon_{n}^{2}}\right\} \leq e^{-0.5 c_{1} n \varepsilon_{n}^{2}}
$$

(iii) for some $c_{1}>0$, and for all sufficiently large $n$,

$$
E_{D^{n}}^{*} \pi\left[d\left(p, p^{*}\right)>\varepsilon_{n} \mid D^{n}\right] \leq e^{-c_{1} n \varepsilon_{n}^{2}}
$$

The above condition on $D(\cdot)=D\left(\bar{r}_{n} \sqrt{n \varepsilon_{n}^{2} \tilde{B}_{n}}\right)$, when considering a specific example of GLM, depends on how $\left|a^{\prime}(h)\right|$ and $\left|\psi^{\prime}(h)\right|$ grow with the linear parameter $h$. We will consider the following examples here.
(a) Poisson regression with $\log$ linear link: mean $\mu=e^{h}, y \in\{0,1,2, \ldots\}$. Then

$$
f(y, h)=\frac{e^{-\mu}}{y!} \mu^{y}=\exp \left\{h y-e^{h}-\ln (y!)\right\}
$$

Here $a(h)=h, a^{\prime}=1, \psi(h)=e^{h}$. So both $\left|a^{\prime}\right|$ and $|\psi|$ grow at most exponentially in $|h|$.
(b) Normal linear regression: mean $\mu=h$; variance $\sigma^{2}=\varphi^{-1} \in \mathfrak{R}^{+}$is assumed to be known for now; $y \in \mathfrak{R}$. Then

$$
\begin{aligned}
f(y, h) & =\frac{1}{\sqrt{2 \pi \sigma^{2}}} e^{\left(-1 /\left(2 \sigma^{2}\right)\right)(y-\mu)^{2}} \\
& =\exp \left\{\varphi h y-\frac{\varphi h^{2}}{2}-\frac{\varphi y^{2}}{2}-\frac{1}{2} \ln \left(2 \pi \varphi^{-1}\right)\right\}
\end{aligned}
$$

Here $a(h)=\varphi h, a^{\prime}=\varphi, \psi(h)=h$. So both $\left|a^{\prime}\right|$ and $|\psi|$ grow at most linearly in $|h|$.
(c) Exponential regression with $\log$ linear link: mean $\mu=e^{h}, y \in(0, \infty)$. Then

$$
f(y, h)=\mu^{-1} e^{-y / \mu}=\exp \left\{-e^{-h} y-h\right\} .
$$

Here $a(h)=-e^{-h}, a^{\prime}=e^{-h}, \psi=e^{h}$. So both $\left|a^{\prime}\right|$ and $|\psi|$ grow at most exponentially in $|h|$.
(d) Binary logistic regression: mean $\mu=e^{h} /\left(1+e^{h}\right), y \in\{0,1\}$. Then

$$
f(y, h)=\mu^{y}(1-\mu)^{1-y}=\exp \left\{h y-\ln \left(1+e^{h}\right)\right\}
$$

Here $a(h)=h, a^{\prime}=1, \psi(h)=e^{h} /\left(1+e^{h}\right)$. So both $\left|a^{\prime}\right|$ and $|\psi|$ are bounded above by 1 .
(e) Binary probit regression: mean $\mu=\Phi(h) \equiv \int_{-\infty}^{h}\left(e^{-z^{2} / 2} / \sqrt{2 \pi}\right) d z$, $y \in\{0,1\}$. Then

$$
f(y, h)=\mu^{y}(1-\mu)^{1-y}=\exp \{y \ln (\Phi(h) /(1-\Phi(h)))+\ln (1-\Phi(h))\} .
$$

Here $a(h)=\ln \left(\Phi(h) /(1-\Phi(h)), a^{\prime}=\left[\Phi(h)^{-1}+\{1-\Phi(h)\}^{-1}\right] \Phi^{\prime}(h), \psi(h)=\right.$ $\Phi(h) \in[0,1]$. By using Mills' ratio, it can be shown that $\left|a^{\prime}(h)\right|$ increases at most linearly with $|h|$.

Using these rates of growth, we can make the condition on $D(\cdot)$ more specific for specific examples of GLM.

The conditions of Theorem 1 also depend on how the eigenvalues of $V_{\gamma}$ and $V_{\gamma}^{-1}$ behave. To be specific, assume that the largest eigenvalues of $V_{\gamma}$ and $V_{\gamma}^{-1}$, for $|\gamma| \leq \bar{r}_{n}$, are both bounded above by some power $\bar{r}_{n}^{v}(v>0)$, for all large enough $\bar{r}_{n}$.

The condition on $r_{n} \ln \bar{B}\left(r_{n}\right)$ then becomes redundant since $r_{n} \ln \bar{B}\left(r_{n}\right) \leq$ $\bar{r}_{n} \ln \tilde{B}_{n} \leq c \bar{r}_{n} \ln \bar{r}_{n} \leq c \bar{r}_{n} \ln K_{n} \prec n \varepsilon_{n}^{2}$ (for some constant $c>0$ and for all large enough $n$ ), since $\tilde{B}_{n}$ is bounded above by a power of $\bar{r}_{n}$.

The condition on $\bar{r}_{n} \ln \left(1 / \varepsilon_{n}^{2}\right)$ also becomes redundant [they are implied by the condition on $\bar{r}_{n} \ln \left(K_{n}\right)$ and $n \varepsilon_{n}^{2} \succ 1$ ] if we assume that $K_{n} \succ n^{\delta}$ for some $\delta>0$.

Consider now the condition on $\bar{r}_{n} \ln D(\cdot)$ for various regression models, depending on the rate of growth $D\left(\bar{r}_{n} \sqrt{n \varepsilon_{n}^{2} \tilde{B}_{n}}\right)$. This condition on $\bar{r}_{n} \ln D(\cdot)$ becomes redundant [it is implied by the condition on $\bar{r}_{n} \ln \left(K_{n}\right)$ ] for normal linear regression, binary logistic regression, and probit regression since $D(\cdot)$ is bounded above by some power of $\sqrt{\bar{r}_{n}^{2+v} n \varepsilon_{n}^{2}}$, which is bounded above by some power of $K_{n}$ (note that $\bar{r}_{n} \leq K_{n}, \varepsilon_{n} \leq 1$ and $K_{n} \succ n^{\delta}$ for some $\delta>0$ ). The condition $B\left(r_{n}\right) \prec n \varepsilon_{n}^{2}$ can be satisfied by requiring $\bar{r}_{n} \prec\left(n \varepsilon_{n}^{2}\right)^{1 / v}$.

For Poisson and exponential regressions with the log-linear link, however, $D(\cdot)$ grows exponentially in $\sqrt{\bar{r}_{n}^{2+v} n \varepsilon_{n}^{2}}$. The condition on $\bar{r}_{n} \ln D(\cdot)$ then cannot be ignored, and it can be satisfied if $\bar{r}_{n} \prec\left(n \varepsilon_{n}^{2}\right)^{1 /(4+v)}$ [which actually implies the later condition $B\left(r_{n}\right) \prec n \varepsilon_{n}^{2}$ and makes it redundant].

These are summarized as follows.
ThEOREM 2. Assume that the prior specification in Section 3 is used, such that $\max \left\{\sup _{\gamma:|\gamma| \leq \bar{r}_{n}} \operatorname{ch}_{1}\left(V_{\gamma}\right), \sup _{\gamma:|\gamma| \leq \bar{r}_{n}} \operatorname{ch}_{1}\left(V_{\gamma}^{-1}\right)\right\} \leq B \bar{r}_{n}^{v}$ for some positive constants $B$ and $v$, for all large enough $\bar{r}_{n}$. Suppose $\left|x_{j}\right| \leq 1$ for all $j$, and $\lim _{n \rightarrow \infty} \sum_{1}^{K_{n}}\left|\beta_{j}^{*}\right|<\infty$, where $K_{n}$ is a nondecreasing sequence in $n$ and $K_{n} \succ n^{\delta}$ for some $\delta>0$.

Let $\varepsilon_{n}$ be a sequence such that $\varepsilon_{n} \in(0,1]$ for each $n$ and $n \varepsilon_{n}^{2} \succ 1$ and assume that the following conditions also hold:

$$
\begin{align*}
\bar{r}_{n} \ln \left(K_{n}\right) & \prec n \varepsilon_{n}^{2},  \tag{10}\\
1 \leq r_{n} & \leq \bar{r}_{n}<K_{n},  \tag{11}\\
1 & \prec r_{n} \prec K_{n},  \tag{12}\\
\Delta\left(r_{n}\right) \equiv \inf _{\gamma:|\gamma|=r_{n}} \sum_{j: j \neq \gamma}\left|\beta_{j}^{*}\right| & \prec \varepsilon_{n}^{2} . \tag{13}
\end{align*}
$$

Also assume that

$$
\begin{equation*}
\bar{r}_{n} \prec\left(n \varepsilon_{n}^{2}\right)^{1 / v} \tag{14}
\end{equation*}
$$

for normal linear, binary logistic and binary probit regression; or assume

$$
\begin{equation*}
\bar{r}_{n} \prec\left(n \varepsilon_{n}^{2}\right)^{1 /(4+v)} \tag{15}
\end{equation*}
$$

for Poisson or exponential regression with log-linear link function. Then the results of Theorem 1 hold.

This result can be used to study the convergence rate $\varepsilon_{n}$ under various situations, depending on how $K_{n}$ grows with $n$, as well as how $\Delta\left(r_{n}\right)=\inf _{|\gamma|=r_{n}} \sum_{j \notin \gamma}\left|\beta_{j}^{*}\right|$ grows with $r_{n}$. Here are some corollaries, which follow by assuming an exponential decay rate of $\Delta(\cdot)$ and checking the conditions of Theorem 2. This includes as a special case only a fixed and finite number of $\left|\beta_{j}^{*}\right|$ 's being nonzero, while also allowing a more realistic setup with many small $\left|\beta_{j}^{*}\right|$ 's, none of which is exactly zero.

Corollary 1. Consider the examples of Poisson regression, exponential regression, normal linear regression, logistic regression or probit regression described before. Assume that the prior specification in Section 3 is used, such that $\max \left\{\sup _{\gamma:|\gamma| \leq \bar{r}_{n}} \operatorname{ch}_{1}\left(V_{\gamma}\right), \sup _{\gamma:|\gamma| \leq \bar{r}_{n}} \operatorname{ch}_{1}\left(V_{\gamma}^{-1}\right)\right\} \leq B \bar{r}_{n}^{v}$ for some positive constants $B$ and $v$, for all large enough $\bar{r}_{n}$. Suppose $\left|x_{j}\right| \leq 1$ for all $j$. Suppose $K_{n} \succ n^{\delta}$ for some $\delta>0$ and $K_{n} \leq e^{C n^{\xi}}$ for some $C>0$ and some $\xi \in(0,1)$, for all large enough $n$. Suppose $\lim _{n \rightarrow \infty} \sum_{j=1}^{K_{n}}\left|\beta_{j}^{*}\right|<\infty$. Also suppose for some $C^{\prime}>0, \Delta\left(r_{n}\right) \leq e^{-C^{\prime} r_{n}}$ for all large enough $n$, and

$$
\begin{equation*}
\left(C^{\prime}\right)^{-1} \ln n \leq r_{n} \leq \bar{r}_{n} \prec(\ln n)^{k} \tag{16}
\end{equation*}
$$

for some $k>1$. Then we can take the convergence rate in Theorem 2 as

$$
\begin{equation*}
\varepsilon_{n} \sim n^{-(1-\xi) / 2}(\ln n)^{k / 2} . \tag{17}
\end{equation*}
$$

REMARK 1 (Good convergence rate). Note that $n^{\alpha} \prec e^{C n^{\xi}}$ for any small $\xi>0$ and large $\alpha>0$. So if $K_{n} \sim n^{\alpha}$ for whatever large power $\alpha$, one can achieve a convergence rate $\varepsilon_{n} \sim n^{-(1-\xi) / 2}(\ln n)^{k / 2} \prec n^{-(1-2 \xi) / 2}$, where $\xi$ can be made arbitrarily close to zero. This gives a rate arbitrarily close to the "finite-dimensional" rate $1 / \sqrt{n}$, despite the large dimension $K_{n}$. We note also that these results suggest slowly growing $r$ and $\bar{r}_{n}$ between powers of $\ln n$, for achieving a near finitedimensional convergence rate. Since these are only sufficient conditions, it may be possible that other ranges of $r$ and $\bar{r}_{n}$ can also lead to a near finite-dimensional rate of convergence. The following result, for example, shows a good convergence rate even when $r$ and $\bar{r}_{n}$ grow slowly in some small power of $n$.

Corollary 2. Consider the setup of Corollary 1. For any $b \in(0, q)$, if instead of (16) we have

$$
\begin{equation*}
\left(C^{\prime}\right)^{-1} \ln n \leq r_{n} \leq \bar{r}_{n} \prec n^{b} \tag{18}
\end{equation*}
$$

then we can take the convergence rate as

$$
\begin{equation*}
\varepsilon_{n} \sim n^{-(1-\xi-b) / 2} \tag{19}
\end{equation*}
$$

Here the power $q=\min \{1-\xi, \delta, \xi /(3+v)\}$ for Poisson and exponential regressions with log-linear link function; $q=\min \{1-\xi, \delta\} I[v \leq 1]+\min \{1-$ $\xi, \delta, \xi /(v-1)\} I[v>1]$ for logistic, probit and normal linear regression.

REMARK 2 (Posterior consistency). The results on posterior consistency can be obtained as a special case by setting $\varepsilon_{n}=\varepsilon$ for any small but fixed $\varepsilon>0$. There is no need to assume a rate for $\Delta\left(r_{n}\right)$ for consistency results to hold, since $\Delta\left(r_{n}\right) \prec \varepsilon^{2}$ as long as $r_{n} \succ 1$ and $\lim _{n \rightarrow \infty} \sum_{1}^{K_{n}}\left|\beta_{j}^{*}\right|<\infty$. The previous Theorem 2 then implies that the following condition on $r_{n}$ and $\bar{r}_{n}$ is sufficient for posterior consistency:

$$
\begin{equation*}
1 \prec r_{n} \leq \bar{r}_{n} \prec \min \left\{K_{n}, n^{1 /(v+4)}, n /\left(\ln K_{n}\right)\right\} . \tag{20}
\end{equation*}
$$

A slightly more relaxed condition for consistency for the special cases of logistic and probit regression can be found in Jiang [15].

REMARK 3 (Normal linear regression with unknown dispersion). So far, for normal linear regression, we have assumed that $y \mid x \sim N\left(E(y \mid x), \varphi^{-1}\right)$ with dispersion parameter (inverse variance) $\varphi(>0)$ known. In practice, $\varphi$ is unknown and a gamma prior is often put on $\varphi$ (e.g., George and McCulloch [7], Kohn, Smith and Chan [17] and Dobra et al. [5]). For example, suppose conditional on model $\gamma, \varphi \mid \gamma \sim G a(\kappa, \rho)$ with prior density $\pi(\varphi \mid \gamma)=\rho^{\kappa} \varphi^{\kappa-1} e^{-\rho \varphi} / \Gamma(\kappa)$, $\beta_{\gamma} \mid \gamma, \varphi \sim N\left(0, \varphi^{-1} V_{\gamma}\right)$ and $\gamma$ follows the prior distribution of Section 3. With this prior specification, it can be shown that the statements regarding normal linear regression in Theorem 2 and Corollaries 1 and 2 are still valid, where we consider bounded covariates standardized such that $\left|x_{j}\right| \leq 1$ for all $j$.
5. Implications of posterior convergence. It is well-known that a convergence statement such as

$$
\begin{equation*}
\lim _{n \rightarrow \infty} P^{*}\left\{\pi_{n}\left[d\left(p, p^{*}\right) \leq \varepsilon_{n} \mid D^{n}\right] \geq 1-e^{-r_{0} n \varepsilon_{n}^{2}}\right\}=1 \tag{21}
\end{equation*}
$$

for some $r_{0}>0$, implies existence of point estimates of $p^{*}$ that have the same convergence rate $\varepsilon_{n}$ in the frequentist sense. Such a point estimate can be obtained by finding the center of an $\varepsilon_{n}$-ball with high posterior probability, or by posterior expectation (e.g., Ghosal, Ghosh and van der Vaart [10]).

A point estimate can also be formed by a generalization of posterior expectation called a "selected posterior estimate" (Jiang [15]). For example,

$$
\begin{equation*}
\hat{p}_{A}=\int p \pi_{A}\left(d p \mid D^{n}\right) \tag{22}
\end{equation*}
$$

where $\pi_{A}\left(d p \mid D^{n}\right)=\pi\left(d p \mid D^{n}, p \in A\right)$, and $p \in A$ is a selection rule, possibly data dependent. A rule of this kind, for example, can be averaging over several of the best models, which are indexed by $\gamma$ 's having the largest marginal posteriors $\pi\left(\gamma \mid D^{n}\right)$. For example, Smith and Kohn [23] considered the use of the best model, and Sha et al. [22] averaged over the ten best models. A rule can also be defined by using the models that include the individually strongest variables. For example, include a model $\gamma$ in the posterior average if $\gamma$ includes a variable $j$ that appears more than 5 percent of time in the posterior distribution [i.e., if $\left.\pi\left(\gamma_{j}=1 \mid D^{n}\right)>0.05\right]$. See, for example, Lee et al. [19].

Suppose a rule $A$ has selection probability $\pi\left\{p \in A \mid D^{n}\right\}>r$ for some constant $r>0$. Then the convergence rate of $\hat{p}_{A}$ can be studied by using the relations

$$
\begin{equation*}
d\left(\hat{p}_{A}, p^{*}\right)^{2} \leq \varepsilon_{n}^{2}+2 \pi\left[d\left(p, p^{*}\right)>\varepsilon_{n} \mid D^{n}\right] / r \tag{23}
\end{equation*}
$$

and

$$
\begin{equation*}
P^{*}\left[d\left(\hat{p}_{A}, p^{*}\right)^{2} \leq \varepsilon_{n}^{2}+2 \delta_{n} / r\right] \geq P^{*}\left[\pi\left(d\left(p, p^{*}\right)>\varepsilon_{n} \mid D^{n}\right) \leq \delta_{n}\right] \tag{24}
\end{equation*}
$$

which follows a familiar treatment based on convexity of $p \mapsto d\left(p, p^{*}\right)^{2}$ (e.g., Ghosal, Ghosh and van der Vaart [10]). The term $\delta_{n} / r$ can usually be taken as $e^{-r_{0} n \varepsilon_{n}^{2}}$ [see result (i) of Theorem 1], which is negligible compared to $\varepsilon_{n}^{2}$ under conditions $\bar{r}_{n} \ln \left(1 / \varepsilon_{n}^{2}\right) \prec n \varepsilon_{n}^{2}$ and $1 \leq \bar{r}_{n}$ of Theorem 1.

For regression purposes, a related mean estimate can be constructed as $\hat{\mu}_{A}(x)=$ $\int y \hat{p}_{A}(y \mid x) v_{y}(d y)$. When binary response $y$ is considered, a classifier can be defined as $\hat{C}_{A}(x)=I\left[\hat{\mu}_{A}(x)>0.5\right]$.

In the general case, there is no relationship bounding the $L_{2}$ distance $E_{x}\left(\hat{\mu}_{A}-\right.$ $\left.\mu^{*}\right)^{2}$ between the estimated mean and the true mean using $d\left(\hat{p}_{A}, p^{*}\right)$, because the latter is bounded but the former is not. However, a weighted $L_{2}$ difference can be bounded as

$$
\begin{equation*}
\int \frac{\left(\hat{\mu}_{A}-\mu^{*}\right)^{2}}{\hat{v}_{A}+v^{*}} v_{x}(d x) \leq 2 d\left(\hat{p}_{A}, p^{*}\right)^{2} \tag{25}
\end{equation*}
$$

where $\nu^{*}(x)=\int y^{2} p^{*}(y \mid x) \nu_{y}(d y)$ and $\hat{v}_{A}(x)=\int y^{2} \hat{p}_{A}(y \mid x) \nu_{y}(d y)$. This is obtained by noting that $\left(\hat{\mu}_{A}-\mu^{*}\right)^{2}=\left\{\int y\left(\sqrt{\hat{p}_{A}}+\sqrt{p^{*}}\right) \cdot\left(\sqrt{\hat{p}_{A}}-\sqrt{p^{*}}\right) \nu_{y}(d y)\right\}^{2}$ and applying the Cauchy-Schwarz inequality.

Since the denominator is at most 2 for binary $y$, the above relation actually leads to a bound for the unweighted $L_{2}$ distance between the means, which further leads to a bound for the classification error due to Corollary 6.2 of Devroye, Györfi and Lugosi [4]. This is summarized below and was used in proving regression and classification consistency in Jiang [15]:

$$
\begin{gather*}
E_{D^{n}}^{*} P_{(x, y)}^{*}\left(\hat{C}_{A}(x) \neq y \mid D^{n}\right)-L^{*} \\
\leq E_{D^{n}}^{*} 2 \sqrt{E_{x}\left(\hat{\mu}_{A}-\mu^{*}\right)^{2}} \tag{26}
\end{gather*}
$$

$$
\begin{align*}
& \leq E_{D^{n}}^{*} 2 \sqrt{4 d\left(\hat{p}_{A}, p^{*}\right)^{2}} \leq 2 \sqrt{4 E_{D^{n}}^{*} d\left(\hat{p}_{A}, p^{*}\right)^{2}}  \tag{27}\\
& \leq 4 \sqrt{\varepsilon_{n}^{2}+2 E_{D^{n}}^{*} \pi\left[d\left(p, p^{*}\right)>\varepsilon_{n} \mid D^{n}\right] / r} \tag{28}
\end{align*}
$$

[The last step is due to (23).] Here $L^{*}=P_{(x, y)}^{*}\left\{C^{*}(x) \neq y\right\}$ is the "Bayes error," where $C^{*}(x)=I\left[\mu^{*}(x)>1 / 2\right]$ is the ideal "Bayes rule" based on the (unknown) true mean function $\mu^{*}$. According to Theorem 1(iii), the term $E_{D^{n}}^{*} \pi\left[d\left(p, p^{*}\right)>\right.$ $\varepsilon_{n} \mid D^{n}$ ] can be made exponentially small (of the form $e^{-c_{1} n \varepsilon_{n}^{2}}$ for some $c_{1}>0$ ), which is negligible when compared to $\varepsilon_{n}^{2}$ as commented earlier. This implies that the error of the classification rule $\hat{C}_{A}(x)$ is at most $5 \varepsilon_{n}$ above that of the optimal Bayes rule, for all large enough $n$. So $\varepsilon_{n}$ also forms a rate of convergence to the optimal Bayes error for the purpose of classification. Here the convergence rate $\varepsilon_{n}$ can be made to be near "finite-dimensional" (nearly $1 / \sqrt{n}$ ) by Bayesian variable selection, despite a high dimension $K_{n} \sim n^{\alpha} \gg n$, in situations commented on earlier (e.g., Corollary 1 and Remark 1).

These convergence rate results show that even in high dimensions with $\operatorname{dim}(x) \gg n$, a good convergence rate can be achieved when the effect of $x$ is "sparse." For such sparse problems Bayesian variable selection can therefore help to alleviate "overfitting" or the "curse of dimensionality."
6. Gaussian variable selection and graphical models. In this section we will assume that $x_{1}^{J_{n}} \equiv\left(x_{1}, \ldots, x_{K_{n}}, y\right)$ are multivariate Gaussian and have been standardized to have $E\left(x_{k}\right)=0, \operatorname{var}\left(x_{k}\right)=1$. Here $y$ is regarded as $x_{J_{n}}$, where $J_{n}=K_{n}+1$. The effects of $x_{k \neq j}$ on $x_{j}$ are summarized by the regression coefficients $\beta_{j \mid k}^{*}$ used in the induced relation $E\left(x_{j} \mid x_{k \neq j}\right)=\sum_{k \neq j} \beta_{j \mid k}^{*} x_{k}$.

In Gaussian graphical models, relations among $x_{1}^{J_{n}}$ are described by a graph, such that a node corresponding to $x_{j}$ is only connected to a "neighborhood" $\left(x_{k}\right)_{k \in n b_{j}}$, where $n b_{j}$ is a subset of $\left\{1, \ldots, J_{n}\right\} \backslash\{j\}$, which indicates selected variables used in regression modeling of $x_{j} \mid x_{k \neq j}$. Therefore, the Bayesian variable selection technique can be used for studying the neighborhood of a variable $x_{j}$ (see, e.g., Dobra et al. [5]). We will consider the situation when none of the effects of $x_{k \neq j}$ on $x_{j}$ is exactly zero. In this case, the usual consistency of selecting the "true graph" (e.g., Meinshausen and Bühlmann [20]) will not be studied here, since the true graph is the saturated graph adopting all $K_{n}$ variables $x_{k \neq j}$ to explain each $x_{j}$. In the high-dimensional case $K_{n} \gg n$, such a "true model" is obviously not very useful. Nevertheless, in such a situation, Bayesian variable selection can still be shown to produce "good" models that are much simpler and yet are still "consistent," if the effects of these $K_{n}$ variables decay sufficiently fast (when ordered in some way). Here "consistency" is in a different sense-these simplified models, picking up only a small number out of all the $K_{n}$ nonzero regression coefficients, will be consistent in terms of producing probability densities "often close" to the true probability density. In this approach, one first uses Bayesian variable selection
to obtain such "good" density estimates, for all $p^{*}\left(x_{j} \mid x_{k \neq j}\right), j=1, \ldots, J_{n}$; then one can construct graphs to summarize the conditional independence structures corresponding to these "good" density estimates. (One can systematically decide to either include or exclude one-sided connections in these graphs (see, e.g., Meinshausen and Bühlmann [20]) when some $x_{k}$ is used in modeling $x_{j} \mid x_{k \neq j}$ but $x_{j}$ is not used in modeling $x_{k} \mid x_{j \neq k}$.)

We are interested in making inference on $J_{n}\left(=K_{n}+1\right)$ (conditional) densities $p^{*}\left(x_{j} \mid x_{k \neq j}\right), j=1, \ldots, J_{n}$, in order to construct a graph. We hope that the $P^{*}$ probability for not reliably estimating each density is small enough so that the $P^{*}$ probability is small for any density to be badly estimated. In other words, we would like to have a bound of $P^{*}$ probability of large errors. For now, pick any $x_{j}$ as the response $y$ and consider its regression on the $x_{k}$ 's $(k \neq j)$. To mimic the regression setup, we can reorder the indices of the $x_{k \neq j}$ 's as $x_{1}^{K_{n}}$. We will use the prior specified in Remark 3. A result as in Theorem 1(ii), obtained when assuming uniformly bounded $\left|x_{k}\right|$ 's, could be used for this purpose of bounding the total error out of the $J_{n}$ regression analyses.

In the current situation of Gaussian graphical models, however, the $x_{k}$ 's are Gaussian instead of being uniformly bounded. In this case, for result (ii) of Theorem 1 to hold, we will change the condition on $\Delta\left(r_{n}\right)=\inf _{|\gamma|=r_{n}} \sum_{k \notin \gamma}\left|\beta_{k}^{*}\right|$ from $\Delta\left(r_{n}\right) \prec \varepsilon_{n}^{2}$ to $K_{n} \Delta\left(r_{n}\right) \prec \varepsilon_{n}^{2}$. This would be satisfied if $\Delta\left(r_{n}\right)$ decays exponentially fast in $r_{n}, r_{n} \succ \ln n$, and $K_{n}$ grows at most polynomially. After taking into account some other conditions, we obtain the following theorem.

THEOREM 3. Consider the prior specification in Remark 3. (When selecting the neighborhood for each $x_{j}$, treat $x_{j}$ as $y$ and $x_{k \neq j}$ as $x_{1}^{K_{n}}$.) Assume that

$$
\max \left\{\sup _{\gamma:|\gamma| \leq \bar{r}_{n}} \operatorname{ch}_{1}\left(V_{\gamma}\right), \sup _{\gamma:|\gamma| \leq \bar{r}_{n}} \operatorname{ch}_{1}\left(V_{\gamma}^{-1}\right)\right\} \leq B \bar{r}_{n}^{v}
$$

for some positive constants $B$ and $v$, for all large enough $\bar{r}_{n}$.
Suppose that, for each $x_{j}$, the effects of the other variables $x_{k \neq j}$ satisfy $\lim _{n \rightarrow \infty} \sum_{k \in \mathcal{K}_{j}}\left|\beta_{j \mid k}^{*}\right|<\infty$, where $\mathcal{K}_{j}=\left\{1, \ldots, K_{n}+1\right\} \backslash\{j\}$. In addition, assume that there exists some $C^{\prime}>0$, such that, for all large enough $n$, $\inf _{\gamma \subset \mathcal{K}_{j},|\gamma|=r_{n}} \sum_{k \in \mathcal{K}_{j} \backslash \gamma}\left|\beta_{j \mid k}^{*}\right| \leq e^{-C^{\prime} r_{n}}$.

Assume that $n^{\delta} \prec K_{n} \prec n^{\alpha}$ for some $\alpha>\delta>0$.
Assume also for some $\xi \in(0,1)$

$$
\begin{equation*}
\ln n \prec r_{n} \leq \bar{r}_{n} \prec n^{b}, \quad \text { where } b<\min \{\delta, \xi, \xi / v\} . \tag{29}
\end{equation*}
$$

Then we have, for some constant $c_{1,2,3}^{\prime}>0$, for all sufficiently large $n$,

$$
\begin{equation*}
P^{*}\left[\pi\left(h_{j} \leq n^{-(1-\xi) / 2} \mid D^{n}\right) \geq 1-e^{-c_{1}^{\prime} n^{\xi}}, j=1, \ldots, K_{n}+1\right] \geq 1-n^{\alpha} e^{-c_{2}^{\prime} n^{\xi}} \tag{i}
\end{equation*}
$$

and
(ii)

$$
P^{*}\left[\hat{h}_{j, A_{j}} \leq c_{3}^{\prime} n^{-(1-\xi) / 2}, j=1, \ldots, K_{n}+1\right] \geq 1-n^{\alpha} e^{-c_{2}^{\prime} n^{\xi}}
$$

Here we define, for $j=1, \ldots, K_{n}+1$,

$$
\begin{gather*}
h_{j}=\left\{\int_{\Re} \Re_{K_{n}+1}\left|p\left(x_{j} \mid x_{k \neq j}\right)^{1 / 2}-p^{*}\left(x_{j} \mid x_{k \neq j}\right)^{1 / 2}\right|^{2} p^{*}\left(x_{k \neq j}\right) d x_{1}^{K_{n}+1}\right\}^{1 / 2}  \tag{30}\\
\hat{h}_{j, A_{j}}=\left\{\int_{\Re} K_{n}\left|\hat{p}_{A_{j}}\left(x_{j} \mid x_{k \neq j}\right)^{1 / 2}-p^{*}\left(x_{j} \mid x_{k \neq j}\right)^{1 / 2}\right|^{2}\right.  \tag{31}\\
\left.\times p^{*}\left(x_{k \neq j}\right) d x_{1}^{K_{n}+1}\right\}^{1 / 2}
\end{gather*}
$$

where $p^{*}$ represents the true density and $\hat{p}_{A_{j}}$ is a selected posterior estimate [as defined in (22)] corresponding to a selection rule $A_{j}$, such that the selection probability $\pi\left(p \in A_{j} \mid D^{n}\right)>r$ for some $r>0$.

Therefore, a near finite-dimensional rate of convergence can be achieved (for some small $\xi>0$ ), jointly for all neighborhoods of $x_{j}, j=1, \ldots, K_{n}+1$, despite the fact that $K_{n}$ can follow a large power of $n$.
7. General prior. In this section we consider the case $\left|x_{j}\right| \leq 1$ for all $j$ and mainly focus on the GLM models as described in Section 2, where $a(h)$ and $b(h)$ contain no additional parameters other than $h$. (Similar conditions and results can be formulated for normal linear regression with unknown error variance.)

Here we consider the general conditions on the prior $\pi\left(\gamma, \beta_{\gamma}\right)$ for producing rate of convergence $\varepsilon_{n}$, which is a sequence in $n$, which we assume to satisfy $\varepsilon_{n} \in(0,1]$ for Conditions ( N ) and ( O ) below.

Condition (N) requires a not too little prior to be placed over a very small neighborhood of the true density $p^{*}$. Condition (O) requires a very little prior to be placed outside of a region that is not too complex in some sense.

CONDITION ( N ) (For prior $\pi$ on an approximation neighborhood). Assume that a sequence of (nonempty) models $\gamma_{n}$ exists such that, as $n$ increases,

$$
\begin{equation*}
\sum_{j \notin \gamma_{n}}\left|\beta_{j}^{*}\right| \prec \varepsilon_{n}^{2} \tag{32}
\end{equation*}
$$

and for any sufficiently small $\eta>0$, there exists $N_{\eta}$ such that, for all $n>N_{\eta}$, we have

$$
\begin{equation*}
\pi\left(\gamma=\gamma_{n}\right) \geq e^{-n \varepsilon_{n}^{2} / 8} \tag{33}
\end{equation*}
$$

and

$$
\begin{equation*}
\pi\left(\beta_{\gamma} \in M\left(\gamma_{n}, \eta\right) \mid \gamma=\gamma_{n}\right) \geq e^{-n \varepsilon_{n}^{2} / 8} \tag{34}
\end{equation*}
$$

where $M\left(\gamma_{n}, \eta\right)=\left(\beta_{j}^{*} \pm \eta \varepsilon_{n}^{2} /\left|\gamma_{n}\right|\right)_{j \in \gamma_{n}}$.

CONDITION (O) (For prior $\pi$ outside of a not-too-complex region). Let $D(R)=1+R \sup _{|h| \leq R}\left|a^{\prime}(h)\right| \cdot \sup _{|h| \leq R}|\psi(h)|$ for any $R>0$. There exist some $C_{n}>0$ and some $\bar{r}_{n}$ satisfying $1 \leq \bar{r}_{n}<K_{n}$, such that

$$
\begin{align*}
\bar{r}_{n} \ln \left(1 / \varepsilon_{n}^{2}\right) & \prec n \varepsilon_{n}^{2},  \tag{35}\\
\bar{r}_{n} \ln K_{n} & \prec n \varepsilon_{n}^{2},  \tag{36}\\
\bar{r}_{n} \ln D\left(\bar{r}_{n} C_{n}\right) & \prec n \varepsilon_{n}^{2} . \tag{37}
\end{align*}
$$

Furthermore, for all large enough $n$, the following two equations hold:

$$
\begin{equation*}
\pi\left(|\gamma|>\bar{r}_{n}\right) \leq e^{-4 n \varepsilon_{n}^{2}}, \tag{38}
\end{equation*}
$$

and for all $\gamma$ such that $|\gamma| \leq \bar{r}_{n}$, for all $j \in \gamma$,

$$
\begin{equation*}
\pi\left(\left|\beta_{j}\right|>C_{n} \mid \gamma\right) \leq e^{-4 n \varepsilon_{n}^{2}} \tag{39}
\end{equation*}
$$

These conditions allow a larger variety of priors. For example, one can use a uniform prior of $\gamma$ over all models with complexity $|\gamma| \leq \bar{r}_{n}$, where $\bar{r}_{n}$ can be taken to follow some rate of growth depending on the convergence rate $\varepsilon_{n}$ desired, and depending on the "bias" rate $\Delta(r)=\inf _{\gamma:|\gamma|=r} \sum_{j \notin \gamma}\left|\beta_{j}^{*}\right|$.

Before we truncated $\pi(\gamma)$ such that $\pi\left[|\gamma|>\bar{r}_{n}\right]=0$. This may not be desirable since we are forbidding the model to be too complex in the prior. We here notice that this truncation is not necessary. We can allow the prior to propose very complicated models with large $|\gamma|$, as long as the prior probability of $|\gamma|>\bar{r}_{n}$ is sufficiently small.

THEOREM 4 (Convergence rate under general prior). For GLM models with bounded covariates $\left|x_{j}\right| \leq 1$ for all $j$, suppose the true regression coefficients satisfy $\lim _{n \rightarrow \infty} \sum_{j=1}^{K_{n}}\left|\beta_{j}^{*}\right|<\infty$.

Let $\varepsilon_{n} \in(0,1]$ be a sequence such that $n \varepsilon_{n}^{2} \rightarrow \infty$. Denote $d\left(p, p^{*}\right)^{2}=$ $\iint\left|p\left(y, x \mid \gamma, \beta_{\gamma}\right)^{1 / 2}-p^{*}(y, x)^{1 / 2}\right|^{2} v_{y}(d y) v_{x}(d x)$. If the prior specification satisfies both Conditions $(\mathrm{N})$ and $(\mathrm{O})$, then we have the following (successively stronger) results:
(i)

$$
\lim _{n \rightarrow \infty} P^{*}\left\{\pi_{n}\left[d\left(p, p^{*}\right) \leq 4 \varepsilon_{n} \mid D^{n}\right] \geq 1-2 e^{-n \varepsilon_{n}^{2} / 4}\right\}=1
$$

(ii) For all sufficiently large $n$,

$$
P^{*}\left\{\pi\left[d\left(p, p^{*}\right)>4 \varepsilon_{n} \mid D^{n}\right] \geq 2 e^{-n \varepsilon_{n}^{2} / 4}\right\} \leq 2 e^{-n \varepsilon_{n}^{2} / 4}
$$

(iii) For all sufficiently large n,

$$
E_{D^{n}}^{*} \pi\left[d\left(p, p^{*}\right)>4 \varepsilon_{n} \mid D^{n}\right] \leq 4 e^{-n \varepsilon_{n}^{2} / 2}
$$

Results (i), (ii) and (iii) of this theorem will be proved by verifying some sufficient conditions for posterior convergence (to be summarized at the beginning of Section 8). These results, with bounded $x_{j}$ 's, will then be used to prove all the previous results on convergence rates, when specific priors as given in Section 3 are used. The only exception is the result in Section 6, where $x_{j}$ 's are jointly normal; they will be obtained by directly verifying the sufficient conditions in Section 8.

These conditions below are based on the Hellinger metric entropy and will be used to obtain posterior convergence rates under the GLM framework. Note that the method involved here is different from that in Jiang [15], who uses the Hellinger bracketing entropy and its upper bound of a parametric covering number (see, e.g., Theorem 3, Lee [18]). That method does not directly apply to modeling unbounded responses such as Gaussian and Poisson responses. (When applied to, e.g., Poisson regression, the upper bound of the bracketing entropy would require a too small restricted parameter space, on which the prior would place a nonnegligible probability.)
8. Proofs. We first use a proposition to summarize a set of sufficient conditions for establishing rates of posterior convergence. These serve as just one possible set of working conditions that we find convenient to use here, through which we have established our results; there exist several other alternatives, possibly with more relaxed conditions, for example, in Ghosal, Ghosh and van der Vaart [10] or Zhang [26].

Suppose $\mathscr{P}_{n}$ is a sequence of sets of probability densities. (For each $n$, denote $\mathscr{P}_{n}^{c}$ as the complement-the set of densities not in $\mathscr{P}_{n}$.) Suppose $\varepsilon_{n}$ is a sequence of positive numbers.

Suppose $N\left(\varepsilon_{n}, \mathscr{P}_{n}\right)$ is the minimal number of Hellinger balls of radius $\varepsilon_{n}$ that are needed to cover $\mathscr{P}_{n}$. [I.e., $N\left(\varepsilon_{n}, \mathscr{P}_{n}\right)$ is the minimum of all $k$ such that there exist $S_{j}=\left\{p: d\left(p, p_{j}\right) \leq \varepsilon_{n}\right\}, j=1, \ldots, k$, such that $\bigcup_{j=1}^{k} S_{j} \supset \mathcal{P}_{n}$, where $d(p, q)=\sqrt{\int(\sqrt{p}-\sqrt{q})^{2}}$ denotes the Hellinger distance between densities $p$ and q.]

Let the components of $D^{n}=\left(w^{(1)}, \ldots, w^{(n)}\right)$ be i.i.d. with true density $p^{*}$, where $\operatorname{dim}\left(w^{(1)}\right)$ and $p^{*}$ can depend on $n$. Denote $\pi(\cdot)$ as the prior (which is allowed to depend on $n$ by using, e.g., an increasing number of parameters to parameterize the density as $n$ increases), $\pi\left(\cdot \mid D^{n}\right)$ as the posterior and $\hat{\pi}(\varepsilon)=$ $\pi\left[d\left(p, p^{*}\right)>\varepsilon \mid D^{n}\right]$ for each $\varepsilon>0$. Define the KL difference as $d_{0}\left(p, p^{*}\right)=$ $\int p^{*} \ln \left(p^{*} / p\right)$. Define also a $d_{t}$ difference as $d_{t}\left(p, p^{*}\right)=t^{-1}\left(\int p^{*}\left(p^{*} / p\right)^{t}-1\right)$ for any $t>0$, which is used in, for example, Wong and Shen [25]. (Note that $d_{t}$ decreases to $d_{0}$ as $t$ decreases toward 0 .)

Denote $P^{*}$ and $E^{*}$ as the respective probability measure and the expectation for the data $D^{n}$.

Define the following conditions:
(a) $\ln N\left(\varepsilon_{n}, \mathcal{P}_{n}\right) \leq n \varepsilon_{n}^{2}$ for all sufficiently large $n$;
(b) $\pi\left(\mathcal{P}_{n}^{c}\right) \leq e^{-2 n \varepsilon_{n}^{2}}$ for all sufficiently large $n$;
(c) for all small enough $\gamma>0$ and $r>0$, there exists $N_{\gamma, r}$ such that for all $n \geq N_{\gamma, r}, \pi\left[p: d_{0}\left(p, p^{*}\right) \leq \gamma \varepsilon_{n}^{2}\right] \geq e^{-r n \varepsilon_{n}^{2}} ;$
(d) $\pi\left[p: d_{t}\left(p, p^{*}\right) \leq \varepsilon_{n}^{2} / 4\right] \geq e^{-n \varepsilon_{n}^{2} / 4}$ for all sufficiently large $n$, for some $t>0$.

Proposition 1. Suppose $n \varepsilon_{n}^{2} \succ 1$. Then, under (a), (b) and (c), we have:
(i)

$$
\lim _{n \rightarrow \infty} P^{*}\left[\hat{\pi}\left(4 \varepsilon_{n}\right)<2 e^{-n \varepsilon_{n}^{2} \min \{1, b / 2\}}\right]=1
$$

under (a), (b) and (d) (for some $t>0$ ), we have
(ii)

$$
P^{*}\left[\hat{\pi}\left(4 \varepsilon_{n}\right) \geq 2 e^{-n \varepsilon_{n}^{2} \min \{1 / 2, t / 4\}}\right] \leq 2 e^{-n \varepsilon_{n}^{2} \min \{1 / 2, t / 4\}},
$$

(iii)

$$
E^{*} \hat{\pi}\left(4 \varepsilon_{n}\right) \leq 4 e^{-n \varepsilon_{n}^{2} \min \{1, t / 2\}}
$$

The proof of this proposition follows the spirit of Ghosal, Ghosh and van der Vaart [10]. The details are omitted here and are included in a technical report (Jiang [14]).

Proof of Theorem 4. We prove result (iii) only, since it implies (ii) by Markov's inequality, which further implies (i). Result (iii) is proven by applying Proposition 1 with $t=1$. The proof is completed by checking conditions (d), (a) and (b) below.

Checking condition (d) for $t=1$. Denote the GLM density as $f(y, h)=$ $\exp \{a(h) y+b(h)+c(y)\}$. Then $p^{*}=f\left(y, h^{*}\right)$, where $h^{*}=x^{T} \beta^{*}=\sum_{j=1}^{K_{n}} x_{j} \beta_{j}^{*}$. Let $p_{\gamma}=f\left(y, h_{\gamma}\right)$, where $h_{\gamma}=x_{\gamma}^{T} \beta_{\gamma}=\sum_{j \in \gamma_{n}} x_{j} \beta_{j}$, where $\gamma_{n}$ is the model in Condition (N).

When $h^{*}$ and $h_{\gamma}$ are close enough, $d_{t}\left(p_{\gamma}, p^{*}\right)$ (for $t=1$ ) can be put in a form $d_{t}\left(p_{\gamma}, p^{*}\right)=E_{x} g\left(h^{i}\right)\left(h^{*}-h_{\gamma}\right)$, by integrating out $y$ and applying a first-order Taylor expansion. Here $g$ is a continuous derivative function in a neighborhood of $h^{*}$ and $h^{i}$ is an intermediate point between $h^{*}$ and $h_{\gamma}$. Note that $\left|h^{i}-h^{*}\right| \leq$ $\left|h_{\gamma}-h^{*}\right| \leq\left|\sum_{j \notin \gamma_{n}} x_{j} \beta_{j}^{*}\right|+\left|\sum_{j \in \gamma_{n}} x_{j}\left(\beta_{j}-\beta_{j}^{*}\right)\right| \leq \Delta_{n}+r_{n} \delta_{n}$, when the $x_{j}$ 's are bounded by 1 and $\beta_{j} \in\left(\beta_{j}^{*} \pm \delta_{n}\right)$ for all $j \in \gamma_{n}$. Here $r_{n}=\left|\gamma_{n}\right|$ and $\Delta_{n}=$ $\sum_{j \notin \gamma_{n}}\left|\beta_{j}^{*}\right|$ is assumed to satisfy $\Delta_{n} \prec \varepsilon_{n}^{2}$.

For sufficiently small $r_{n} \delta_{n},\left|g\left(h^{i}\right)\right|$ is bounded since $\left|h^{i}\right| \leq\left|h^{*}\right|+\left|h^{i}-h^{*}\right| \leq$ $B_{0}+\Delta_{n}+r_{n} \delta_{n}$ is bounded, where $B_{0}=\lim _{n \rightarrow \infty} \sum_{j=1}^{K_{n}}\left|\beta_{j}^{*}\right|$. Then $d_{t}\left(p_{\gamma}, p^{*}\right) \leq$ $C\left(\Delta_{n}+r_{n} \delta_{n}\right)$ for some constant $C$, for all small enough $r_{n} \delta_{n}$.

We will take $\delta_{n}=\eta \varepsilon_{n}^{2} /\left|\gamma_{n}\right|$ for some small enough $\eta>0$. This will make $d_{t} \leq$ $\varepsilon_{n}^{2} / 4$ for all large enough $n$, since $\Delta_{n} \prec \varepsilon_{n}^{2}$.

This implies that the set of densities $S=\left\{p\left(\cdot \mid \gamma_{n}, \beta\right): \beta \in\left(\beta_{j}^{*} \pm \delta_{n}\right)_{j \in \gamma_{n}}\right\}$ is contained in $T=\left\{p: d_{t}\left(p, p^{*}\right) \leq \varepsilon_{n}^{2} / 4\right\}$. The conditions on $\pi\left(\gamma_{n}\right)$ and $\pi(\beta \in$ $\left.\left(\beta_{j}^{*} \pm \delta_{n}\right)_{j \in \gamma_{n}} \mid \gamma_{n}\right)$ then imply that $\pi(T) \geq \pi(S) \geq e^{-n \varepsilon_{n}^{2} / 4}$ for all large $n$, confirming condition (d).

Checking condition (a). Each density $p$ is labeled by a model index $\gamma$ and the corresponding regression coefficients $\beta_{j}$. We will define $\mathcal{P}_{n}$ as the set of densities that can be represented with $|\gamma|$ (the number of nonzero regression parameters) being at most $\bar{r}_{n}$, and with each parameter $\left|\beta_{j}\right| \leq C_{n}$.

The corresponding space of regression parameters can be covered by small $\ell_{\infty}$ balls of the form $B=\left(v_{j} \pm \delta\right)_{j=1}^{K_{n}}$, of radius $\delta>0$. For each model $\gamma$ in $\mathcal{P}_{n}$, there are $|\gamma|$ nonzero components of $\beta_{j}$, valued in $\pm C_{n}$. It takes at most $\left[\left(2 C_{n}\right) /(2 \delta)+\right.$ $1]^{|\gamma|}$ balls to cover the parameter space of model $\gamma$ in $\mathscr{P}_{n}$. [The centers of these balls can be taken inside the parameter space of model $\gamma$, so that each center $v=\left(v_{j}\right)_{1}^{K_{n}}$ has components satisfying $v_{j}=0 \forall j \notin \gamma$ and $\left|v_{j}\right| \leq C_{n} \forall j \in \gamma$.]

There are at most $K_{n}^{r}$ models of size $|\gamma|=r$, and $r=0,1,2, \ldots, \bar{r}_{n}$. These show that $N(\delta)$, the number of size- $\delta$ balls needed to cover the space of regression parameters for $\mathcal{P}_{n}$, is at most $\sum_{r=0}^{\bar{r}_{n}} K_{n}^{r}\left[\left(2 C_{n}\right) /(2 \delta)+1\right]^{r}$, which is bounded above by $\left(\bar{r}_{n}+1\right)\left(K_{n}\left(C_{n} / \delta+1\right)\right)^{\bar{r}_{n}}$.

Given any density in $\mathscr{P}_{n}$, it can be represented by a set of regression parameters $\left(u_{j}\right)_{1}^{K_{n}}$ falling in one of these $N(\delta)$ balls, say, ball $B=\left(v_{j} \pm \delta\right)_{j=1}^{K_{n}}$, where $u_{j}$ and $v_{j}$ are zero for the same set of components $\gamma$, where $|\gamma| \leq \bar{r}_{n}$.

Consider the corresponding GLM densities $p_{u, v}=\exp \left\{y a\left(h_{u, v}\right)+b\left(h_{u, v}\right)+\right.$ $c(y)\}$, where $h_{u}=\sum_{j=1}^{K_{n}} u_{j} x_{j}$ and $h_{v}=\sum_{j=1}^{K_{n}} v_{j} x_{j}$. Then the Hellinger distance $d\left(p_{u}, p_{v}\right) \leq\left\{d_{0}\left(p_{u}, p_{v}\right)\right\}^{1 / 2}$, where the KL difference $d_{0}\left(p_{u}, p_{v}\right)=$ $E_{x} \int p_{v}\left(\ln p_{v}-\ln p_{u}\right) v_{y}(d y)$. After integration in $y$, one can apply a Taylor expansion and show that $d_{0}\left(p_{u}, p_{v}\right) \leq E_{x}\left(a^{\prime}\left(h^{i}\right) \psi\left(h_{v}\right)+b^{\prime}\left(h^{i}\right)\right)\left(h_{v}-h_{u}\right)$, where $\psi=-b^{\prime} / a^{\prime}$ and $h^{i}$ is an intermediate point between $h_{v}$ and $h_{u}$. Note that $u$ and $v$ both have components bounded in value by $C_{n}$ and they have zero components out of a same set, say, $\gamma$, such that $|\gamma| \leq \bar{r}_{n}$. Therefore, $h_{v}$ and $h_{u}$ (and therefore, also $h^{i}$ ) are bounded above by $\bar{r}_{n} C_{n}$. Note also that $\left|h_{v}-h_{u}\right|=\left|\sum_{j \in \gamma} x_{j}\left(v_{j}-u_{j}\right)\right| \leq$ $\bar{r}_{n} \delta$, since $\left|x_{j}\right| \leq 1,\left|v_{j}-u_{j}\right| \leq \delta$ and $|\gamma| \leq \bar{r}_{n}$. Therefore,

$$
\begin{equation*}
d_{0}\left(p_{u}, p_{v}\right) \leq 2 \sup _{|h| \leq \bar{r}_{n} C_{n}}\left|a^{\prime}(h)\right| \sup _{|h| \leq \bar{r}_{n} C_{n}}|\psi(h)| \bar{r}_{n} \delta \tag{40}
\end{equation*}
$$

and

$$
\begin{equation*}
d\left(p_{u}, p_{v}\right) \leq\left\{2 \sup _{|h| \leq \bar{r}_{n} C_{n}}\left|a^{\prime}(h)\right| \sup _{|h| \leq \bar{r}_{n} C_{n}}|\psi(h)| \bar{r}_{n} \delta\right\}^{1 / 2} \tag{41}
\end{equation*}
$$

So $d\left(p_{u}, p_{v}\right) \leq \varepsilon_{n}$ if $\delta=\varepsilon_{n}^{2} /\left\{2 \sup _{|h| \leq \bar{r}_{n} C_{n}}\left|a^{\prime}(h)\right| \sup _{|h| \leq \bar{r}_{n} C_{n}}|\psi(h)| \bar{r}_{n}\right\}$. Therefore, density $p_{u}$ in $\mathscr{P}_{n}$ falls in a Hellinger ball of size $\varepsilon_{n}$, centered at $p_{v}$. There are at most $N(\delta)$ such balls, because each center $p_{v}$ is the density corresponding to the parameter $v$, which is the center of $B$, one of the at most $N(\delta)$ balls used to cover the restricted parameter space.

Therefore, the Hellinger covering number

$$
N\left(\varepsilon_{n}, \mathscr{P}_{n}\right) \leq N(\delta)
$$

$$
\begin{align*}
& \leq\left(\bar{r}_{n}+1\right) K_{n}^{\bar{r}_{n}}\left(1+2 \varepsilon_{n}^{-2} \sup _{|h| \leq \bar{r}_{n} C_{n}}\left|a^{\prime}(h)\right| \sup _{|h| \leq \bar{r}_{n} C_{n}}|\psi(h)| \bar{r}_{n} C_{n}\right)^{\bar{r}_{n}}  \tag{42}\\
& \leq\left(2 K_{n}^{2} D\left(\bar{r}_{n} C_{n}\right) / \varepsilon_{n}^{2}\right)^{\bar{r}_{n}},
\end{align*}
$$

if $0<\varepsilon_{n} \leq 1$ and $1 \leq \bar{r}_{n}<K_{n}$. Therefore, the conditions in Condition (O) guarantee that $\ln N\left(\varepsilon_{n}, \mathscr{P}_{n}\right) \prec n \varepsilon_{n}^{2}$ for all large enough $n$, proving condition (a).

Checking condition (b). For the $\mathcal{P}_{n}$ defined above, the prior on the complement $\pi\left(\mathscr{P}_{n}^{c}\right) \leq \pi\left[|\gamma|>\bar{r}_{n}\right]+\sum_{\gamma:|\gamma| \leq \bar{r}_{n}} \pi[\gamma] \pi\left(\bigcup_{j \in \gamma}\left[\left|\beta_{j}\right|>C_{n}\right] \mid \gamma\right)$, which is at most $\pi\left[|\gamma|>\bar{r}_{n}\right]+\max _{\gamma:|\gamma| \leq \bar{r}_{n}} \pi\left(\bigcup_{j \in \gamma}\left[\left|\beta_{j}\right|>C_{n}\right] \mid \gamma\right)$. This is, due to Condition (O), at most $\left(1+\bar{r}_{n}\right) e^{-4 n \varepsilon_{n}^{2}}=e^{\ln \left(1+\bar{r}_{n}\right)-4 n \varepsilon_{n}^{2}} \leq \exp \left(-4 n \varepsilon_{n}^{2} / 2\right)$ for all large enough $n$. Here we have used $1 \leq \bar{r}_{n}<K_{n}$, so that $\ln \left(1+\bar{r}_{n}\right) \leq \bar{r}_{n} \ln K_{n} \prec n \varepsilon_{n}^{2}$ due to Condition (O). This proves condition (b).

Proof of Theorem 1. We apply Theorem 4 with $\varepsilon_{n}$ replaced by $\varepsilon_{n}^{\prime}$, so that the Hellinger neighborhood will take a radius $4 \varepsilon_{n}^{\prime}$. This can be later rescaled to obtain the results in Theorem 1 concerning a radius $\varepsilon_{n}$, by setting $\varepsilon_{n}=4 \varepsilon_{n}^{\prime}$ or $\varepsilon_{n}^{\prime}=\varepsilon_{n} / 4$.

For Condition (O): with the prior in Section 3, the condition on $\pi\left[|\gamma|>\bar{r}_{n}\right]$ is trivially satisfied, since it is zero due to truncation. We will take $C_{n}=\sqrt{\tilde{B}_{n} n \varepsilon_{n}^{2}}$ so that the condition on $\bar{r}_{n} \ln D\left(\bar{r}_{n} C_{n}\right)$ is satisfied. The condition on $\pi\left[\left|\gamma_{j}\right|>C_{n} \mid \gamma\right]$ is checked by using Mills' ratio. It is at most $2 e^{-C_{n}^{2} /\left(2 \tilde{B}_{n}\right)} / \sqrt{2 \pi C_{n}^{2} / \tilde{B}_{n}}$, which is therefore less than $e^{-n \varepsilon_{n}^{2} / 4}=e^{-4 n\left(\varepsilon_{n}^{\prime}\right)^{2}}$ for all large enough $n$, as required by Condition (O). Here $\tilde{B}_{n}$ is an upper bound on the prior variance of $\beta_{j}$ under model $\gamma$ with $|\gamma| \leq \bar{r}_{n}$, and $n \varepsilon_{n}^{2} \succ 1$. All other conditions in Condition (O) are satisfied.

For Condition (N): Take the sequence of models $\gamma_{n}$ such that, for each $n$, $\gamma=\gamma_{n}$ reaches its infimum in $\Delta\left(r_{n}\right)=\inf _{\gamma:|\gamma|=r_{n}} \sum_{j: j \notin \gamma}\left|\beta_{j}^{*}\right|$. Then $\sum_{j \notin \gamma_{n}}\left|\beta_{j}^{*}\right|=$ $\Delta\left(r_{n}\right) \prec n \varepsilon_{n}^{2}$.

For the condition on the prior $\pi\left[\beta \in\left(\beta_{j}^{*} \pm \eta \varepsilon_{n}^{2} / r_{n}\right)_{j \in \gamma_{n}} \mid \gamma_{n}\right]$, use the normality of the prior and obtain the lower bound $\left|2 \pi V_{\gamma_{n}}\right|^{-1 / 2} e^{-0.5 \beta^{T} V_{\gamma_{n}}^{-1} \beta}\left(\eta \varepsilon_{n}^{2} / r_{n}\right)^{r_{n}}$ for some intermediate value $\beta$ achieving the infimum of the density over $\left(\beta_{j}^{*} \pm \eta \varepsilon_{n}^{2} / r_{n}\right)_{j \in \gamma_{n}}$.

Note that $\beta^{T} V_{\gamma_{n}}^{-1} \beta \leq\|\beta\|^{2} B\left(r_{n}\right) \leq\left(\sum_{j \in \gamma_{n}}\left|\beta_{j}\right|\right)^{2} B\left(r_{n}\right) \leq C_{1} B\left(r_{n}\right)$ for some constant $C_{1}>0$, since the eigenvalues of $V_{\gamma_{n}}^{-1}$ are at most $B\left(r_{n}\right)$ (for all large enough $n$ ), and the Euclidean norm $\|\beta\| \leq \sum_{j \in \gamma_{n}}\left|\beta_{j}\right| \leq \lim _{n \rightarrow \infty} \sum_{j=1}^{K_{n}}\left|\beta_{j}^{*}\right|+$ $r_{n} \eta \varepsilon_{n}^{2} / r_{n}$ is bounded. Note also that $\left|2 \pi V_{\gamma_{n}}\right|^{-1 / 2} \geq e^{-C_{2} r_{n}-C_{3} r_{n} \ln \bar{B}\left(r_{n}\right)}$ for some constant $C_{2}$ and some constant $C_{3}>0$, due to the eigenvalues of $V_{\gamma_{n}}$ being bounded above by $\bar{B}\left(r_{n}\right)$ (for all large enough $n$ ).

Therefore,

$$
\begin{align*}
\pi[\beta & \left.\in\left(\beta_{j}^{*} \pm \eta \varepsilon_{n}^{2} / r_{n}\right)_{j \in \gamma_{n}} \mid \gamma_{n}\right] \\
& \geq \exp \left\{-C_{2} r_{n}-C_{3} r_{n} \ln \bar{B}\left(r_{n}\right)-0.5 C_{1} B\left(r_{n}\right)-r_{n} \ln \left(r_{n} /\left(\eta \varepsilon_{n}^{2}\right)\right)\right\} \tag{43}
\end{align*}
$$

This will be greater in order than any $e^{-c n \varepsilon_{n}^{2}}(c>0)$, satisfying a requirement of Condition (N), since $r_{n}, r_{n} \ln \bar{B}\left(r_{n}\right)$ and $B\left(r_{n}\right)$ are all smaller than $n \varepsilon_{n}^{2}$ in order, and so are $r_{n} \ln r_{n} \leq \bar{r}_{n} \ln K_{n}$ and $r_{n} \ln \left(1 / \varepsilon_{n}^{2}\right) \leq \bar{r}_{n} \ln \left(1 / \varepsilon_{n}^{2}\right)$.

Now consider the condition on $\pi\left(\gamma_{n}\right)$. Note that the $\gamma_{n}$ chosen is such that $\left|\gamma_{n}\right|=r_{n}$, where $r_{n}\left(\leq \bar{r}_{n}\right)$ is the expected size of the model $\tilde{\gamma}=\tilde{\gamma}_{1}^{K_{n}}$ proposed by the prior before truncation. The prior specification of $\gamma_{1}^{K_{n}}$ (in Section 3) is i.i.d. binary with $\pi\left(\tilde{\gamma}_{j}=1\right)=r_{n} / K_{n}$. For the condition on $\pi\left(\gamma=\gamma_{n}\right)$ to hold, it suffices for us to show that $(*)$ for any $c>0, \pi\left(\tilde{\gamma}=\gamma_{n}\right)>e^{-c n \varepsilon_{n}^{2}}$ for all large enough $n$. This is because $\pi\left(\gamma=\gamma_{n}\right)$ cannot be smaller, since it is obtained by truncation of $\tilde{\gamma}$, and truncation increases the probability of all allowed configurations (note that $\left.\left|\gamma_{n}\right| \leq \bar{r}_{n}\right)$.

Now $\left|\gamma_{n}\right|=r_{n} \leq \bar{r}_{n}$ implies that there are $r_{n}$ out of $K_{n} \tilde{\gamma}_{j}$ 's equal to 1 , with the rest being 0 . The probability is therefore $\pi\left(\tilde{\gamma}=\gamma_{n}\right)=\left(r_{n} / K_{n}\right)^{r_{n}}(1-$ $\left.r_{n} / K_{n}\right)^{K_{n}-r_{n}}$. Since $r_{n} / K_{n} \prec 1$, we have $\ln \pi\left(\tilde{\gamma}=\gamma_{n}\right) \sim r_{n} \ln \left(r_{n} / K_{n}\right) \geq-r_{n} \ln K_{n}$ ( $r_{n} \geq 1$ ), where $r_{n} \ln K_{n} \prec n \varepsilon_{n}^{2}$. This leads to claim (*).

Proof of Theorem 3. It suffices for us to prove (**) result (ii) of Proposition 1 in a regression setup for normal dispersion models with $K_{n}$ Gaussian covariates. We will take $\varepsilon_{n} \sim n^{-(1-\xi) / 2}$ for some $\xi \in(0,1)$. Then result (i) of Theorem 3 can be obtained by a union bound over $K_{n}+1$ regressions, treating each of the $K_{n}+1 x_{j}$ 's in turn as the response $y$, and $x_{k \neq j}$ as $x_{1}^{K_{n}}$. Result (ii) of Theorem 3 can be obtained by using bounds of the form (24).

We prove $\left({ }^{* *}\right)$ by directly applying Proposition 1 and verifying conditions (a), (b) and (d) (for $t=1$ ). The details are omitted here and are included in a technical report (Jiang [14]) in order to save space.
9. Discussion. Bayesian variable selection (BVS) handles high-dimensional regression by using a suitable prior to propose lower-dimensional models which select a few explanatory variables out of the many $\left(K_{n}\right)$ candidates. For generalized linear models, we have shown that (see, e.g., Remark 1) a near finite-dimensional
convergence rate $\varepsilon_{n}$ can be obtained, even when the number of candidate variables $K_{n}$ grows as any high power $n^{\alpha}$ of the sample size $n$. Such a good rate $\varepsilon_{n}$ is derived assuming an exponentially decaying tail $\Delta\left(r_{n}\right)=\inf _{\gamma:|\gamma|=r_{n}} \sum_{j \neq \gamma}\left|\beta_{j}^{*}\right|$. This includes as a special case the situation when only a fixed and finite number of true regression coefficients ( $\beta_{j}^{*}$ 's) are nonzero. On the other hand, it also allows more realistic situations with many small $\left|\beta_{j}^{*}\right|^{\prime}$ s, none of which is exactly zero. The rates we obtain here are infinitesimally weaker than the finite-dimensional rate $n^{-1 / 2}$. We suspect that the exact rate $n^{-1 / 2}$ cannot be achieved in the setup that we consider, since the priors we use need to propose models of dimension $r_{n}$ increasing to infinity as $n$ increases (even though $r_{n} \ll n$ ). This is for the purpose of being able to approximate a true model to any precision. With such increasing model dimensions, we suspect that the exact $n^{-1 / 2}$-rate cannot be achieved in any way.

Although we have only considered in detail the situation with an exponentially decaying $\Delta(\cdot)$, the more general framework of, for example, Theorem 2 allows us to treat other situations of $\Delta(\cdot)$ as well. For example, when $\Delta(\cdot)$ follows an inverse power law, the convergence rate $\varepsilon_{n}$ can be somewhat slower. However, even in such situations, BVS can still exhibit some "resistance against overfitting" when $K_{n}$ is large. Not only can posterior consistency be still achieved when $\lim _{n \rightarrow \infty} \sum_{j=1}^{K_{n}}\left|\beta_{j}^{*}\right|<\infty$, but also the convergence rates will not be directly linked to the large dimension $K_{n}$-they will be related to the sizes of the $\left|\beta_{j}^{*}\right|$ 's instead.

An Associate Editor raised the interesting question whether the sparseness conditions for the true regression coefficients can be extended to a form of $\ell_{k}$-summability for $k>1$ (such as $\ell_{2}$ ). We do not have a general answer, except in an analytically-friendly special case as follows: The true model is $y \sim N\left(x^{T} \beta^{*}, 1\right)$ (it can be extended to allow a dispersion parameter), such that $E x x^{T}$ forms an identity matrix, or more generally, $E x x^{T}$ and its inverse both have bounded eigenvalues. The prior proposes fitted models of the form $y \sim N\left(x_{\gamma}^{T} \beta_{\gamma}, 1\right)$, according to the prior specification in Section 3. For this example, by a treatment parallel to the current paper, we can accommodate $\beta^{*}$ that is $\ell_{2}$-summable but not $\ell_{1}$-summable, such as $\beta^{*}=\left(j^{-1}\right)_{1}^{K_{n}}$, resulting in a possibly slower rate for posterior convergence. On the other hand, when $\beta^{*}=\left(j^{-1 / 2}\right)_{1}^{K_{n}}$, which is $\ell_{k}$-summable for $k>2$ but not $\ell_{2}$-summable, the current approach does not work. Roughly speaking, we would need to use very complicated fitted models of size $|\gamma| \sim K_{n}$ to approximate the true density, in order to obtain a nonzero prior probability over a small neighborhood of the true model. Then the complexity/entropy conditions [e.g., equation (10)], which imply $|\gamma| \ln K_{n} \prec n$, could not be satisfied for such fitted models of size $|\gamma| \sim K_{n}$ in a high-dimensional setting $K_{n}>n$.

Although the topic of our paper is Bayesian, it is noted that the use of $\ell_{1}$-type conditions here is related to some other work in the frequentist approach. Our paper is closer to Bühlmann [2] in the sense that both assume a true model satisfying some $\ell_{1}$ summability condition, while the fitted model (boosting for Bühlmann [2]
and BVS for the current paper) does not use an $\ell_{1}$ constraint or penalization. The fitted models in this paper are proposed according to a prior that uses i.i.d. binary distributions (with a small selection probability) when selecting the candidate variables. This may be regarded as a nondeterministic way of penalizing the $\ell_{0}$ norm of $\beta$ (or the number of nonzero regression coefficients) of the fitted models. On the other hand, in Greenshtein and Ritov [12], Greenshtein [11] and Meinshausen and Bühlmann [20], the fitted models (instead of the true models) are subject to an $\ell_{1}$ constraint or penalization. In the more general framework of persistence in Greenshtein and Ritov [12] and Greenshtein [11], the true models actually do not need to satisfy the $\ell_{1}$ summability condition.

The current paper focuses on fitting a density with Bayesian variable selection (BVS). A referee raised some interesting questions about the use of BVS when the main goal is selecting the variables. In some sense the current paper does prove that the method of BVS will provide "good" sets of variables, based on which good predictive performance, for example, in classification, can be achieved; see discussion in Section 5. The paper focused on the more realistic situation when there is no simple true model with many zero regression coefficients. All variables may have some effects, more or less. So the problem is not to select a "true" model (which would be the full model) but a "good" model (possibly much simpler than the full model) that achieves good performance for prediction, regression or density estimation. In this sense the paper does address variable selection and shows that BVS provides "good" sets of variables with high posterior probability. What will happen when there does exist a small true model, for example, when some regression coefficients are bounded away from zero, while the rest are exactly zero? We conjecture that, with high probability, BVS will select all the "relevant" variables with nonzero regression coefficients, but it may also include some "irrelevant" variables, with small regression coefficients proposed by the posterior. A truncation scheme similar to thresholding may be used to screen out the "irrelevant" variables, if necessary. However, we leave this as an open question, since such a scenario, being more idealized but still very interesting, is not within the main scope of the current paper.

Another future work may be to consider the (generalized) linear structure of the fitted models in a misspecified framework such as in Kleijn and van der Vaart [16], so that the true model may be nonlinear. On the other hand, one should note that nonlinearity may be treated even under the linear framework of the true model. This can be done by including higher order terms, interactions, regression spline terms with various knot-locations, and so on (see, e.g., Smith and Kohn [23]).

Acknowledgments. I wish to thank Professor Eitan Greenshtein for providing a related technical report. I also wish to thank the reviewers and the Editors for their insightful questions and comments.

## REFERENCES

[1] Bickel, P. J. and Levina, E. (2004). Some theory for Fisher's linear discriminant function, "naive Bayes," and some alternatives when there are many more variables than observations. Bernoulli 10 989-1010. MR2108040
[2] Bühlmann, P. (2006). Boosting for high-dimensional linear models. Ann. Statist. 34 559-583. MR2281878
[3] Clyde, M. and DeSimone-Sasinowska, H. (1997). Accounting for model uncertainty in Poisson regression models: Particulate matter and mortality in Birmingham, Alabama. ISDS Discussion Paper 97-06. Available at ftp.isds.duke.edu/pub/WorkingPapers/97-06. ps.
[4] Devroye, L., Györfi, L. and Lugosi, G. (1996). A Probabilistic Theory of Pattern Recognition. Springer, New York. MR1383093
[5] Dobra, A., Hans, C., Jones, B., Nevins, J. R., Yao, G. and West, M. (2004). Sparse graphical models for exploring gene expression data. J. Multivariate Anal. 90 196-212. MR2064941
[6] FAN, J. and Li, R. (2006). Statistical challenges with high dimensionality: Feature selection in knowledge discovery. International Congress of Mathematicians 3 595-622. European Math. Soc., Zürich. MR2275698
[7] GEorge, E. I. and McCulloch, R. E. (1997). Approaches for Bayesian variable selection. Statist. Sinica 7 339-373.
[8] Ghosal, S. (1997). Normal approximation to the posterior distribution for generalized linear models with many covariates. Math. Methods Statist. 6 332-348. MR1475901
[9] Ghosal, S. (1999). Asymptotic normality of posterior distributions in high-dimensional linear models. Bernoulli 5 315-331. MR1681701
[10] Ghosal, S., Ghosh, J. K. and van der Vaart, A. W. (2000). Convergence rates of posterior distributions. Ann. Statist. 28 500-531. MR1790007
[11] Greenshtein, E. (2006). Best subset selection, persistence in high-dimensional statistical learning and optimization under $\ell_{1}$ constraint. Ann. Statist. 34 2367-2386. MR2291503
[12] Greenshtein, E. and Ritov, Y. (2004). Persistence in high-dimensional linear predictor selection and the virtue of overparametrization. Bernoulli 10 971-988. MR2108039
[13] Hastie, T., Tibshirani, R. and Friedman, J. (2001). The Elements of Statistical Learning: Data Mining, Inference and Prediction. Springer, New York. MR1851606
[14] Jiang, W. (2005). Bayesian variable selection for high dimensional generalized linear models. Technical Report 05-02, Dept. Statistics, Northwestern Univ. Available at newton.stats. northwestern.edu/~jiang/tr/glmone2.tr.pdf.
[15] Jiang, W. (2006). On the consistency of Bayesian variable selection for high dimensional binary regression and classification. Neural Comput. 18 2762-2776. MR2256703
[16] Kleijn, B. J. K. and Van der Vaart, A. W. (2006). Misspecification in infinite-dimensional Bayesian statistics. Ann. Stat. 34 837-877. MR2283395
[17] Kohn, R., Smith, M. and Chan, D. (2001). Nonparametric regression using linear combinations of basis functions. Statist. Comput. 11 313-322. MR1863502
[18] Lee, H. K. H. (2000). Consistency of posterior distributions for neural networks. Neural Networks 13 629-642.
[19] Lee, K. E., Sha, N., Dougherty, E. R., Vannucci, M. and Mallick, B. K. (2003). Gene selection: A Bayesian variable selection approach. Bioinformatics 19 90-97.
[20] Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. Ann. Statist. 34 1436-1462. MR2278363
[21] Nott, D. J. and Leonte, D. (2004). Sampling schemes for Bayesian variable selection in generalized linear models. J. Comput. Graph. Statist. 13 362-382. MR2063990
[22] Sha, N., Vannucci, M., Tadesse, M. G., Brown, P. J., Dragoni, I., Davies, N., Roberts, T. C., Contestabile, A., Salmon, M., Buckley, C. and Falciani, F. (2004). Bayesian variable selection in multinomial probit models to identify molecular signatures of disease stage. Biometrics $\mathbf{6 0}$ 812-819. MR2089459
[23] Smith, M. and Kohn, R. (1996). Nonparametric regression using Bayesian variable selection. J. Econometrics 75 317-343.
[24] WANG, X. and George, E. I. (2004). A hierarchical Bayes approach to variable selection for generalized linear models. Techinical report. Available at www.cs.berkeley.edu/~russell/ classes/cs294/f05/papers/wang+george-2004.pdf.
[25] Wong, W. H. and Shen, X. (1995). Probability inequalities for likelihood ratios and convergence rates of sieve MLEs. Ann. Statist. 23 339-362. MR1332570
[26] Zhang, T. (2006). From $\varepsilon$-entropy to KL-entropy: Analysis of minimum information complexity density estimation. Ann. Statist. 34 2180-2210. MR2291497
[27] Zhou, X., LIU, K.-Y. and WONG, S. T. C. (2004). Cancer classification and prediction using logistic regression with Bayesian gene selection. J. Biomedical Informatics 37 249-259.

Department of Statistics
NORTHWESTERN UniVERSITY Evanston, IlLinois 60208 USA
E-MAIL: wjiang@ northwestern.edu


[^0]:    Received December 2005; revised October 2006.
    AMS 2000 subject classifications. Primary 62F99; secondary 62J12.
    Key words and phrases. Convergence rates, generalized linear models, high dimensional data, posterior distribution, prior distribution, sparsity, variable selection.

